

Copyright
by
Balaji Rengarajan
2009

The Dissertation Committee for Balaji Rengarajan
certifies that this is the approved version of the following dissertation:

**Self Organizing Networks: Building Traffic and
Environment Aware Wireless Systems**

Committee:

Gustavo de Veciana, Supervisor

Jeff Andrews

Constantine Caramanis

John Hasenbein

Sanjay Shakkottai

**Self Organizing Networks: Building Traffic and
Environment Aware Wireless Systems**

by

Balaji Rengarajan, B.E.; M.S.

DISSERTATION

Presented to the Faculty of the Graduate School of
The University of Texas at Austin
in Partial Fulfillment
of the Requirements
for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

August 2009

Acknowledgments

“I can no other answer make, but, thanks, and thanks.”

William Shakespeare

This dissertation would not exist if it were not for the guidance, encouragement, and support I received and continue to receive from faculty, friends, and family members. I take this opportunity to express my altogether inadequate thanks.

I’ve been especially fortunate to be advised by Prof. Gustavo de Veciana. I’ve always left meetings with him, brimming with ideas and optimism. He has provided me with scientific opportunities, and economic support. The countless lessons I have learned from him about research, academia, and life have helped me grow both professionally and as a person. Not only was Gustavo always readily available for me, as he so generously is for all of his students, but he always reviewed my writing more quickly than I could have hoped and offered many painstaking comments. In every sense, none of this work would have been possible without him. Thank you for being my role model and my mentor.

I would like extend a special thank-you to the members of my committee: Prof. Sanjay Shakkottai, whose class and infectious enthusiasm sparked my interest in the analysis of communication networks; Prof. Constantine

Caramanis who provided tremendous support and guidance as I struggled with understanding and applying the principles of convex optimization; Prof. Jeff Andrews who taught me all about wireless communication and whose incisive comments helped shape my dissertation; and Prof. John Hasenbein whose encouraging comments and insightful observation during my qualifying exam helped develop my research. The faculty at the University of Texas at Austin have provided me with a tremendous graduate education. They have taught me how to think scientifically and rigorously apply reasoning to difficult problems. I will always be in their debt.

Anyone who has ever been a student has probably run into the mounds of paperwork involved in a degree program. From the application for program acceptance to the application for graduation, my path has been smoothed by the efficient and ever cheerful staff of this department. I am grateful for having been able to rely on Melanie Gulick, Janet Preuss, and so many others. I also gratefully acknowledge the generous support from Intel and the National Science Foundation that enabled my work. I am also indebted to Intel, and Bell Labs for providing me internship opportunities that helped me understand how research in the real world works.

A penultimate thank-you goes to my parents, Vasantha and Rengaranjan. For always being pillars of support, and believing in me when I scarcely believed in myself, they deserve far more credit than I can ever give them. They have always stressed the importance of education, and all that is good in me comes from them. For instilling in me the spirit to ask questions and a

thirst for knowledge, I will be eternally grateful. Last and surely not least, I want to acknowledge my wonderful friends and colleagues. Many of them were willing guinea pigs as I rehearsed my talks multiple times and were the best walls to bounce ideas off of. In particular, I thank Jean, Kalpana, and Umesh for keeping me sane, and making my stay in Austin so much more enjoyable.

It would be impossible to thank individually, everyone who facilitated the successful completion of this project. To those I did not specifically name, I hope this acknowledgment will be accepted as a token of my sincere gratitude.

Self Organizing Networks: Building Traffic and Environment Aware Wireless Systems

Publication No. _____

Balaji Rengarajan, Ph.D.
The University of Texas at Austin, 2009

Supervisor: Gustavo de Veciana

This dissertation investigates how to optimize flow-level performance in interference dominated wireless networks serving dynamic traffic loads. The schemes presented in this dissertation adapt to long-term (hours) spatial load variations, and the main metrics of interest are the file transfer delay or average flow throughput and the mean power expended by the transmitters.

The first part presents a system level approach to interference management in an infrastructure based wireless network with full frequency reuse. The key idea is to use loose base station coordination that is tailored to the spatial load distribution and the propagation environment to exploit the diversity in a user population's sensitivity to interference. System architecture and abstractions to enable such coordination are developed for both the down-link and the uplink cases, which present differing interference characteristics. The basis for the approach is clustering and aggregation of traffic loads into

classes of users with similar interference sensitivities that enable coarse grained information exchange among base stations with greatly reduced communication overheads. The dissertation explores ways to model and optimize the system under dynamic traffic loads where users come and go resulting in interference induced performance coupling across base stations. Based on extensive system-level simulations, I demonstrate load-dependent reductions in file transfer delay ranging from 20-80% as compared to a simple baseline not unlike systems used in the field today, while simultaneously providing more uniform coverage. Average savings in user power consumption of up to 75% are achieved. Performance results under heterogeneous spatial loads illustrate the importance of being traffic and environment aware.

The second part studies the impact of policies to associate users with base stations/access points on flow-level performance in interference limited wireless networks. Most research in this area has used static interference models (i.e., neighboring base stations are always active) and resorted to intuitive objectives such as load balancing. In this dissertation, it is shown that this can be counter productive, and that asymmetries in load can lead to significantly better performance in the presence of dynamic interference which couples the transmission rates experienced by users at various base stations. A methodology that can be used to optimize the performance of a class of coupled systems is proposed, and applied to study the user association problem. It is demonstrated that by properly inducing load asymmetries, substantial performance gains can be achieved relative to a load balancing policy (e.g., 15 times

reduction in mean delay). A novel measurement based, interference-aware association policy is presented that infers the degree of interference induced coupling and adapts to it. Systematic simulations establish that both the optimized static and interference-sensitive, adaptive association policies substantially outperform various proposed dynamic policies and that these results are robust to changes in file size distributions, channel parameters, and spatial load distributions.

Table of Contents

Acknowledgments	iv
Abstract	vii
List of Figures	xiv
Introduction	1
Part I System-Level Coordination	9
Chapter 1. Overview of Proposed Approach and Related Work	10
1.1 Related Work	11
1.1.1 Centralized Approaches	12
1.1.2 Labeling-Based Approaches	13
1.1.3 Time Reuse Partitioning Approaches	15
1.1.4 Joint Coding Across Base Stations	16
1.1.5 Approaches Using Coordinated Power Control	16
1.1.6 A Major Shortcoming: Dynamic Traffic Loads	19
1.2 The Central Idea: Exploiting Population Diversity	20
1.3 Contributions	22
Chapter 2. The System Model and Traffic Abstractions	25
2.1 System Model	25
2.2 Simulation Model	26
2.3 Traffic Abstractions - User Classes	27

Chapter 3. The Downlink Case	32
3.1 System Abstractions	32
3.2 Optimizing the Decoupled Model	37
3.2.1 Static Scheduling	37
3.2.1.1 Estimating Class Capacities	38
3.2.1.2 Matching Capacity and Load	42
3.2.1.3 Delay Optimal Scheduling	45
3.2.2 Dynamic Inter-Class Scheduling	47
3.3 Optimizing the Coupled Model	49
3.4 Further Benefits of Coordination	55
3.5 The Importance of Being Traffic Aware	57
Chapter 4. The Uplink Case	62
4.1 System Abstractions	64
4.1.1 Determining user power levels under interference profiles	65
4.1.2 Estimating class rates	65
4.1.3 Optimizing the schedule	65
4.2 User Performance	67
Chapter 5. Final Observations	70
5.1 Cost of Base Station Coordination	70
5.2 Reducing Coordination Overheads	71
5.3 Integration with Physical Layer Methods and Future Research Avenues	74
Part II Optimizing User Association Policies	76
Chapter 6. Motivation and Related Work	77
6.1 Related Work	77
6.2 Is Load Balancing Always Optimal?	80
6.3 Contributions	82

Chapter 7. Analyzing a Coupled Queuing System	84
7.1 Prior Work	84
7.2 System Model and Notation	86
7.3 Performance Bounds	88
7.3.1 The Moments Based Approach	89
7.3.2 A Semidefinite Relaxation	92
7.4 Fidelity of the Bounds	94
7.5 Some Other Instances of Coupled Queuing Systems	97
7.5.1 A Generalized Processor Sharing System	98
7.5.2 A Multiple Access System	104
Chapter 8. Static User Association Policies	109
8.1 System Model	109
8.2 Simulation Model	111
8.3 Optimal Static Policies	112
8.4 Optimal Threshold Trends	121
8.5 Optimizing the Threshold	124
8.5.1 Determining Optimal Thresholds	125
8.6 Performance Comparison	127
8.6.1 Comparing Static Policies	127
8.6.2 Optimized Policy vs. Dynamic Strategies	127
8.6.3 Three Base Station Network	129
8.6.3.1 Weighted Signal Strengths	130
8.6.3.2 Pairwise Optimization	132
Chapter 9. The Interference-Sensitive Adaptive Policy	136
9.1 Measuring the impact of interference	137
9.2 Algorithm to determine the serving BS	140
9.2.1 Two base station scenario	140
9.2.2 Multi base station scenario	141
9.3 Performance Evaluation	143
9.3.1 Two base station scenario	143
9.3.2 Three (or More) Base Station Network:	146

9.3.3 Performance Sensitivity	147
9.4 Summing Up	150
Conclusion	151
Bibliography	155
Vita	169

List of Figures

1	Variation of interference levels within a cell	4
1.1	Examples of concurrent transmission scenarios with strong and weak interference.	21
1.2	Capacity to users within a cell: Diversity in sensitivity to interference from neighboring base stations.	22
2.1	Traffic abstractions to enable efficient coordination.	28
2.2	Building user classes.	30
3.1	Illustration of a joint transmission profile.	33
3.2	System abstractions for downlink coordination.	35
3.3	Comparing the different estimates for class capacity	43
3.4	Average file transfer delays under capacity maximizing static schedules.	45
3.5	Performance of capacity maximizing vs. delay optimal static, decoupled schedules with 2 classes per sector.	47
3.6	Average file transfer delays under delay minimizing static, decoupled schedules complemented by inter class scheduling with 2 classes per sector.	49
3.7	Average user throughput under delay minimizing static, decoupled schedules complemented by inter class scheduling with 2 classes per sector.	50
3.8	Average file transfer delays under schedules factoring inter-base station coupling, with 2 classes per sector.	54
3.9	Average user throughput under schedules factoring inter-base station coupling, with 2 classes per sector.	54
3.10	Average power consumed at the base stations.	55
3.11	Distribution of user-perceived delay	56
3.12	CDF of user delay	56
3.13	A clustered user population.	58
3.14	A scenario with non-homogeneous spatial load	60

4.1	Difference between uplink and downlink scenarios.	63
4.2	Uplink delay performance under delay minimizing schedules that account for inter-base station coupling.	67
4.3	Average throughput under schedules factoring inter-base station coupling.	68
4.4	Power-Delay tradeoffs on the uplink.	69
5.1	Performance comparison schemes with no coordination, with coordination, with interference cancellation and then coordination and interference cancellation.	75
6.1	The user association problem in a two base station scenario.	81
7.1	Service Rates in the two queue system.	95
7.2	Mean delay in the two queue system.	96
7.3	Three Queue system	97
7.4	Mean overall system queue length for a 2 Queue GPS system.	99
7.5	Mean queue length at one of the queues in the two queue GPS system.	100
7.6	Mean system queue length	101
7.7	Mean queue length at queue 1	102
7.8	Mean queue length at queue 2	103
7.9	Mean queue lengths resulting from the different objective functions.	103
7.10	Optimal weights for the different objective functions	104
7.11	Optimal probability of contention.	106
7.12	Mean delay comparison.	107
7.13	Mean throughput comparison.	107
8.1	A sub-optimal load allocation policy.	114
8.2	Example coupling construction for arrivals/departures based on realization of Z	117
8.3	Impact of load split threshold on delay performance.	122
8.4	Goodness of optimized thresholds.	126
8.5	Performance of the optimized policy vs. load balancing.	128
8.6	Mean Delay - Static vs. Greedy	129

8.7	Mean Delay - Static vs. Repacking	130
8.8	Spatial delay distribution - Static vs. Greedy system	131
8.9	Load division after projecting down to the two dimensional hyperplane	132
8.10	Three Base Station Network: Mean Delay Performance under the weighted signal strength policy	134
9.1	ISAP: Delay performance under heterogeneous spatial loads	145
9.2	ISAP: Delay performance in a three base station network with homogeneous load.	147
9.3	Sensitivity to channel model: Optimal load split thresholds	148
9.4	Sensitivity to channel model: Delay performance	148
9.5	Sensitivity to file size distribution: Delay performance	149

Introduction

Multiple network technologies are emerging to support high data-rate applications, making bandwidth a limited commodity. At the same time, the push for mobility has spurred the need for smaller devices that are more powerful, have more capabilities and longer battery lifetimes. Network deployments with increased numbers of base stations/access points can successfully deliver high capacity and energy efficiency. By decreasing the distance between users and their base stations, one can drastically increase capacity while reducing transmission energy requirements. Of course, this comes at a significant increase in infrastructure and management costs. There are also deleterious implications in terms of the operational regime of such networks. In particular, the proportion of users whose capacity is limited by interference from their neighbors grows. Also, as the number of base stations serving an area is increased, the coverage area and the number of users served by individual base stations decreases. This has the undesirable side effect of reducing the network's capability for statistical multiplexing and increases the 'burstiness' of the offered load. Thus we are faced with operating wireless systems in a highly dynamic, interference limited regime. In the presence of high temporal variability, spatial heterogeneity and interference, the goals of efficiency, high capacity and user performance require system operation that is adapted to the traffic and the environment.

A simple way to ease this problem is to carefully plan the deployment

of the wireless network. Consider the scenario shown in Fig. 1, with two neighboring base stations transmitting concurrently. The extent to which a user is affected by interference depends on the location of the user relative to the serving and interfering base stations, and the propagation characteristics. Users close to the serving base station typically see a strong received signal and weak interference, while users close to the cell boundary see a comparatively weak signal and stronger interference. If base stations are deployed such that the majority of the users are close to their serving base stations, the impact of inter-cell interference becomes less significant. However, dynamic populations of mobile terminals, which offer unpredictable spatial traffic loads make this difficult to accomplish. The interference problem is likely to be further exacerbated as traditional macrocellular systems are enhanced with repeaters for range enhancement and hierarchical cellular architectures are incrementally deployed with limited RF planning to meet changing usage patterns.

Interference is a phenomenon that is significantly different from noise. Noise can be combated simply by increasing the transmit power. Increasing transmit power across the board, however, also results in increased interference. Therefore, it is necessary to employ other techniques to tackle interference. Next generation wireless networks are expected to support a wide variety of data and multimedia services in addition to traditional voice traffic. These applications bring fundamentally different traffic characteristics and quality of service requirements as compared to traditional voice services. Policies in these networks must be designed to maximize spectral efficiency and sys-

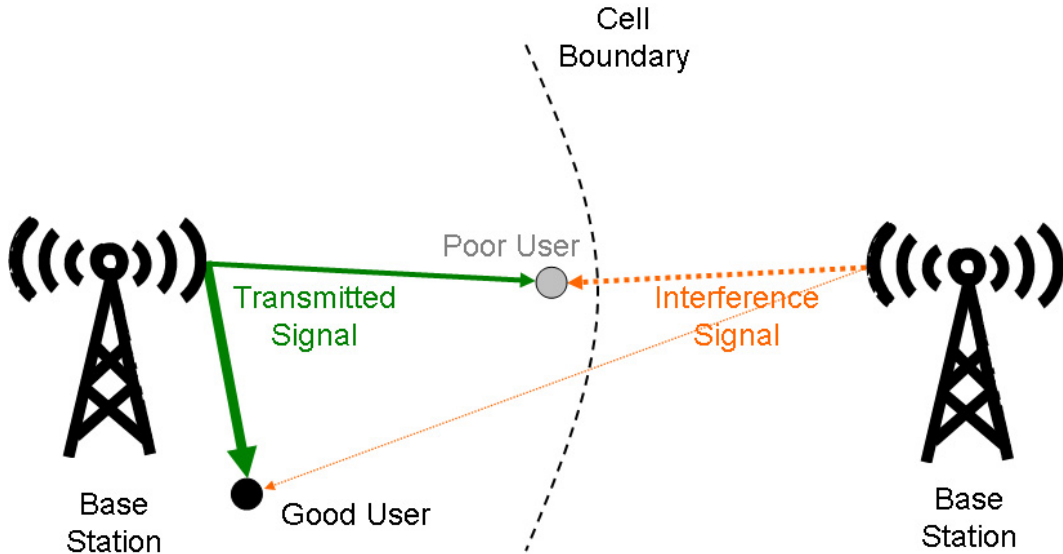


Figure 1: Variation of interference levels within a cell

tem capacity along with increasing the reliability of the airlink and enhancing coverage. Although next generation wireless systems such as OFDMA based WiMAX/3GPP(LTE) use time and frequency assignments to orthogonalize intra-cell transmissions i.e., among users in the same cell, inter-cell interference will continue to be an important factor impacting performance in such systems.

Traditional approaches for mitigating interference across base stations in an infrastructure based wireless network partition resources, e.g., frequency, so that concurrent transmissions can be realized with minimal interference. Such approaches are simple and do reduce the effective interference seen by users, thus enhancing the coverage area of a base station. However, this reduction in interference is achieved at the expense of significantly diminished

individual peak and overall system capacity. For example, with a reuse factor of $1/3$, the bandwidth available for transmission at a base station is reduced by a factor of three, with a similar reduction in the interference seen at the receivers. Yet, since the achievable capacity decays roughly linearly with the available bandwidth, but only grows logarithmically with SINR, this tradeoff is typically not favorable except at locations experiencing exceedingly high interference, and low SINR levels. Reusing the entire frequency spectrum in every cell can allow us to achieve very large network capacities, provided inter-cell interference is effectively managed. Even in the case of WLANs with frequency reuse, high densities of users in large scale networks could lead to high interference due to the limited number of orthogonal frequencies available under the present standards, making interference inevitable.

Adapting to the Spatial Load: Base stations in wireless networks do not serve a fixed set of users concentrated at a single location, but rather serve users that tend to be distributed geographically throughout the service area. Users at diverse locations typically see very different channel gains to the neighboring base stations in the network. Thus, users are impacted differently by interference depending on their spatial location, as exhibited in Fig. 1. The spatial distribution of the load in the network has a significant effect on the capacity of the network, and the performance experienced by users.

In addition to short-term, unpredictable variations in the load caused by individual user arrivals and departures, there are predictable long-term variations in the aggregate traffic load depending on the day-of-week, hour-of-

day, etc. A scheme that is adapted to these long-term traffic patterns, and not dependent on the instantaneous loads or short-term variations is potentially simpler than dynamic schemes which require knowledge of the instantaneous loads being served. Our research hypothesis is that one can reap substantial benefits through system-level optimization of wireless network by:

1. measuring the characteristics of the environment and traffic loads seen by base stations/access points.
2. based on the measurements, perform traffic-aware optimization of network functions across base stations.

As part of this dissertation, I will examine the impact of heterogeneous spatial traffic loads on network performance, and demonstrate that policies adapted to the spatial loads indeed outperform even conventional dynamic ones.

Impact of dynamic system loads: In a realistic scenario, data requests from users are generated at random times, and the users leave when their service requirements have been met. Such dynamic systems are, in general, hard to analyze and have not been studied as extensively as their static counterparts, i.e., serving a fixed set of backlogged users. The load dynamics translate to time varying interference that couples user capacities across base stations in a complex manner. For such coupled systems, stability is fairly difficult to establish, and performance is particularly hard to optimize. The performance that users perceive in such dynamic systems can be very different

from the performance predicted by a saturated model that assumes that transmitters are always on. The dynamics of the offered loads plays a substantial role when there is interference coupling among base stations, and schemes that are optimal in a static setting can be sub-optimal for the dynamic setting.

In the static setting, researchers have typically proposed schemes that attempt to achieve acceptable signal to interference plus noise ratios at all user locations or ones that maximize the static network capacity. In the dynamic setting, the network capacity is inextricably linked to the distribution of the spatial load being served. Additionally, these metrics do not adequately capture the flow-level performance experienced by best effort users, e.g., file transfers and web browsing. The policies developed as part of this dissertation directly optimize user perceived flow level performance. The main metrics of interest are the file transfer delay or average flow throughput and the mean power expended by the transmitters. A further goal is to develop a low complexity, system-level approach that substantially improves performance perceived by best effort users without requiring high channel measurement and estimation, communication, and computational overheads.

The focus of this work is on understanding the type of dynamic coupling one might expect among wireless nodes which are subject to dynamic loads, and interfere with each other. I examine two network functions that when optimized in a traffic aware manner can substantially improve performance in interference dominated wireless networks serving dynamic user populations:

1. system level coordination of base stations/access points
2. user association with base stations/access points.

There is a lot of interplay between the policies that tackle these problems, and ideally these policies would be jointly optimized. However, as we will see in the sequel, the impact of dynamic inter-cell interference can be difficult to characterize. So, for ease of analysis, this dissertation is divided into two parts, each focusing on one of the above problems.

Part I

System-Level Coordination

Chapter 1

Overview of Proposed Approach and Related Work

In the following chapters, a system-level approach for dynamic, environment and traffic aware network operation that can harness the temporal variability and spatial heterogeneity across users is developed. The approach is a relatively low-complexity one based on loose coordination among neighboring base stations to exploit diversity in users' sensitivity to interference from different base stations. Here, I consider primarily best effort traffic as against voice or other delay sensitive traffic. Best effort traffic can be distinguished from traditional voice traffic in that it is not as sensitive to delays, and users can be scheduled to take advantage of varying channel conditions. This fact has been used to opportunistically schedule [13, 61, 74] best effort users that see time varying channels, so as to improve performance. The scheme presented in this proposal is similar, but applies this principle in a scenario with multiple base stations to create and exploit favorable conditions for transmission to different sets of users. The metric we use to evaluate the proposed scheme is also different from the traditional outage based metrics used in the case of delay sensitive traffic, and focuses instead on the performance perceived by typical users in the system.

1.1 Related Work

Various *physical layer* approaches have been proposed in the literature that use advanced signal processing techniques to mitigate the effect of interference. The major drawback of most physical layer schemes is the high complexity resulting from the increased computational requirements at the receiver and transmitter as well as the need for instantaneous channel estimates. [27, 85, 87] are examples of receiver centric schemes, and [22, 49, 70] are some transmitter based techniques. A detailed overview of various physical layer approaches to combat cochannel interference can be found in [4, 41]. As an alternative, some researchers have also considered systems with distributed antennas instead of a single centrally located base station [25, 39, 71]. While such a topology has the advantage of reducing the distance between the access point and a typical user, it requires a significant infrastructure overhaul at a potentially high cost.

System level approaches to interference management typically attempt to increase the signal to interference ratio at the receivers through mechanisms that reduce the net interference power they see. An example of a simple method for co-channel interference management is the one proposed in [86] that employs dynamic power control based only on desired-path signal level measurements. The scheme proposed aims to improve the outage capacity of cellular systems by reducing the variance of the signal to interference ratio at the receiver while leaving the average unchanged. Several other schemes with varying degrees of complexity have been proposed to deal with cochannel

interference and some representative examples are detailed below.

1.1.1 Centralized Approaches

Centralized schemes to manage co-channel interference are explored in [28, 48]. The centralized scheduler proposed in [48] is assumed to be aware of all packet queue states, and path gains between every transmitter-receiver pair. The scheduler determines the subset of active base stations and their target users at every time slot, such that the signal to interference ratio at all the receivers is adequate to achieve satisfactory bit error performance. The scheduler simultaneously attempts to achieve the maximum possible throughput with good delay performance. However, the optimal centralized scheduler is found to be NP-hard. Different heuristic schedulers that achieve good throughput performance are presented, and their performance is evaluated through simulation.

The proposed scheme in [28] combines load balancing and interference avoidance. Users can potentially be served by any of the base stations under consideration, and further, a user's serving base station can vary from time slot to time slot. The base stations are restricted to either transmit to a user at full power, or to stay silent. Two variations of a centralized, coordinated scheduler are proposed. In the first one, the scheduler has the same information as the scheduler in [48], and chooses a subset of users to be served and the associated base stations every time slot. The other scheduler presented, called the two-tier scheduler, utilizes a centralized scheduler that is only aware of

long-term average signal strengths, but knows the instantaneous queue lengths. The centralized scheduler determines the active base station set, and each individual base station subsequently selects the particular user.

The disadvantage of the above schedulers is the large amount of information that has to be conveyed to the centralized scheduler. Even in the two-tier scheduler proposed in [28], the instantaneous queue lengths have to be communicated to the scheduler every time-slot, in addition to the long-term signal strengths for each user. The other schedulers impose even higher burdens on the backhaul. A TDMA based scheduler is proposed in [23, 24], where only the user with the best channel among all users in the coordinating cluster of cells is chosen for transmission. The gains from such a scheme in an asymptotic regime with a large number of users is quantified. This scheme is potentially suboptimal and still requires base stations to share information every time slot. Also, mechanisms to ensure fairness across users will be required in a practical system.

1.1.2 Labeling-Based Approaches

We use the term labeling-based approaches to describe schemes that utilize some system planning or engineering to guarantee that the transmissions of different base stations are coordinated. These schemes exploit the geometry induced by the typical hexagonal cell layout, and sectoring schemes used in large wireless networks. While these schemes typically assume that the various base stations are synchronized, they do not require a central controller.

An algorithm called the staggered resource allocation (SRA) method is presented in [32], that identifies the major sources of interference for each sector and schedules transmissions to avoid them. The interference sources are determined assuming idealized base station locations and a uniform propagation environment. Time is slotted into frames, and each frame is further divided into multiple subframes. The sectors are labeled assuming a reuse factor of six, and each sector progressively schedules transmissions in the subframes based on a predefined set of priorities that is designed to avoid the worst interferers. The maximum number of subframes that can be used by each sector is fixed, and is used to determine the degree of concurrent transmissions.

An enhanced version of the SRA algorithms is proposed in [58], that takes into account variations in the reception quality at the terminals. The terminals are classified based on the subset of interfering base stations that can be active while ensuring that the received signal to interference ratio is above the outage threshold. The subframes introduced by the SRA algorithm are further divided into mini-frames, and a fixed subset of base stations is allowed to transmit in a particular mini-frame. Transmissions to terminals are then scheduled based on the terminal's ability to tolerate cochannel interference. The proposed scheme does not take into account the possibility of using adaptive modulation and coding schemes based on the received Signal to interference ratios at the terminals. Another scheme similar to SRA is the one proposed in [76]. The scheme is adapted to work even with irregular cell layouts where the number of neighbors might vary across sectors.

1.1.3 Time Reuse Partitioning Approaches

Examples for time reuse partitioning approaches include [34, 68]. These methods deal with cochannel interference and the disparity in received signal to interference ratio across users by allowing multiple reuse patterns to coexist in the time domain. The method proposed in [68] uses two reuse patterns, one-cell and four-cell reuse. The frame is divided into two sections, one section consisting of shared slots that use one-cell reuse and the other section consisting of dedicated slots that use four-cell reuse. Users in each cell are classified into groups based on their average signal to interference ratio using a threshold. Users with average SIR below the threshold are preferentially scheduled in the dedicated slots to improve performance. A similar approach is used in [34], where the cell is divided into concentric region, each associated with a different reuse factor. The base stations estimate the relative load in each zone in the cell, and try to assign time periods to each zone such that all users achieve equal throughput.

[52, 59] propose approaches in the context of an OFDMA system where the available frequency spectrum is divided into orthogonal subchannels. [59] considers the case where only the dominant interferer to each user is taken into account. Subchannels are assigned to users in each cell through a centralized combinatorial optimization problem. [52] uses a centralized approach to specify different frequency reuse factors for the various subchannels as well as a heuristic with reduced complexity. The proposed scheme still requires careful frequency planning, and both the above schemes require an optimization

problem to be solved centrally every time a new user request arrives or leaves.

1.1.4 Joint Coding Across Base Stations

Highly fine-grained coordination among base stations has been examined, with base stations coherently coordinating transmissions. [42, 75] explore an extension of dirty paper coding to a cellular MIMO scenario where all interfering base stations jointly encode and coordinate transmissions. If the interference can be pre-calculated by the transmitters, such joint encoding can suppress interference and improve performance. However, the computational burden associated with the successive encoding and decoding is extremely high. [33, 45, 46, 50] study the gains that can be achieved through techniques that reduce the computational burden, such as zero-forcing transmissions where each users signal vectors are projected away from the others and MMSE precoding. [93] additionally examines an extension to TDMA, where one user is served at a time jointly by all base stations. Such schemes do improve performance, but at the cost of high computational complexity and greatly increased inter-base station communication requirements. The exact and instantaneous channel knowledge needed makes these schemes infeasible for practical systems.

1.1.5 Approaches Using Coordinated Power Control

These methods coordinate the transmit power used by neighboring base stations so as to create favorable conditions for the users in different parts of the cells. One such algorithm called quasi-static resource allocation is proposed

in [21] for sectored broadband networks. The central idea is for base stations to periodically turn off each of their sectors for a brief period of time, producing predictable variations in the channel conditions seen by various users. The users track the varying channel conditions and report their preferred transmission periods to the base stations. This information can be used by the base stations to schedule users such that all users in the cell experience acceptable performance levels. A distributed, measurement-based algorithm is presented that allows base stations to find a sequence in which sectors have to be turned off to ensure that all users perceive favorable channel conditions in some time period. Simulation results presented show moderate gains in the average bit error rates with this on-off strategy.

Another power-control based interference management scheme is proposed in [60]. The proposed scheme divides the available carriers into two sets, primary carriers that are used to transmit and receive packets at the maximum allowed power and the secondary carriers that are forced to use a reduced power. The users can be served using either the primary or the secondary set of carriers or using both sets. This is determined through an optimization that depends on the distribution of load in the cell, subject to some fairness constraint. Under certain conditions, it is found that the optimal strategy in a scenario with two cells is to serve near users using only the secondary carriers (low power) and far users using the primary carriers (high power). A similar approach to interference management that varies transmit power across time so as to improve performance is proposed in [88]. The proposed method, called

the cell-breathing scheme, varies the maximum allowable transmit power at a slow pace such that a cell transmits at high power when the other neighboring cells transmit at low power. The users track the varying channel conditions and this information is used by the base station to effectively schedule transmissions. An antenna-based version using the same concept has been explored in [84].

The concept of fractional frequency reuse (FFR) in the context of OFDMA systems has appeared in technical contributions submitted to cellular network standardization forums [82, 83]. Fractional frequency allows base stations to use different transmit powers in different frequency sub bands in order to mitigate interference. [66, 77] propose schemes in the context of voice-like traffic, where each user has a threshold SINR requirement that has to be met. The objective of the schemes is to minimize the overall transmit power used in order to reduce interference. A scheme to dynamically create FFR patterns for best effort traffic is proposed in [78]. While these algorithms respond to changes on fast time scales (seconds), they do not adapt to slow changes in the long-term load. The focus of the above schemes is to ensure that users perceive signal to interference ratios that maximize utility. However, this metric does not fully describe the flow-level performance experienced by best effort users, e.g., file transfers and web browsing.

1.1.6 A Major Shortcoming: Dynamic Traffic Loads

In a realistic scenario, data requests from users are generated at random times, and the users leave when their service requirements have been met. Such dynamic systems are, in general, hard to analyze and have not been studied as extensively as their static counterparts, i.e., serving a fixed set of backlogged users. The load dynamics translate to time varying interference that couples user capacities across base stations, and even the stability of such systems is difficult to verify, see [20]. The performance that users perceive in such dynamic systems can be very different from the performance predicted by the static model; e.g., the flow level performance of opportunistic scheduling was studied in a dynamic setting in [13], and it was demonstrated that schemes that are optimal in a static setting are sub-optimal for the dynamic setting. All the schemes discussed above concentrate on the static case, where the base stations serve a fixed set of users.

Potential capacity gains from inter-cell coordination in a dynamic setting were characterized in [10]. Capacity in this case, is analyzed in a dynamic setting where users arrive according to a random process, and leave when their service requirements have been met. In this scenario, capacity is defined in terms of the maximum amount of traffic that can be supported for a given spatial traffic pattern. It is shown that evaluating the network capacity is equivalent to finding the optimal static scheduling strategy that is independent of the network state. The capacity gain from inter-cell scheduling is quantified for two-cell networks, and for some very special cases of multi-cell networks,

e.g. perfectly symmetric networks. The scheme that is developed as part of this dissertation addresses the case with dynamic, spatially heterogeneous spatial traffic loads. The scheme is based on coarse-grained coordination between base stations, and I develop abstractions to keep the communication and computational overheads low. Moreover, this is the the first scheme to directly optimize user perceived performance (file transfer delays and throughput) in this setting.

1.2 The Central Idea: Exploiting Population Diversity

Base stations in wireless networks do not serve a fixed set of users concentrated at a single location, but rather serve users that tend to be distributed geographically throughout the service area. Users at diverse locations typically see very different channel gains to the neighboring base stations in the network. The key idea in this paper is to take advantage of this diversity in users' sensitivity to interference originating from the adjoining cells.

Fig. 1.1 illustrates this idea in the case of downlink transmissions in a two base station network. When both base stations are transmitting concurrently, they interfere with each other. The degree to which this is a problem depends on the locations of, and propagation characteristics to the mobile terminals of interest relative to both base stations. For example, as shown on the left in Fig. 1.1, if the mobile terminals are roughly at the midpoint between the two base stations, they see weaker received signals and stronger interference from the neighboring base station. In contrast, a node close to

its associated base station will see a strong received signal and weaker interference. Assuming each base station is likely to serve a large population of spatially distributed users, it would be advantageous to coordinate base stations so that they operate in favorable transmission scenarios such as the one on the right. In such a scenario, the base station on the right could use a low power to transmit to the nearby user while suffering minimal loss in capacity. This would considerably improve the conditions perceived by the user at the cell edge which can now be served at much higher capacities.




Figure 1.1: Examples of concurrent transmission scenarios with strong and weak interference.

Fig. 1.2 depicts an area served by three base stations, with each user (location) being served by the base station with the strongest signal. The rates at which users can be served when all the base stations are transmitting at maximum power is shown on the left, while the figures in the middle and on the right show the perceived capacity when one of the base stations is turned off. The diversity in the users' sensitivity to interference from different base stations can be easily seen. Users close to their serving base station are relatively unaffected by interference, while those at the cell edge see a significant increase in capacity when the interfering base station near them is turned off.

To effectively operate the system, this increase in capacity has to necessarily be balanced against the loss in capacity from turning a base station off. This scenario can be extended to one with a larger number of permissible power levels at which the base stations can operate. The challenge lies in coordinating transmissions and choosing transmit power levels that achieve good user performance while keeping complexity and overheads low. The novelty of my work lies in the development of new abstractions, a network architecture, and associated optimizations that make such coordination practical, and efficient.




Figure 1.2: Capacity to users within a cell: Diversity in sensitivity to interference from neighboring base stations.

1.3 Contributions

I propose a measurement-based coordination scheme that is tailored to the spatial load distribution served by the network, as well as the particular propagation environment. The proposed scheme only requires coarse grained information to be communicated among base stations over slow time scales, resulting in greatly reduced demands on the backhaul. I evaluate performance

in a dynamic setting where users come and go, and the main metrics of interest are the file transfer delay or average flow throughput and the mean power expended by the transmitters. My contributions are highlighted as follows:

First, I develop an approach to measure and classify a spatial population of users into a small number of user *classes* that capture average system loads, characteristics of the propagation environment, and interference sensitivities. These user classes are a critical abstraction towards reducing the complexity of the system-level optimization. To enable the optimization of *class-level* coordination schedules, one needs to properly represent the service rates that classes will see in a dynamic system. I propose an effective approximation for this which factors the intra-class variability in service capacity across users.

Second, I investigate the optimization of a coarse-grained coordination schedule for both the uplink and the downlink scenarios. I consider various scenarios from high to low loads. Key differences arise due to the degree of dynamic interference, i.e., neighboring base stations may not always be on, and the extent to which this impacts the optimized schedule's performance may vary. I propose and evaluate various approaches to incorporate such dynamics.

Third, through extensive analysis and simulation, I illustrate the significant gains that can be achieved in terms of delay performance, power consumption at the transmitter, and substantially enhanced spatially homogeneous service to users. I further demonstrate the impact that the spatial traffic distribution can have on user performance, illustrating the importance

of a scheme that is traffic aware.

Chapter 2

The System Model and Traffic Abstractions

2.1 System Model

Best-effort file transfers are considered on both the uplink and the downlink. User requests are assumed to arrive to the coverage area \mathcal{X} as a Poisson process with a location dependent intensity $\lambda(x), x \in \mathcal{X}$. For simplicity, each user is assumed to be associated with the base station that provides the strongest signal, e.g., the geographically closest base station in the absence of shadowing. User requests arrive at random, and leave the system when the associated data transfer is completed. In the downlink scenario, base stations are assumed to transmit at the specified power level when there are active associated users present, and turn off otherwise. A natural consequence of this assumption is that base stations interfere with transmissions in the neighboring cells only when serving associated users. Similarly, in the uplink scenario, no interference is generated by a cell with no active users.

In a wireless cellular network with full reuse, it is typically transmissions in the neighboring cells that generate most of the interference. In a small network, all the base stations could potentially be coordinated. Larger networks can be split into a number of independent coordinated clusters, such

that the cells/sectors whose performance is tightly coupled through mutual interference are grouped together. Let N denote the number of neighboring base stations/sectors being coordinated, indexed by $b = 1, \dots, N$. Let \bar{F}_{bk} , denote the mean size in bits of a file that is transferred on the downlink or the uplink. The long-term, average channel between base stations and users is assumed to be reciprocal. Let h_i^b denote the average channel gain between base station b and user i , and $\vec{h}_i = (h_i^b | b = 1, \dots, N)$ represent a collection of channel gain vectors.

2.2 Simulation Model

In the sequel I will describe different methods to coordinate base station transmissions, and use extensive uplink and downlink simulations to compare their performance. In the simulations, I consider three facing sectors in a hexagonal layout of base stations with cell radius 250m - see Fig. 2.2a. Users associate themselves to the geographically closest base station. A carrier frequency of 1GHz, and a bandwidth of 10MHz are assumed. The maximum transmit power is limited to 10W. The base stations are assumed to be able to transmit at three different power levels: 0, 5, and 10W. Additive white Gaussian noise with power -55dBm is assumed. A log distance path loss model [69] with path loss exponent 2 is considered. Shadowing, and fading are not considered in these results, but the addition of shadowing does not fundamentally change the characteristics of the proposed measurement driven scheme, see Sec. 2.3. File sizes are assumed to be log normally distributed, with mean

2MB. The data rate at which users are served is calculated based on the perceived SINR using Shannon's capacity with rates quantized to 0, 1, 2, 5, 10, 20, and 30Mbps. Users arrive according to a Poisson process, and except in Sec. 3.5 are assumed to be distributed uniformly within the simulated area. In Sec. 3.5, I explore the impact of non-homogeneous user populations, and the spatial traffic model is described in detail therein. In all simulations, mean user perceived delay is estimated within a relative error of 1%, at a confidence level of 95%.

2.3 Traffic Abstractions - User Classes

Exploiting the diversity in user populations to mitigate the effects of interference requires base stations to adapt not only to user distributions within their cells, but also to distributions in neighboring cells. Sharing information between base stations on a per user basis would result in extremely high communication costs. So, in this section I propose to use aggregates, see Fig. 2.1, that allow base stations to efficiently share information about the spatial loads and sensitivities to interference.

In addition to short term, unpredictable variations in the load caused by individual user arrivals and departures, there are predictable long-term variations in the aggregate traffic load depending on the day of week, hour-of-day, etc. [16, 18]. Consider monitoring a user population sharing a wireless system over a long period of time, say a few hours. I shall assume that during this period, the average rate of user requests arriving at any location x remains

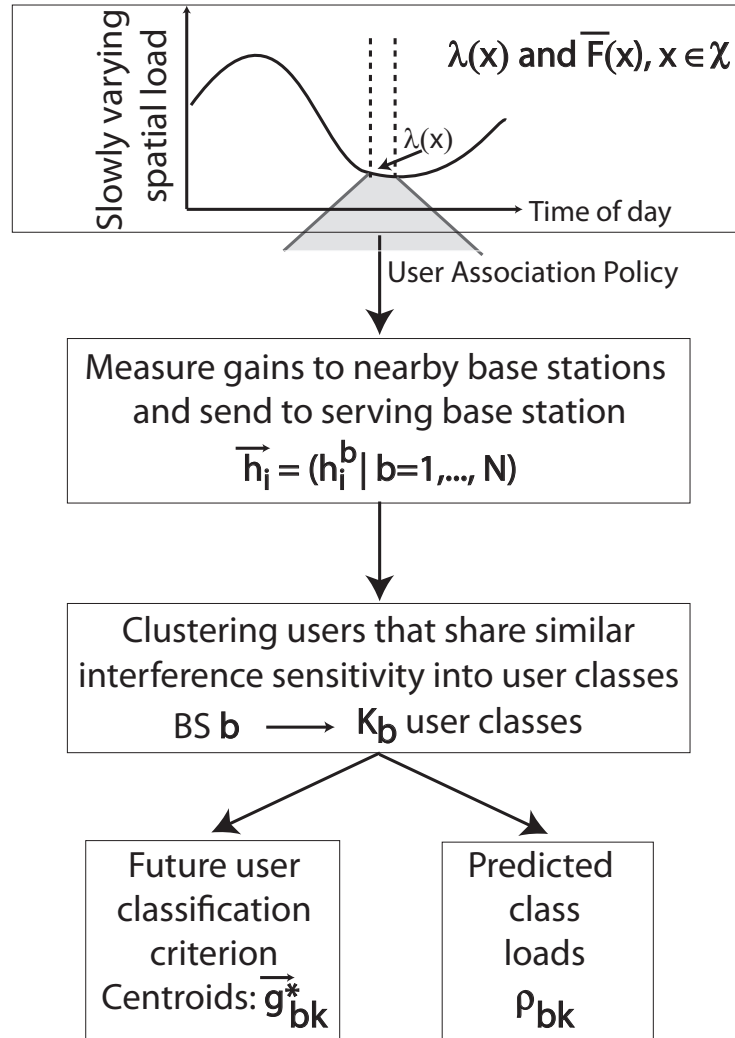



Figure 2.1: Traffic abstractions to enable efficient coordination.


constant, i.e., $\lambda(x)$ denotes the long-term rate of arrival of user requests at location x . Note that the traffic load might still be spatially heterogeneous. For each base station/sector b , I define K_b *user classes* that will abstract the key characteristics of the load distribution and the propagation environment. They enable base stations to measure, aggregate, and share coarse grained information about the traffic loads they support. They also drive the system-level optimization methodology described in the sequel.

User classes and class loads aggregate users (locations) that share similar sensitivity to interference from neighboring base stations. A simple way to capture these environmental conditions is to measure the average channel gains between users and neighboring base stations – this is already done in practice to facilitate handoffs. Users feedback the measured channel gain vectors \vec{h}_i to their serving base stations. Fig. 2.2a depicts the measurements made by each user when coordinating three facing sectors in a hexagonal layout of base stations.

Users sharing similar gain vectors, \vec{h}_i , have similar susceptibility to interference from neighboring base stations on the downlink, and cause similar levels of interference at the neighboring base stations in the uplink case. Yet, in an interference limited regime, Shannon’s capacity formula suggests that users transmission rates vary as the logarithm of the ratio of the received signal power to interference. Thus, for each user measurement, I define a logarithmically distorted gain vector $\vec{g}_i = (g_i^b | b = 1, \dots, N)$, where $g_i^b = \log(h_i^b)$. A k -means clustering algorithm [3, 43] is used to cluster measured log-gain vectors



(a) An example scenario for coordination.



(b) Example class definitions.

Figure 2.2: Building user classes.

into a fixed number of user classes. Specifically, the algorithm partitions users associated with base station b into K_b clusters with centroids \vec{g}_{bk}^* , $k = 1, \dots, K_b$, such that the mean Euclidean distance between the log-gain vectors and the centroids is minimized. Given a clustering, and the resulting centroid vectors, future users can be classified based on which centroid its log-distorted gain vector is closest to. In the sequel, I address the question of how many classes are used per base station, and the associated tradeoffs.

Fig. 2.2b exhibits a clustering for a sector in our example scenario where three neighboring base stations are to be coordinated. The points in the figure represent individual users, while the sets reflect their division into classes. The users near the serving base station are minimally impacted by interference, leading to radially symmetric classes that are influenced mainly by the path loss to the serving base station. Interference plays a significant part in trans-

missions involving users at the cell edge, and the asymmetric classes reflect the resulting discrimination based on which neighbor has the most impact. Note that in practice, due to shadowing and real environment obstructions, user classes will not result in the clean spatial partition exhibited in this example. In fact, they would instead reflect the character of the environment as well as the typical locations where the user population dwells.

With classes defined, estimating the average loads for each class under a given spatial traffic load is a simple task. Arrivals to class $k = 1, \dots, K^b$ associated with base station/sector b are thus Poisson, with rate denoted by λ_{bk} . Define $\rho_{bk} = \lambda_{bk} \bar{F}_{bk}$ to be the mean traffic (bits per second) arriving at class k in base station b . Let $\vec{\rho} = (\rho_{bk} : b = 1, \dots, N, k = 1, \dots, K^b)$ denote the expected offered load vector. The classes may have different offered loads, capturing in part the spatial distribution of traffic supported by the system. The expected offered load vector can thus be exchanged between base stations infrequently (on the order of hours), drastically reducing communication overheads.

Chapter 3

The Downlink Case

3.1 System Abstractions

For simplicity, I will initially focus on the downlink scenario. A *joint transmission profile* represents one of the various modes in which the network can be operated. As illustrated in Fig. 3.1, it specifies a power profile, i.e., the transmit power level for each base station, and the associated user classes to be jointly served. The base stations are assumed to be able to transmit at one of P discrete power levels, including 0, corresponding to no transmission. The N -dimensional column vectors \vec{p}^i and \vec{c}^j specify the power levels and classes to be served by the base stations under power profile i and class combination j . The b^{th} component of these vectors, p_b^i and c_b^j , specify the transmit power to be used by base station b and the class to be served. The number of different power profiles is denoted by $U = P^N$, the number of class combinations by $V = \prod_{b=1}^N K_b$, and thus the number of joint transmission profiles is $L = UV$. Let $\mathcal{P} := \{\vec{p}^1, \dots, \vec{p}^U\}$ and $\mathcal{C} := \{\vec{c}^1, \dots, \vec{c}^V\}$ denote the sets of admissible joint power profiles and class combinations respectively for the N base stations, and \mathcal{L} the set of joint transmission profiles. Thus, each joint transmission profile l where $l = 1, \dots, L$ is two vectors: $\vec{p}(l) \in \mathcal{P}$ and $\vec{c}(l) \in \mathcal{C}$.




Figure 3.1: Illustration of a joint transmission profile.

A joint transmission schedule corresponds to the fractions of time $\vec{\alpha} = (\alpha_l : l = 1, \dots, L)$ for which the network operates in each transmission profile. The base stations are assumed to be synchronized, as is the case in modern broadband systems like GSM and WiMAX, and they can (quickly) cycle through these profiles. Note that the resource that is subdivided for the purpose of coordination could also be frequency, or even a combination of time and frequency in an OFDMA-like system. In general, this schedule will be picked to optimize a chosen performance measure, $f(\vec{\alpha})$, through an optimization of the form:

Problem 3.1.1. *A generic optimization problem to determine a coordination schedule:*

$$\min_{\vec{\alpha}} f(\vec{\alpha})$$

such that

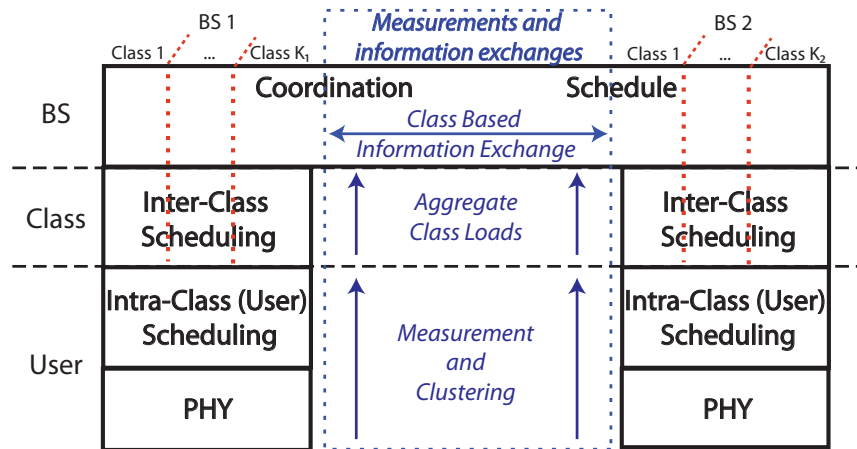
$$\rho_{bk} \leq R_{bk}(\vec{\alpha}), \forall b, k, \quad (3.1)$$

$$\sum_{l=1}^L \alpha_l \leq 1 \quad \text{and} \quad \alpha_l \geq 0, l = 1, \dots, L. \quad (3.2)$$

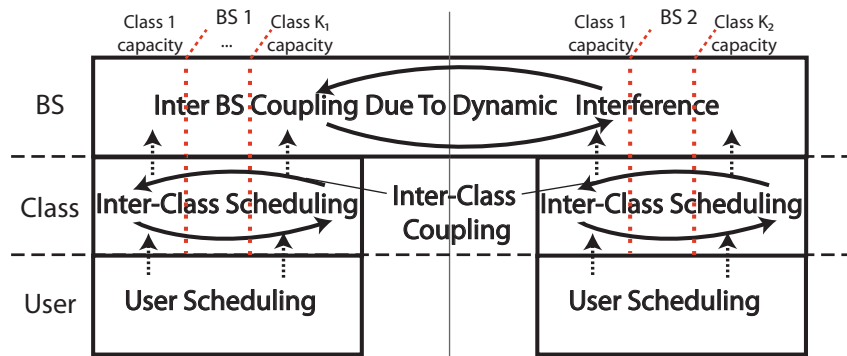
Here, $R_{bk}(\vec{\alpha})$ denotes the capacity allocated to class k at base station b under schedule $\vec{\alpha}$. Eq. (3.1) constrains the rate allocation across classes to be one that stabilizes the network. Eqs. (3.2) ensure that the coordination schedule is valid.

Fig. 3.2a exhibits the overall system architecture used for coordination. Each user reports signal strength measurements from neighboring base stations to its serving base station. Each base station uses this information to aggregate users into classes that capture the spatial load being served, and the nature of the propagation environment. Coarse grained information about traffic loads and achievable transmission rates are exchanged between base stations at the level of user classes. The base stations can then determine the optimized coordination schedule (common to all coordinating base stations), i.e., the fraction of time each transmission profile is used. The user classes are the key abstraction that allows such cross-layer, cross-base station optimizations to be carried out while keeping communication and computational overheads manageable.

Note that the transmission profiles are *not* a specification of which user to serve, only a restriction on the transmit power to be used at the base station and a ‘recommended’ class that might be beneficially served. Base stations can



(a) System architecture for coordination.



(b) Modeling performance coupling.

Figure 3.2: System abstractions for downlink coordination.

independently devise complimentary dynamic *inter-class* scheduling policies that serve classes other than the recommended one. Since the choice of class does not affect the interference levels observed at the neighboring cells, such

inter-class scheduling does not violate the coordination schedule. Further, base stations can use any *intra-class* scheduling policy to serve users within the selected class(es). In this paper, I assume that base stations use processor sharing (or an approximation thereof) to serve the active users in the chosen class(es). Fig. 3.2a summarizes the relationship between the various elements of the architecture.

In order to solve Problem 3.1.1, accurate estimates for $R_{bk}(\vec{\alpha})$ are needed. However, the dynamics of the system make this a difficult task, see Fig. 3.2b. User performance is coupled across base stations as the capacity to users is impacted significantly by the state (transmitting or not) of neighboring base stations. A further degree of coupling, intra-base station coupling, can be introduced depending on how individual users and classes are served within each base station. If inter-class scheduling depends, for instance, on the instantaneous loads in the classes, the performance of the different classes at a base station will be coupled together. The choice of user and inter-class scheduling policy also affect the activity level of the base stations, thus impacting neighboring base stations through interference driven coupling.

Determining the exact capacity allocated by a schedule to each class when the activity levels and performance of neighboring base stations are coupled corresponds to analyzing a set of spatially coupled (through interference) queues. Systems of coupled queues have been analyzed in the past [14, 15, 31, 36], but the problem is extremely difficult and closed form expressions are only available for simple scenarios with only two queues. So, for

simplicity, I assume in Sec. 3.2 that the performance of the various base stations are decoupled by assuming base stations are always on; i.e., the performance (interference) seen by users does not depend on the traffic at other base stations.

As a further approximation, I study policies under which a base station is restricted to serving only the class specified by the transmission profile. If the chosen class has no active users, the base station does not opt to serve another class. Such a policy is denoted *static scheduling*. Thus, there is no inter base station or inter class coupling. Subsequently, I evaluate the performance of a policy that allows base stations to share the excess capacity from empty scheduled classes among other associated users and thus introduces coupling among classes. In Sec. 3.3, I will drop the assumption of decoupled base stations and present approximations for optimizing the coupled systems.

3.2 Optimizing the Decoupled Model

3.2.1 Static Scheduling

As the number of user classes is increased, the fidelity of the gathered information increases. However, communication overheads, and the computational complexity associated with the proposed coordination scheme also grow. Problem 3.1.1, for example, has a number of constraints and decision variables which respectively grow linearly and polynomially (of degree N) in the number of classes. Therefore, it is advantageous to use a relatively small number of classes. However, in this case, there may be large disparities in transmission

rates of users in the same class. In order to optimize the schedule, we first need to develop good estimates for the class capacities, $R_{bk}(\vec{\alpha})$ which themselves depend on the schedule $\vec{\alpha}$. As will be seen in the following section, this is not a simple matter even for a decoupled model, yet good approximations that make the optimization problem convex can be found to make this tractable.

3.2.1.1 Estimating Class Capacities

Let the random variable I denote a randomly selected user from the system's load distribution, i.e., the distribution of user requests; thus $I = i$ corresponds to a location, and assume user i stays there until its request is completed. Let $b(i)$, and $k(i)$ be user i 's base station and class respectively. Finally, let R_i^l denote the peak rate at which user i can be served under profile l , assuming all base stations are active. Note that R_i^l is zero, if a class other than $k(i)$ is served by base station $b(i)$ under profile l .

Proposition 3.2.1. *Consider the downlink queue associated with class k at base station b . It sees an offered load of ρ_{bk} bits/sec., and a time varying capacity that depends on $\vec{\alpha}$. Suppose the rate at which base stations switch among profiles is fast compared to the time scale of the user dynamics, and the base station uses processor sharing to serve users in each class, then the queue is stable if $u_{bk} = \frac{\rho_{bk}}{R_{bk}^H(\vec{\alpha})} < 1$, where*

$$R_{bk}^H(\vec{\alpha}) = \frac{1}{\mathbf{E} \left[\frac{1}{\sum_{l=1}^L \alpha_l R_l^l} \mid b(I) = b, k(I) = k \right]}. \quad (3.3)$$

Further, when the queue is stable, the mean number of active users in the class

is given by $\frac{u_{bk}}{1-u_{bk}}$.

Proof. If the rate at which base stations switch between the different transmission profiles is infinitely fast, the variations in rate perceived by users become negligible, and the system corresponds to a processor sharing queue operating in a ‘fluid’ regime similar to the approximation used in [11]. In this regime, a typical user I is served at the average transmission rate given by $\sum_{l=1}^L \alpha_l R_I^l$ if it is the only active user in the class. In this case, the time to serve user I is $\frac{\bar{F}_{bk}}{\sum_{l=1}^L \alpha_l R_I^l}$. The mean time to serve a user in the class is given by $\mathbf{E} \left[\frac{\bar{F}_{bk}}{\sum_{l=1}^L \alpha_l R_I^l} \right] = \frac{\bar{F}_{bk}}{R_{bk}^H(\vec{\alpha})}$. The total normalized load offered by the class is then given by $u_{bk} = \frac{\rho_{bk}}{R_{bk}^H(\vec{\alpha})}$. The fact that this processor sharing queue is stable when $u_{bk} < 1$ follows from the results in [11, 13], and the mean queue length of the system can be computed to be $\frac{u_{bk}}{1-u_{bk}}$ using the expression for the queue length distribution from [11]. \square

Note that $R_{bk}^H(\vec{\alpha})$ is the harmonic mean of the average transmission rates seen by the different users in class k associated with base station b . Henceforth, I shall refer to this as the capacity allocated to the class under schedule $\vec{\alpha}$. Unfortunately, estimating this for each $\vec{\alpha}$ requires knowledge (estimates) of the complete spatial distribution of users versus simple descriptive statistics, e.g., means and variances, and thus increased communication and computational overheads.

The arithmetic and geometric mean of the average transmission rate perceived by users are alternative estimates for the class capacity. The arith-

metric mean approximation is given by:

$$\begin{aligned} R_{bk}^A(\vec{\alpha}) &= \mathbf{E} \left[\sum_{l=1}^L \alpha_l R_I^l \mid b(I) = b, k(I) = k \right] \\ &= \sum_{l=1}^L \alpha_l \mathbf{E}[R_I^l \mid b(I) = b, k(I) = k]. \end{aligned} \quad (3.4)$$

The geometric mean approximation for class capacity is given by:

$$R_{bk}^G(\vec{\alpha}) = \exp(E[\log(\sum_{l=1}^L \alpha_l R_I^l) \mid b(I) = b, k(I) = k]).$$

Note that the arithmetic mean is simple to compute: it depends only on the mean rates observed by users in the class under each profile, and is linear in $\vec{\alpha}$. However, it can be shown that $R_{bk}^H(\vec{\alpha}) \leq R_{bk}^G(\vec{\alpha}) \leq R_{bk}^A(\vec{\alpha})$, whence the geometric mean is the better estimate for the harmonic mean [38]. Unfortunately, the geometric mean is also burdensome to compute, making it unsuitable.

An approximation for the geometric mean based on moments was derived in [91], and empirical studies presented in [44] show that the approximation yields accurate results. I propose using this approximation, truncated to the first and second moments, to effectively capture intra-class diversity in transmission rates. Let Σ_{bk} be the covariance matrix of the transmission rates to the users in class k in base station b , $\sigma_{bk}(l, m) = \mathbf{Cov}[R_I^l, R_I^m \mid b(I) = b, k(I) = k]$. The rate allocated to class k in base station b is approximated as

$$\begin{aligned} R_{bk}^G(\vec{\alpha}) &\approx R_{bk}^A(\vec{\alpha}) - \frac{\mathbf{Var} \left[\sum_{l=1}^L \alpha_l R_I^l \mid b(I) = b, k(I) = k \right]}{2R_{bk}^A(\vec{\alpha})} \\ &= R_{bk}^A(\vec{\alpha}) - \frac{\vec{\alpha}^T \Sigma_{bk} \vec{\alpha}}{2R_{bk}^A(\vec{\alpha})}. \end{aligned} \quad (3.5)$$

Thus, the capacity allocated to all classes can be estimated with the coordinating base stations exchanging only the class means, and covariances of the transmission rates under the different profiles.

However, the estimate in Eq. (3.5) does not lead to constraint (3.1) being a provably convex function of $\vec{\alpha}$. The following approximation to Eq. (3.5) is used to model the allocated rates:

$$R_{bk}^{GA}(\vec{\alpha}) = R_{bk}^A(\vec{\alpha}) - \frac{\vec{\alpha}^T \Sigma_{bk} \vec{\alpha}}{c_{bk}}. \quad (3.6)$$

Here, $\vec{c} = (c_{bk}, b = 1, \dots, N, k = 1, \dots, K^b)$ is a positive vector that is appropriately chosen, to yield a good estimate for the class capacity.

Fact 3.2.1. $\left(R_{bk}^A(\vec{\alpha}) - \frac{\vec{\alpha}^T \Sigma_{bk} \vec{\alpha}}{c_{bk}}\right)^{-1}$ is a convex function of $\vec{\alpha}$, when it is positive, and \vec{c} is any positive vector.

Proof. $\frac{\vec{\alpha}^T \Sigma_{bk} \vec{\alpha}}{c_{bk}}$ is convex in $\vec{\alpha}$, since the covariance matrix and thus the Hessian is positive semidefinite. Also, $-R_{bk}^A(\vec{\alpha})$ is a linear function of $\vec{\alpha}$. Thus, $-R_{bk}^A(\vec{\alpha}) + \frac{\vec{\alpha}^T \Sigma_{bk} \vec{\alpha}}{c_{bk}}$ is also convex in $\vec{\alpha}$. This implies that $R_{bk}^A(\vec{\alpha}) - \frac{\vec{\alpha}^T \Sigma_{bk} \vec{\alpha}}{c_{bk}}$ is a positive concave function. Since the reciprocal of a positive, concave function is convex, $\left(R_{bk}^A(\vec{\alpha}) - \frac{\vec{\alpha}^T \Sigma_{bk} \vec{\alpha}}{c_{bk}}\right)^{-1}$ is a convex function of $\vec{\alpha}$. \square

I examine the actual achieved class capacities, and compare it to the estimates developed above in an example scenario with three sectors and users classified into two classes per sector using the method described in Sec. 2.3. Fig. 3.3 exhibits the class capacities for a fixed transmission schedule. The classification process results in classes of uneven sizes, and the classes with

higher load in Fig. 3.3 correspond to larger fractions of the uniformly distributed users with larger intra-class variance in the rates observed by users. The schedule used allocates a larger share of the capacity to these larger classes. Both the arithmetic and geometric mean approximations are optimistic in estimating the capacity allocated to classes, but the geometric mean is much more accurate as it takes into account the variability within a class. As can be seen from the figure, this larger variance results in the arithmetic mean being too optimistic for the larger classes, and overestimates the capacity allocated to the classes by up to 20% compared to the geometric mean estimate. The simulation results also indicate that the geometric mean approximation yields considerably better estimates for the class capacities, compared to the arithmetic mean. In the case of a fast fading environment, if the base station uses an opportunistic policy for user scheduling, the estimates should additionally capture the effective class capacities allocated by the scheduler as well as the channel state dependent, time-varying rates at which users will be served. The derivation of such capacity estimates which will necessarily depend on the characteristics of the chosen user level scheduler is a topic for future study. Next, two different strategies for optimizing user performance are discussed.

3.2.1.2 Matching Capacity and Load

The first schedule optimization approach that is considered to determine the joint transmission schedule is as follows:

Problem 3.2.1. *Determine a static, capacity maximizing, decoupled schedule*

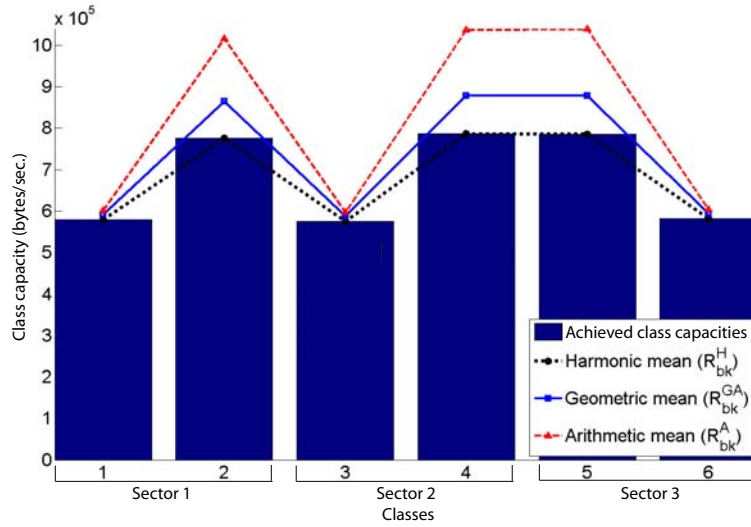


Figure 3.3: Comparing the different estimates for class capacity

based on:

$$\min_{\vec{\alpha}} \left\{ \sum_{l=1}^L \alpha_l \mid \rho_{bk} \leq R_{bk}(\vec{\alpha}), \forall b, k, ; \alpha_l \geq 0, l = 1, \dots, L. \right\}$$

The optimal schedule maximizes the fraction of time that the system is idle, which is a natural starting point. The optimal transmission schedule $\vec{\alpha}^*$ associated with Problem 3.2.1 assigns capacity to each class in proportion to the offered load. This formulation is similar to the idealized case considered in [10], and the optimal schedule stabilizes the network, if possible, for any load distribution proportional to $\vec{\rho}$ when $R_{bk}(\vec{\alpha})$ is exact, i.e., $R_{bk}(\vec{\alpha}) = R_{bk}^H(\vec{\alpha})$.

The geometric approximation from Eq. (3.6) is used to estimate class capacities. To determine the constants, c_{bk} , optimization Problem 3.2.1 is first solved with $R_{bk}(\vec{\alpha}) = R_{bk}^A(\vec{\alpha})$ to find $\vec{\alpha}^{A*}$. Let c_{bk} be the arithmetic mean

approximation of the rate allocated using schedule $\vec{\alpha}^{A*}$, $c_{bk} = R_{bk}^A(\vec{\alpha}^{A*})$. Then, problem 3.2.1 is re-solved with the geometric mean approximation.

The graph in Fig. 3.4 shows the average downlink file transfer delays vs. offered load obtained under three schemes: uncoordinated transmissions at the maximum power, and two static approximations with two and three user classes per base station. At higher loads, coordination performs extremely well, improving delay performance over the scheme with no coordination by over 80%. However, this is not uniformly the case, and at very low loads, the coordination scheme increases mean delays by around 50% compared to the non-coordinated scheme. Under low loads, coordinating across base stations to mitigate interference is less of a concern because the probability that neighboring base stations are simultaneously transmitting is low. Therefore, one might as well allow base stations to transmit at higher power without coordination. Also, since a static schedule is being used, the probability that there are no active users in the class scheduled at a base station is high at low loads. This leads to the base station unnecessarily wasting time while users wait their turn to get served. This is also the reason for the coordination scheme with only two classes per sector outperforming the scheme with three classes until the offered load is high enough. A larger number of classes results in base stations wasting more time when using a static schedule, as the scope for statistical multiplexing is further reduced. Splitting the load and the resources into independent small chunks results in reduced capacity for sharing, and incurs a statistical multiplexing loss. At low loads, the gains from reduced interference

levels resulting from careful coordination across base stations are not sufficient to compensate for this statistical multiplexing loss.

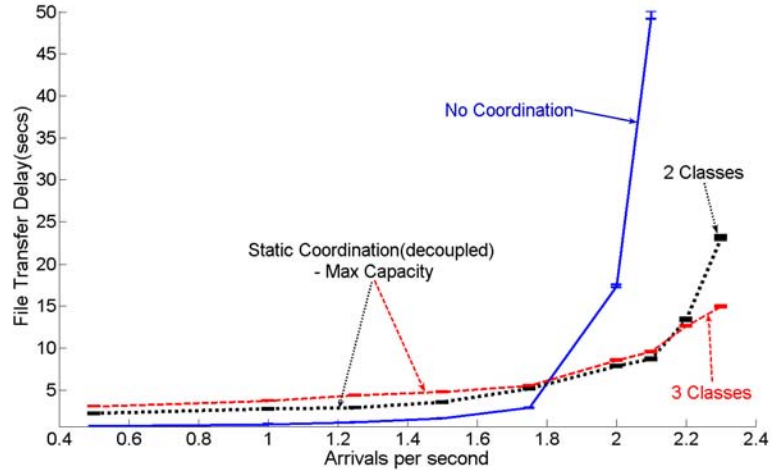


Figure 3.4: Average file transfer delays under capacity maximizing static schedules.

3.2.1.3 Delay Optimal Scheduling

When the load offered by different user classes is very different, allocating capacity proportionally to the load does not result in optimal delay performance. Classes with a larger number of users share the allocated capacity more effectively due to statistical multiplexing within the class vs. ‘smaller’ classes. Therefore, delay performance can be further improved by allocating more than a proportional share of the capacity to the smaller classes, and less to the larger classes. The following optimization minimizes the mean sum queue length across all the classes, assuming each class corresponds to a M/GI/1-PS queue, thus minimizing user-perceived delay. I continue to assume

that the different base stations are decoupled.

Problem 3.2.2. *Determine a static, delay minimizing, decoupled schedule based on:*

$$\min_{\vec{\alpha}} \sum_{b=1}^N \sum_{k=1}^{K_b} \frac{\frac{\rho_{bk}}{R_{bk}(\vec{\alpha})}}{1 - \frac{\rho_{bk}}{R_{bk}(\vec{\alpha})}}$$

such that

$$\begin{aligned} \rho_{bk} &\leq R_{bk}(\vec{\alpha}), \forall b, k, \\ \sum_{l=1}^L \alpha_l &\leq 1 \quad \text{and} \quad \alpha_l \geq 0, l = 1, \dots, L. \end{aligned}$$

Proposition 3.2.2. $\sum_{b=1}^N \sum_{k=1}^{K_b} \frac{\frac{\rho_{bk}}{R_{bk}(\vec{\alpha})}}{1 - \frac{\rho_{bk}}{R_{bk}(\vec{\alpha})}}$ is a convex function of $\vec{\alpha}$, if $\frac{\rho_{bk}}{R_{bk}(\vec{\alpha})}$ is convex.

Proof. Let $u_{bk}(\vec{\alpha}) = \frac{\rho_{bk}}{R_{bk}(\vec{\alpha})}$. Then, $\frac{\frac{\rho_{bk}}{R_{bk}(\vec{\alpha})}}{1 - \frac{\rho_{bk}}{R_{bk}(\vec{\alpha})}} = \frac{u_{bk}(\vec{\alpha})}{1 - u_{bk}(\vec{\alpha})}$. $\frac{u_{bk}(\vec{\alpha})}{1 - u_{bk}(\vec{\alpha})}$ is a convex non-decreasing function of u_{bk} , and $u_{bk}(\vec{\alpha})$ is a convex function of $\vec{\alpha}$. Since the composition of a convex, non-decreasing function and a convex function is convex, $\frac{u_{bk}(\vec{\alpha})}{1 - u_{bk}(\vec{\alpha})}$ is a convex function of $\vec{\alpha}$. Therefore, the sum $\sum_{b=1}^N \sum_{k=1}^{K_b} \frac{u_{bk}(\vec{\alpha})}{1 - u_{bk}(\vec{\alpha})}$ is also convex. \square

Note that one can also consider other convex objective functions to capture other QoS metrics such as blocking rate, or other metrics such as power consumption at the base stations.

Fig. 3.5 exhibits the performance of the capacity maximizing schedule developed earlier vs. the above delay minimizing approach under a static

schedule. Both scenarios utilize two classes per base station along with the geometric approximation in Eq. (3.6) to estimate the class capacities, and three transmit power levels. The queue length-minimizing approach clearly outperforms the first approach where we allocated capacity proportionally to the class loads. This is mainly because this approach takes into account the potential each class has for statistical multiplexing. This queue length-minimizing approach will be used as the basis for developing further improved joint transmission schedules in the sequel.




Figure 3.5: Performance of capacity maximizing vs. delay optimal static, decoupled schedules with 2 classes per sector.

3.2.2 Dynamic Inter-Class Scheduling

As noted earlier, for downlink transmissions, the capacity perceived by users in neighboring base stations is independent of the user/class that a base station serves and depends only on the transmit power levels used by the

various base stations. Thus, when there are no active users in the class picked by the static schedule, the base station can dynamically pick an alternate class to serve without adversely affecting any of the cooperating base stations, i.e., without increasing the interference levels perceived by users. This class can be chosen by the base station based on different criteria, such as maximizing transmission rates, or serving the class with the largest number of active users. This is referred to as inter-class scheduling.

Definition 3.2.1. *An inter-class scheduler that upon exhausting the scheduled class performs processor sharing scheduling across all active users, is referred to as a dynamic processor sharing inter-class scheduler.*

I found through simulations that the delay performance of this strategy compared favorably to other policies. Note that this strategy allocates a proportionally larger rate to user classes that have a large number of active flows. When the traffic offered by all classes share similar characteristics, the optimized static schedule balances the expected number of active users in each class. Thus, this dynamic scheduling strategy attempts to align the available capacity to the particular realization of the offered load.

As can be seen in Figs. 3.6 and 3.7, inter-class scheduling significantly improves user delay performance and throughput, especially at light to moderate loads where mean delays are reduced by up to 40% as compared to the static scheme. At very low loads, it is still true that a scheme that transmits at maximum power without any coordination outperforms the coordination

scheme. Attempting to coordinate transmissions at low loads results in base stations needlessly using a lower power, thus transmitting at a lower rate even when the neighboring base stations are idle. Since the probability of simultaneous transmissions occurring is minimal at low loads, coordinating is not worthwhile.

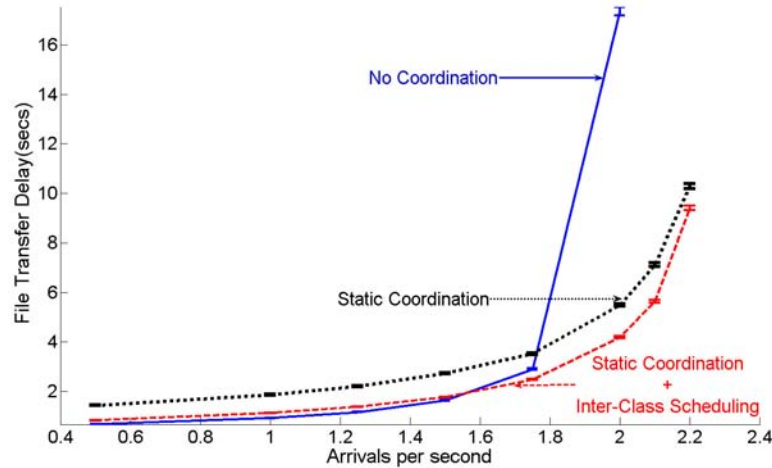


Figure 3.6: Average file transfer delays under delay minimizing static, decoupled schedules complemented by inter class scheduling with 2 classes per sector.

3.3 Optimizing the Coupled Model

The coordination schedules thus far have not taken into account the utilization of the neighboring base stations, and the performance coupling resulting from inter-cell interference. This is responsible for the poor performance at low loads. Determining the exact utilizations of the mutually coupled network of base stations for a particular joint transmission schedule is a difficult

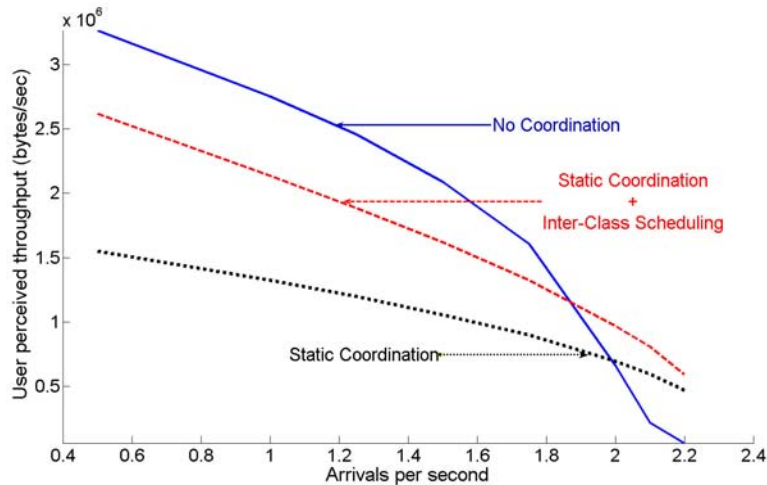


Figure 3.7: Average user throughput under delay minimizing static, decoupled schedules complemented by inter class scheduling with 2 classes per sector.

problem. However, if the utilizations can be estimated, the actual capacity perceived by classes in the dynamic, coupled system can be approximately determined. This would, in turn, allow us to pick better coordination schedules that explicitly take into account the degree to which the base stations are coupled.

Consider again the static coordination scheduling policies introduced in Sec. 3.2.1. Let $\vec{u}(\vec{\alpha}) = (u_{bk}(\vec{\alpha}) : b = 1, \dots, N, k = 1, \dots, K^b)$, where $u_{bk}(\vec{\alpha})$ is the resulting utilization of class k in base station b . As the base stations switch among different transmission profiles, a base station might not transmit in a designated profile if there are no active users at that base station. As a result, users in neighboring base stations can be served at enhanced rates. This effect can be modeled as a correspondence between a profile chosen as part of the joint transmission schedule, and a number of *induced* profiles in

which the network actually operates depending on class utilizations.

A base station remaining idle, with no users to serve just corresponds to using a transmit power level equal to zero, which corresponds to a valid joint transmission profile. When N base stations are being coordinated, each transmission profile can, in actual operation, result in one of up to 2^N profiles depending on which base stations are busy, or idle. Note that, these induced profiles are still a subset of \mathcal{L} . Let $\vec{\beta} = (\beta_m : m = 1, \dots, L)$ be the fractions of time actually spent in each profile when the transmission schedule specified by $\vec{\alpha}$ is followed. $\vec{\beta}$ is approximated as a function of $\vec{\alpha}$ and \vec{u} as follows.

$$\beta_m(\vec{\alpha}, \vec{u}) = \sum_{l=1}^L \alpha_l q_l^m(\vec{u}),$$

where $q_l^m(\vec{u})$ denotes the probability that, given the transmission profile l is chosen by the schedule, the network actually operates in profile m because the corresponding set of base stations are inactive. The vector $\vec{s}^{lm} = (s_b^{lm} : b = 1, \dots, N)$ is defined to take binary values as follows: $s_b^{lm} = 1$ if $p_b(l) = p_b(m)$, and 0 otherwise. $q_l^m(\vec{u})$ is then estimated assuming that the busy periods of the queues corresponding to the classes in different base stations are independent, i.e.,

$$q_l^m(\vec{u}) = \begin{cases} 0 & \text{if } \vec{c}(l) \neq \vec{c}(m), \\ 0 & \text{if } \vec{p}(m) \cdot (\vec{p}(l) - \vec{p}(m)) \neq 0, \\ \prod_{b=1}^N (u_{bc_b(l)})^{s_b^{lm}} (1 - u_{bc_b(l)})^{(1-s_b^{lm})} & \text{otherwise.} \end{cases}$$

Note that the network can only operate in a transmission profile m that allocates the same transmit power level as l , (or zero) to the base stations. This

is captured by the second case in the equation above. The fraction of time actually spent by the network in each induced profile can be computed in a similar fashion in the case of the dynamic coordination policy, except that q_l^m depends on the probability that there are no active users in any of the classes associated with a base station. — A joint transmission schedule optimizing users' delay performance is computed while taking into account the coupling across base stations iteratively. Let u_{bk}^z, R_{bk}^z represent the utilization, and rate estimates for the classes used in iteration z . Here, $\vec{\beta}^z = (\beta_m^z : m = 1, \dots, L)$ denotes the computed resultant schedule induced by the choice of time fractions $\vec{\alpha}^z = (\alpha_l^z : l = 1, \dots, L)$ in iteration z , and is a function of u_{bk}^z , and $\vec{\alpha}^z$. Let $\vec{\alpha}^{z*}$ denote the optimal coordination schedule found in iteration z , and $\vec{\beta}^{z*}$ the resultant induced schedule. Initially, $u_{bk}^1 = 1, \forall b, k$, and $R_{bk}^1 = R_{bk}^{A^*}$, and

$$\begin{aligned}\beta_m^z(\vec{\alpha}^z, \vec{u}^z) &= \sum_{l=1}^L \alpha_l^z q_l^m(\vec{u}^z) \\ u_{bk}^{z+1} &= \frac{\rho_{bk}}{R_{bk}^{(z)}(\vec{\beta}^{(z)*})}, \forall b, k.\end{aligned}$$

The optimization problem solved at each iteration is:

Problem 3.3.1. *Determining a delay minimizing schedule for the coupled network:*

$$\min_{\vec{\alpha}^z} \sum_{b=1}^N \sum_{k=1}^{K_b} \frac{\frac{\rho_{bk}}{R_{bk}^z(\vec{\beta}^z)}}{1 - \frac{\rho_{bk}}{R_{bk}^z(\vec{\beta}^z)}}$$

such that

$$\begin{aligned} \rho_{bk} &\leq R_{bk}^z(\vec{\beta}^z), \forall b, k, \\ \sum_{l=1}^L \alpha_l^z &\leq 1, \\ \alpha_l^z &\geq 0, l = 1, \dots, L. \end{aligned}$$

In the simulations that follow, the geometric rate approximation based on Eq. (3.6) is used:

$$R_{bk}^z(\vec{\beta}^z) = R_{bk}^{GA}(\vec{\beta}^z) = R_{bk}^A(\vec{\beta}^z) - \frac{\vec{\beta}^z T \Sigma_{bk} \vec{\beta}^z}{2R_{bk}^{(z-1)*}(\vec{\beta}^{(z-1)*})}$$

The objective function, and constraints in Problem 3.3.1 are convex, since $\vec{\beta}^z$ is a linear function of $\vec{\alpha}$, and the composition of a convex function and an affine function preserves convexity. This ensures that the problem can be efficiently solved at each iteration.

Fig. 3.8 illustrates the reduction in the average user-perceived delays that is achieved using two iterations in the above formulation. Here, the delay performance of the scheme with no coordination is not shown for clarity. Fig. 3.9 shows the increased user throughputs achieved from this coordination scheme, and also compares against the non-coordinated case. Now, at low loads, the coordinated transmission schedule does not penalize performance by restricting the transmit power level used by the base stations. The coordinated schedule performs as well as random scheduling at very low loads, when the probability of simultaneous transmissions at neighboring base stations is extremely low. At moderate to high loads, an optimized coordinated

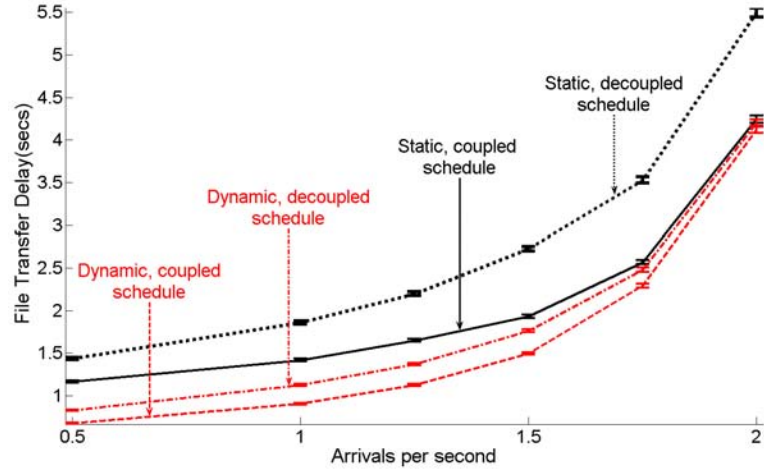


Figure 3.8: Average file transfer delays under schedules factoring inter-base station coupling, with 2 classes per sector.

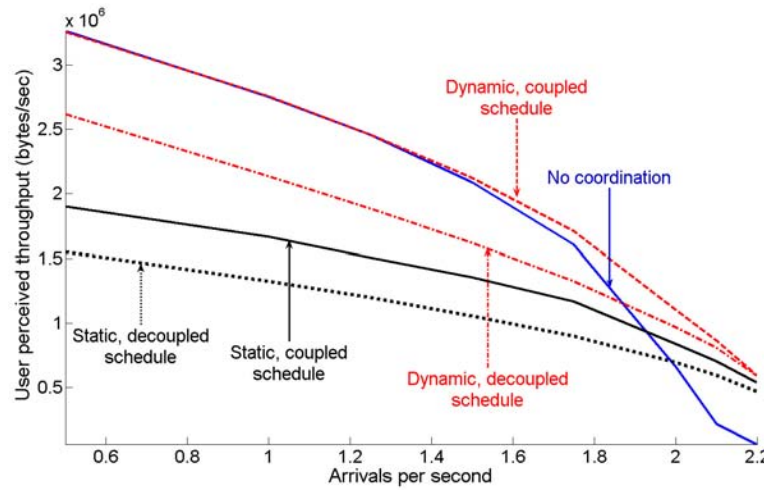


Figure 3.9: Average user throughput under schedules factoring inter-base station coupling, with 2 classes per sector.

scheduling scheme factoring the effect of coupling across base stations considerably outperforms the non-coordinated network, decreasing mean delays by over 80% as compared to a non-coordinated scheme. This ensures that the coordination scheme achieves good delay performance irrespective of the load on the network.

3.4 Further Benefits of Coordination

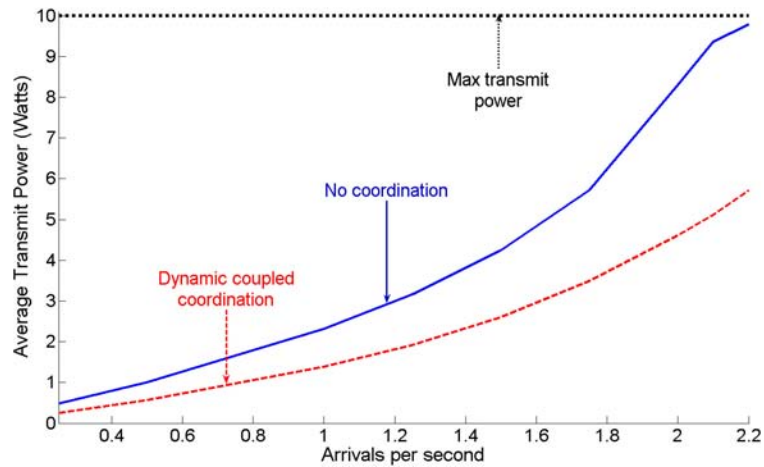


Figure 3.10: Average power consumed at the base stations.

In addition to improving delay performance and capacity, coordination has further benefits. As shown in Fig. 3.10, the average power expended by the base station is substantially reduced when coordination is used, e.g., 45% when the arrival rate is 2 users per second. This suggests a reduction in cooling costs at the base station, and also indicates that we can further improve delay performance if the base stations were subject to mean power constraints, and

could transmit at higher peak power levels.

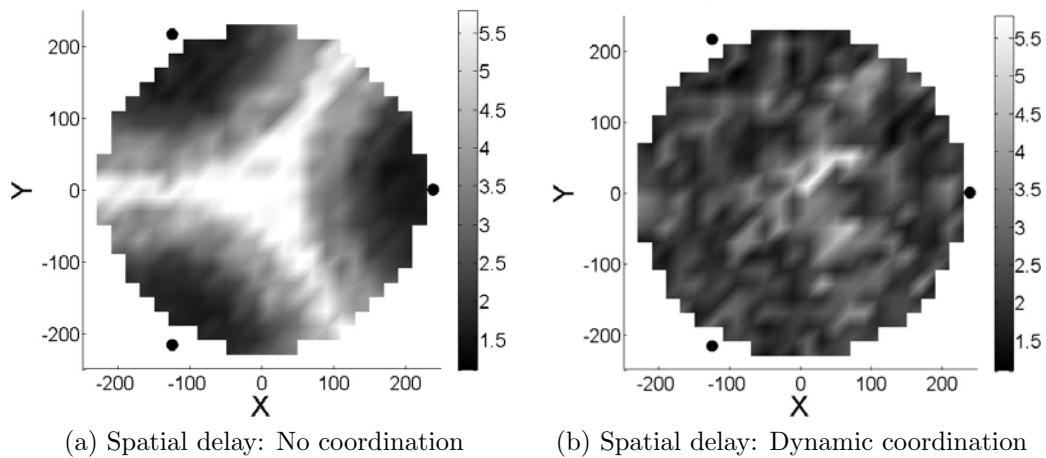


Figure 3.11: Distribution of user-perceived delay

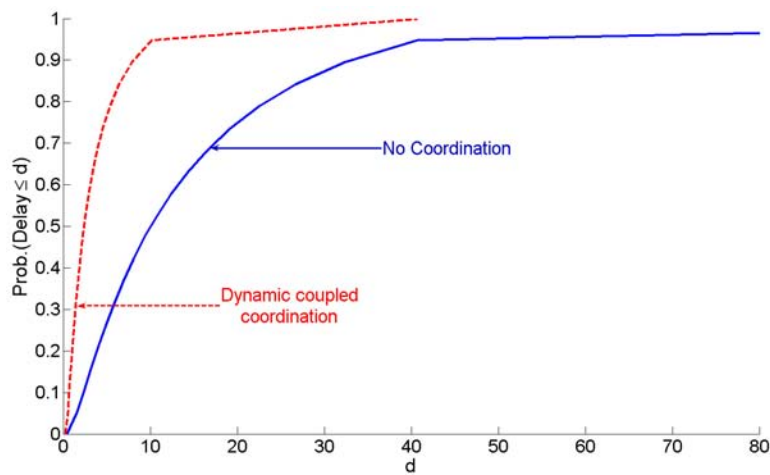


Figure 3.12: CDF of user delay

Figs. 3.11a, and 3.11b show the spatial delay distribution induced when no coordination is used, and the coordination scheme that minimizes the over-

all queue length, with $\lambda = 1.75$. As shown in Fig. 3.11b, when coordination is used, the average delays seen by users at different locations are much more spatially homogeneous. In particular, with no coordination users at the edge experience very poor performance. Under coordination, users' experience is virtually decoupled from their location in the coverage area.

Fig. 3.12 exhibits the distribution of delay across all users, when $\lambda = 2$. Coordination improves delay performance for all users, not just the ones at the edge. This is because the coordination scheme increases the probability that there are no active users at a base station. Thus, even though users close to the base stations are potentially served using lower transmit power levels, they benefit from lower interference levels.

3.5 The Importance of Being Traffic Aware

In a real-world wireless network, the traffic load is unlikely to be spatially homogeneous and may exhibit significant variations over time. For example, at different times of the day, one might see concentrations in different regions, e.g., coffee houses, lunch spots, public transportation, or depending on congestion patterns, etc. This chapter explores the potential gains from coordination in such scenarios. In particular, the focus is on understanding the degree to which optimizing for a particular load is beneficial. For example, if a fixed interference mitigation scheme such as a static fractional frequency reuse pattern is used, a natural choice is to optimize for a uniform distribution of users. The performance of our dynamic coordination scheme is first evaluated

when optimized for a uniform load, and is then compared to the case where it is tuned to the particular spatial traffic load.




Figure 3.13: A clustered user population.

The clustered traffic model used is as follows. User locations are constrained to a subset of the simulated area determined by the realization of a Boolean germ-grain model [80]. The grains of the Boolean model are discs of fixed radius, while the germs are distributed uniformly within the simulated area. The probability that an arrival's location falls in any of the discs is equal. The density of users at various points within the cell depends on the number of grains covering it. The density of users in areas covered by multiple grains is high, resembling a hotspot. Fig. 3.13 exhibits a realization of the spatial load with 70 germs, and discs with radius equal to one fifth the radius of the cell are used. Note that there are regions within the cell with sparse

user densities, and others where users tend to cluster. As the number of germs increases, the arrivals process converges to a homogeneous Poisson process. A small number of germs represents a user population that is highly clustered, with large variations in user densities within the coverage area.

In the simulations that follow, the number of germs is varied from 10 to 10,000 to simulate various degrees of clustering in the spatial load. For each case, we investigate the performance in twenty different realizations for the Boolean model. As explained previously, the actual load on the system is highly dependent on the spatial characteristics of the traffic. To roughly evaluate performance under vastly different spatial loads, we normalize the overall arrival rate so that the actual loads are comparable. Specifically, we choose the arrival rate that results in the base stations being 95% utilized, assuming all base stations transmit at maximum power all the time even if they have nothing to send. This operating point is computed using the harmonic mean, as described in Sec. 3.2.

Fig. 3.14a depicts the reduction in delay achieved by the two schemes compared to the non-coordinated case. It is clear that when the actual traffic being served is highly clustered (few germs), the traffic-independent coordination scheme performs much worse. In fact the average delays experienced by users are more than doubled vs the case with no coordination. As the number of germs is increased, and the spatial distribution of users approaches the uniform distribution, the traffic-independent scheme performs better than the non-coordinated one, and eventually catches up to the traffic-aware scheme.

The reduction in delay achieved by the traffic aware scheme appears independent of clustering in the loads. Note however, that the normalization used is imperfect, and in fact the measured loads were lower for scenarios subject to clustered loads. Thus I conjecture that subject to the same system load the gain achieved by the traffic aware will increase if the spatial load exhibits higher random clustering.

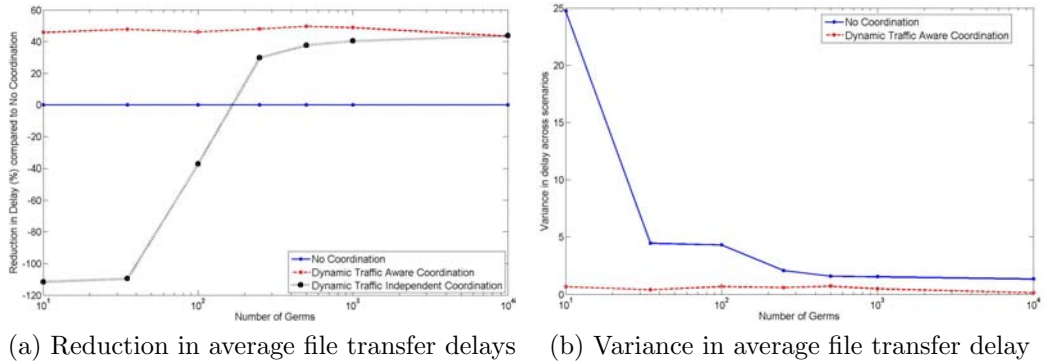


Figure 3.14: A scenario with non-homogeneous spatial load

Fig. 3.14b exhibits the variance across the scenarios under the traffic-aware scheme and the case where no coordination is used. This variance is induced by the sensitivity to inter-cell interference, and because different locations are affected very differently by interference. A non-coordinated system that serves a varying, non-homogeneous spatial distribution of users is prone to excessive variations in user perceived performance, and can experience very poor delay performance during time periods when it has to support a user population that is “poorly situated”. The traffic aware coordination scheme is

successful in shielding users from varying spatial loads, and achieves relatively homogeneous performance independent of where a user population lies. This decoupling of performance from both the variable spatial distribution of load, and the location of the users is a significant benefit.

Chapter 4

The Uplink Case

The impact of interference on the uplink is less pronounced than on the downlink. The key difference between the downlink and uplink cases results from the shift in the source of interference. On the downlink, the interference perceived by a user is independent of the particular user that is scheduled at the neighboring base stations. On the uplink, users' transmissions create interference at the neighboring base stations and a change in the location of the user can drastically alter the resulting degree of interference. This automatically modulates the interference caused at a base station as users at different locations are scheduled.

Compared to the downlink case where edge users always see high interference from active neighbors, the number of scenarios where users are severely limited by interference on the uplink is reduced. Consider the scenario depicted in Fig. 4.1 when all transmissions are at full power. In the downlink case shown in Fig. 4.1a, the edge user receives interference that is very close to the strength of the received signal. However, the interference at BS B is very low on the uplink as shown in Fig. 4.1b and the edge user's rate is impacted much less by interference. BS A does perceive higher interference in the uplink

scenario than the near user does on the downlink. However, the interference is still much weaker than the received signal at BS A, and the nearby user can be served at high rates.




Figure 4.1: Difference between uplink and downlink scenarios.

Users close to their serving base station typically have high channel gains to the serving base station and low channel gains to the neighboring base stations. Such users cause very low interference at their neighbors, and due to high received signal strengths are not severely affected by interference themselves. Users at the cell edge cause very high interference at the neighboring base stations, and additionally have to cope with low channel gains to the serving base station. Thus, the strength of interference seen by a base station depends both on the transmit power chosen by the users transmitting in the neighboring cells as well as the channel from the interfering user to the base station. While knowledge of the transmit power level at the neighboring base stations was sufficient to predict the interference received by a user or user class on the downlink, a different coordination mechanism is required for uplink coordination.

4.1 System Abstractions

The abstractions used for uplink coordination are similar to the downlink case. The difference is that *power profiles* are replaced by *joint interference profiles*. As noted earlier, fixing the transmit power used in all the cooperating cells/sectors is not sufficient to predict the interference at the base stations. The interference profiles directly bound the average interference that each cell is allowed to cause on neighboring base stations, so now an uplink transmission profile is specified by the combination of an interference profile and a class vector.

The maximum interference caused by transmissions at a cell to a neighboring base station is set at one of Q discrete levels, including 0. The $N \times N$ matrix $\vec{\mathbf{q}}^i$ specifies bounds on the interference each sector can cause at each of its neighboring base stations under interference profile i . The maximum average interference that transmissions at sector b can cause at sector m under interference profile i is denoted $q_{b,m}^i$. Note that $q_{b,b}^i = \infty$, for $b = 1, \dots, N$. The number of different interference profiles is denoted by $U' = Q^{(N(N-1))}$, and the number of joint transmission profiles is $L' = U'V$. Let $\mathcal{Q} := \{\vec{\mathbf{q}}^1, \dots, \vec{\mathbf{q}}^{U'}\}$ and $\mathcal{C} := \{\vec{c}^1, \dots, \vec{c}^V\}$ denote the sets of admissible joint interference profiles and class combinations respectively for the N base stations. Thus, each joint transmission profile l where $l = 1, \dots, L'$ is given by: $\vec{\mathbf{q}}(l) = \vec{\mathbf{q}}^i \in \mathcal{Q}$ and $\vec{c}(l) = \vec{c}^j \in \mathcal{C}$. A joint uplink transmission schedule corresponds to the fractions of time $\vec{\alpha} = (\alpha_l: l = 1, \dots, L')$ for which the network uses each transmission profile.

4.1.1 Determining user power levels under interference profiles

Consider a user u served by base station/sector b' . Recall that \vec{h}_u is the channel gain vector corresponding to user u . Users choose the largest transmit power that ensures that the average interference caused at all the neighboring base stations meets the constraints. The transmit power chosen by user u under interference profile i is given by $\min_{b=1,\dots,n} \frac{q_{b',b}^i}{h_u^b}$.

4.1.2 Estimating class rates

Each user can calculate the minimum rate achieved under each transmission profile after calculating the transmit power and using the upper bounds on received interference specified in the corresponding interference profile. Thus, the harmonic mean of the user rates provides a lower bound on the effective rate at which users in any class are served. Any of the estimates presented in Sec. 3.2.1.1 can then be used as an approximation of the class rates under a particular coordination schedule. In the simulation results presented in the sequel, the geometric mean rate approximation is used.

4.1.3 Optimizing the schedule

In order to optimize the user perceived delay performance, we use the methodology described in Sec. 3.3, with the transmission profiles defined as the combination of an interference profile and a class vector. The optimization problem solved at iteration z is:

Problem 4.1.1. *Determine a delay minimizing schedule for the coupled net-*

work based on:

$$\min_{\vec{\alpha}^z} \sum_{b=1}^N \sum_{k=1}^{K_b} \frac{\frac{\rho_{bk}}{R_{bk}^z(\vec{\beta}^z)}}{1 - \frac{\rho_{bk}}{R_{bk}^z(\vec{\beta}^z)}}$$

such that

$$\begin{aligned} \rho_{bk} &\leq R_{bk}^z(\vec{\beta}^z), \forall b, k, \\ \sum_{l=1}^{L'} \alpha_l^z &\leq 1 \quad \text{and} \quad \alpha_l^z \geq 0, l = 1, \dots, L'. \end{aligned}$$

While optimizing user performance is valuable, reducing the energy consumption while maintaining acceptable user performance is likely to be an important concern on the uplink. The arithmetic mean of the users' transmit powers under each transmission profile provides an approximation of the power consumed. Let $J_b(l)$ denote the average power consumed at base station b under transmission profile l . The average power consumption under a joint transmission schedule $\vec{\alpha}$ can then be estimated as $\sum_{l=1}^{L'} \sum_{b=1}^N \alpha_l^z J_b(l)$. The coordination framework presented above can be modified to minimize the weighted sum of the mean user delay and the mean power consumption under a coordination schedule. The weight chosen, denoted by γ , represents the relative importance of conserving energy versus minimizing delay, and is a parameter that can be adjusted. The objective function to be minimized at each iteration in the methodology of Sec. 3.3 is given by

$$\sum_{b=1}^N \sum_{k=1}^{K_b} \frac{\frac{\rho_{bk}}{R_{bk}^z(\vec{\beta}^z)}}{1 - \frac{\rho_{bk}}{R_{bk}^z(\vec{\beta}^z)}} + \gamma \sum_{l=1}^{L'} \sum_{b=1}^N \alpha_l^z J_b(l).$$

In the sequel, processor sharing is again used as the intra-class scheduling

policy and dynamic processor sharing as the inter-class scheduling discipline - see Def. 3.2.1.

4.2 User Performance

Uplink file transfers in the three sector scenario shown in Fig. 2.2a are simulated with user requests distributed as a homogeneous Poisson process in space. Fig. 4.2 exhibits plots of the mean delay performance, while Fig. 4.3

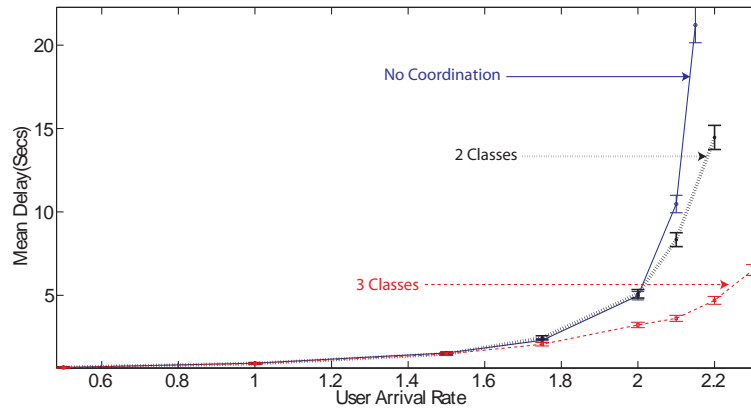


Figure 4.2: Uplink delay performance under delay minimizing schedules that account for inter-base station coupling.

shows plots of the average throughput achieved under the delay-minimizing joint uplink transmission schedule. As a result of using a methodology that accounts for inter-base station coupling, we see that the delay and throughput performance always equals or improves on those obtained for the uncoordinated scheme. At high loads, mean delay is improved by about 40% when 2 classes are used per base station, and by up to 80% when 3 classes are used per base station. The average throughput is increased by up to 27% when

2 classes are used per sector and by up to 90% when 3 classes are used. At moderate to high loads, using a finer grain classification of users results in significant performance gains. Since individual users adjust their transmit powers to satisfy the constraints imposed by the interference profile, the variability in rates across users in an interference profile is increased. Users near the boundary have low channel gains to their serving base station, and are additionally forced to use lower power levels in order to limit the interference that they cause. Using a larger number of classes improves the estimates for the class rates, and also enables the schedule to accurately differentiate between users at different locations.

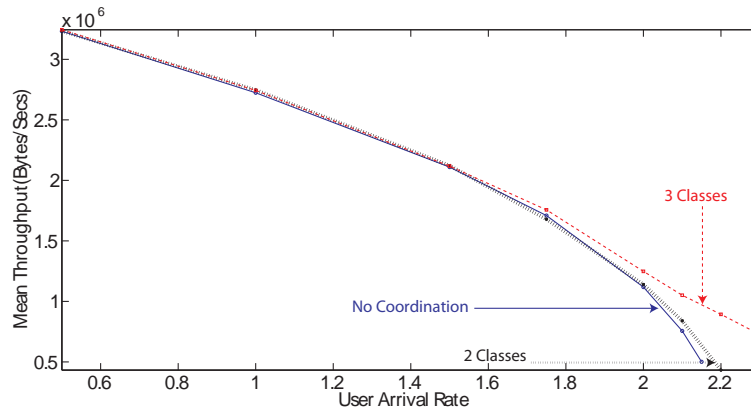


Figure 4.3: Average throughput under schedules factoring inter-base station coupling.

Fig. 4.4 exhibits plots of the mean delay achieved against the average power consumption under the non-coordinated system as well as the coordinated schedule that minimizes a weighted sum of mean delay and mean power. The overall rate at which users arrive into the system is fixed at 2.1 arrivals

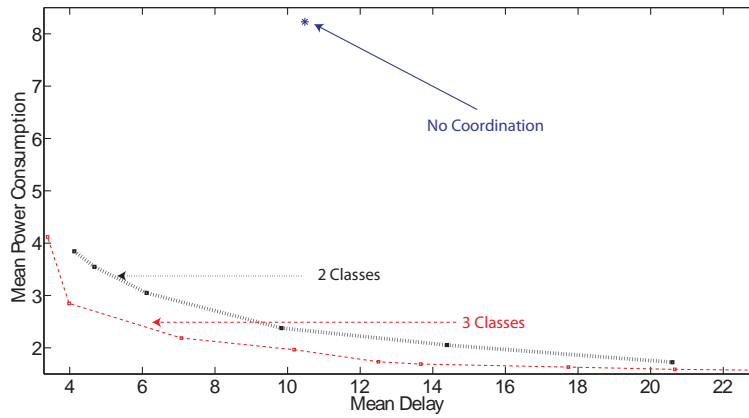


Figure 4.4: Power-Delay tradeoffs on the uplink.

per second and they are assumed to be spatially homogeneous. The weight γ is varied to demonstrate the trade-offs between energy saving and performance that can be achieved through coordination. Average energy consumption can be decreased by up to 75% through coordination while achieving delay performance identical to the non-coordinated system. Note that even when the coordination scheme is tuned to minimize user perceived delay performance, the average power consumption is lowered by approximately 50% relative to the non-coordinated system. This improvement in performance and energy efficiency is achieved while simultaneously ensuring that the average delays seen by users at different locations are much more spatially homogeneous relative to the case with no coordination, similar to the downlink case (graph omitted).

Chapter 5

Final Observations

In order to determine the joint transmission schedule, base stations in a coordinating cluster need to exchange user class loads as well as information required to compute class capacities. Two different architectures could potentially be used to exchange information across base stations. In a centralized version, base stations could use backhaul communication to exchange information. Base stations could also leverage the users in the system to relay these system messages. Two coordinating base stations will likely find edge users that can communicate with both of them. Such users could forward messages between base stations allowing a completely distributed mechanism. Also, the transmission schedule could be computed centrally and communicated to the base stations, or each base station could individually solve the identical optimization problem and determine the schedule when tie-breaking mechanisms and an ordering of the profiles has been previously agreed upon.

5.1 Cost of Base Station Coordination

Consider the case where each of the N coordinating base stations in the cluster use K user classes. In order to exchange spatial load informa-

tion, the total communication required among base stations within the cluster is $KN(N - 1)$ data points. If the arithmetic mean based approximation, see Eq. (3.4), is used for computing class capacities, only the mean capacity to each class under each power profile needs to be exchanged. Thus, the total information exchanged among base stations is $KN(N - 1)U$ data points. However, when the geometric mean approximation, see Eq. (3.6), is used for the class capacities, the second moment has to be exchanged in addition to the mean, and the total amount of data to be exchanged is given by $KN(N - 1)U + N(N - 1)KU^2$. Thus the total information required grows quadratically in the number of coordinating base stations, linearly in the number of classes, and linearly or quadratically in the number of power profiles depending on the method used to approximate class capacities. Computing the joint transmission schedule requires the solution of a convex optimization problem with UV variables, and $(NK + UV + 1)$ constraints. Note that this information exchange and schedule computation happens when the long-term average spatial load changes (on the order of hours). For example, in a scenario with three coordinating base stations, three transmit power levels, and three classes per base station, the total information exchanged per base station when the geometric mean approximation is used is: 4542 data points.

5.2 Reducing Coordination Overheads

Reducing the communication and computational overhead associated with the coordination scheme is of high importance. To this end, it would be

advantageous to eliminate some of the redundant transmission profiles. One relatively simple approach to pruning the number of profiles we have to deal with is eliminating the power vectors that are dominated by at least one other power vector.

Definition 5.2.1. *The power vector indexed by i_0 is said to be dominated by another vector indexed by i_1 if $\vec{r}^{i_1} \geq \vec{r}^{i_0}$, where $l_0 \equiv \{i_0, j\}$ and $l_1 \equiv \{i_1, j\}$ for any class vector j .*

Theorem 5.2.1. *Under a simple physical interference model, the power vector \vec{p} is dominated by the power vector $\mu\vec{p}$, if $\mu \geq 1$.*

Proof. The only assumption made about the interference model is that the received signal strength that a user sees is proportional to the transmit power. Denote the received power from base station b at the user of interest by q_b . Consider an user in base station b_1 . The signal transmitted by base station b_1 is the desired signal, and the transmissions from all other base stations constitute interference.

When \vec{p} is used, the transmit power used at base station b_1 is p_{b_1} , and the received signal is q_{b_1} . The interference that the user sees is $\sum_{b=1, b \neq b_1}^N q_b$. If we denote the noise power by p_{noise} , the SINR that the user sees is given by
$$\frac{q_{b_1}}{\sum_{b=1, b \neq b_1}^N q_b + p_{\text{noise}}}.$$

Now, when $\mu\vec{p}$ is used, the transmit power used at base station b_1 is μp_{b_1} , and the received signal is μq_{b_1} . The interference that the user sees is $\sum_{b=1, b \neq b_1}^N \mu q_b$. If we denote the noise power by p_{noise} , the SINR that the user

sees is given by $\frac{\mu q_{b_1}}{\sum_{b=1, b \neq b_1}^N \mu q_b + p_{\text{noise}}} = \frac{q_{b_1}}{\sum_{b=1, b \neq b_1}^N q_b + \frac{p_{\text{noise}}}{\mu}}$. For $\mu \geq 1$, the SINR in this case is clearly higher than the SINR that the user sees when \vec{p} is used. Thus, the rate that the user sees is also higher. \square

Thus, as long as minimizing delay or maximizing throughput is our only concern, we can eliminate any vector \vec{p} if a vector $\mu\vec{p}$ is available, $\mu \geq 1$. Note that eliminating a power vector eliminates all the transmission profiles obtained as a combination of that power vector and any class vector. As a consequence of the above theorem, we can also conclude that, when we have a large number of power vectors and the granularity of the vectors is very fine, only the power vectors with at least one base station operating at maximum power need to be considered.

A more general approach to eliminating redundant transmission profiles is to examine the structure of the rate region generated by the discrete set of power profiles and class vectors. The convex hull of the UV vectors in $\sum_{b=1}^N K_b$ dimensions is the rate region. Only the profiles that are the vertices of the convex hull are necessary to obtain the region. The other profiles are superfluous. Algorithms exist to find the convex hull of a finite set of vectors in any dimension. However, the problem is complicated when the dimension is greater than 3. In the case of 3 dimensions or less, the problem of finding the convex hull can be solved efficiently, with complexity $x \log x$, where x is the number of vectors [35, 67]. Each class vector can be picked one by one, and for each class vector, the profiles that are dominated can then be eliminated. If we consider groups of three cells, this problem is in three dimensions and

can be solved efficiently. The complexity of the algorithm to eliminate the superfluous vectors is $VU \log U$.

5.3 Integration with Physical Layer Methods and Future Research Avenues

A low complexity, system-level approach that substantially improves performance perceived by best-effort users without requiring high channel measurement and estimation, communication, and computational overheads was presented. The proposed approach simultaneously achieved spatially homogeneous performance while also reducing the transmit power requirements. Future wireless networks could include physical layer techniques such as interference cancellation. Such techniques are likely to be imperfect due to associated measurement and estimation errors. One can view the proposed approach as complementary to an imperfect physical layer interference mitigation scheme, and use it as an overlay that takes into account traffic and environmental characteristics and spatial diversity under imperfect interference cancellation. Fig. 5.1 exhibits the performance of the same three base station system, without coordination, and coordination and dynamic scheduling as considered earlier, but then with no coordination under an idealized regime where 50% of all interference is canceled, and with both coordination and interference cancellation. These results are an indication that even under an aggressive regime where one is able to cancel 50% of all interference seen by users irrespective of their location, that system-level coordination will provide

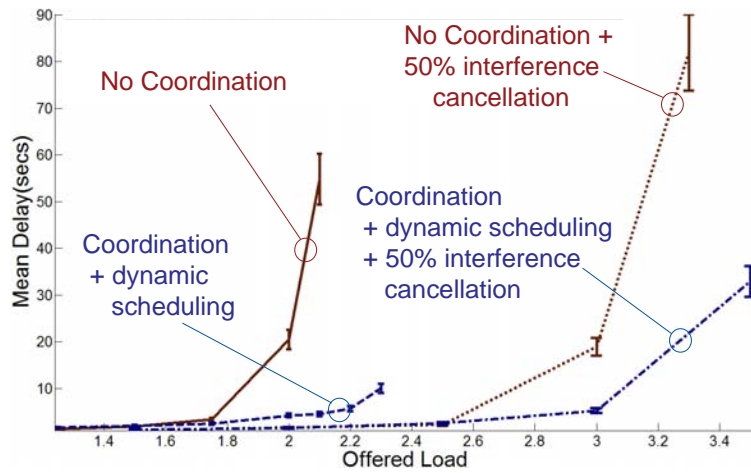


Figure 5.1: Performance comparison schemes with no coordination, with coordination, with interference cancellation and then coordination and interference cancellation.

additional performance benefits - albeit somewhat reduced.

System-level coordination can also be profitably used in the case of (packet) delay sensitive traffic, as long as suitable complementary dynamic user scheduling schemes are developed to meet users' QoS requirements. A factor that we have not considered in this paper is user mobility. Mobile users simply transition from one class to another as they move about within the network, and can potentially be treated as premature departures from a class arriving at another. These topics are left for future research.

Part II

Optimizing User Association Policies

Chapter 6

Motivation and Related Work

A basic problem in a wireless network is to decide which base station should serve a new user request. Clearly, the policy used to assign serving base stations to user requests greatly influences the performance experienced by users in the network. The simulation results presented in Part I assume that user requests are served by the base station that provides the strongest signal, which is not unlike policies in use in networks today. In the following chapters, I study the impact of policies to associate users with base stations/access points on the downlink flow-level performance in interference limited wireless networks. To aid our understanding of the problem and for ease of analysis, base station coordination is not considered in this part.

6.1 Related Work

Without having access to good performance models, many researchers have resorted to intuitive objectives such as load balancing across system resources. Load balancing policies for wireless systems were first studied in the context where the traffic carried by the network is voice, and frequency reuse is used to combat interference. Load balancing algorithms to minimize outage

probability were presented in [26, 51, 55]. The idea in a circuit switched network is to direct a new call to the base station with the greatest number of available channels, so that the probability that future incoming calls will be blocked because of lack of resources is minimized. [8, 89] focus on packet based voice networks, and consider balancing schemes that account for the variation in service capacity to users at different locations.

This philosophy has been extended to address the case of best effort traffic. When the wireless network is subject to spatially heterogeneous traffic loads, emphasis has been placed on the development of schemes that try to balance the load across base stations. Schemes that maximize base station and network utility respectively using joint base station assignment and scheduling were presented in [57, 64]. The schemes use a utility function that is decreasing in base station load to divert traffic away from heavily loaded cells. The scheme proposed in [64] is a centralized one that incurs excessive communication and computational overheads.

In [28], a load balancing scheme which requires much less coordination is considered. The scheme tries to explicitly balance the load across base stations, taking into consideration both the long term rate at which users can be served, and their load. Another idea which was proposed, in [28, 37], is to lower the strength of the pilot signals that heavily loaded base stations broadcast, so as to discourage users from joining them. Such an approach has also been adopted in Qualcomm's systems [1]. However, such manipulation could potentially interfere with the estimation procedures used by the users

to accurately estimate the channel to support adaptive modulation, which requires stable pilot power.

A joint power control and cell selection scheme for CDMA networks, referred to as cell breathing was developed in [37], and a very similar scheme was developed in [90]. Users' power levels are adjusted and they are switched among cell sites depending on the load levels, to balance loads. Cell breathing refers to the notion of cells expanding or contracting in space, to equalize loads. A scheme that is similar in spirit is proposed in [73], called MAC-cell breathing, that attempts to balance the load in all base stations. Users attempt to join the base station that provides the highest throughput, taking into account channel conditions as well as cell congestion. A centralized controller instructs heavily loaded base stations to reduce the rates allocated to users at the edge, thus encouraging them to join a neighboring base station.

The focus in these schemes is to ensure that the load being served by different base stations in a neighborhood is as similar as possible. The implicit assumption made by all the above schemes is that fewer users in a cell corresponds to increased capacity to them. While this is certainly true if the users were seeing constant interference, we will see that this is surprisingly not the case in the presence of dynamic interference. The following sections demonstrate that such load balancing, be it greedily done by users or across the system, may be counter productive when there is dynamic coupling due to interference.

The case of dynamic traffic, with the associated bursty interference,

has not been extensively studied. In [10], the effect of equalizing the load in neighboring base stations was studied through simulation, and it was observed that load balancing did not make much of a difference under heavy load. This problem is also studied in [18], but under the assumption that transmissions are orthogonal. The impact of dynamic interference was also demonstrated in [92], wherein the problem of load balancing in a hybrid wireless local area/wide area network was studied using approximations proposed in [9].

The stability region of a dynamic system with interacting servers under load balancing strategies was examined in [17]. The stability region was explicitly characterized in the case of a two server system, and a lower bound on the stability region was obtained for systems with multiple servers. The stability region in the case of static load balancing policies and a class of dynamic policies was also studied in [47]. While the above papers address the question of determining the network capacity, they do not provide insight into designing user association policies to optimize performance perceived by users in a system serving a load that is in the interior of the stability region. In contrast, the focus here is on flow level performance, i.e., the actual file transfer delays experienced by users.

6.2 Is Load Balancing Always Optimal?

Consider the user association problem exhibited in Fig. 6.1a. Assume again that the base stations share the same spectrum, so they interfere with each other when they are concurrently active, which in turn reduces the maxi-

imum transmission rate to users. For simplicity, assume user requests to download files arrive uniformly between Base Stations 1 and 2. If both the network and traffic demands are *symmetric*, one might intuitively expect that a static policy that associates arrivals with the closest base station, i.e., the one that delivers the strongest signal, and thus balances the offered load, would be ‘optimal’.

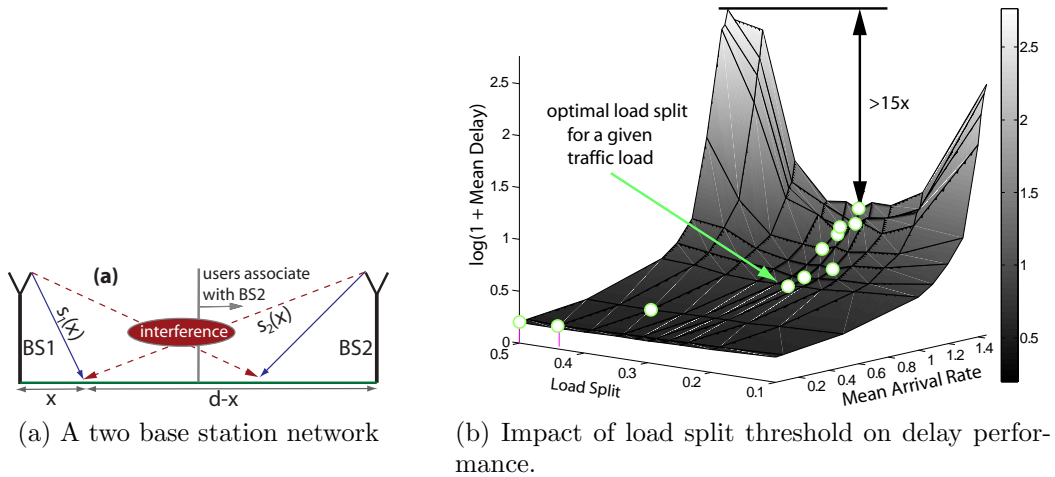


Figure 6.1: The user association problem in a two base station scenario.

Fig. 6.1b shows the simulated delay performance (explained in more detail in the sequel) when the load split between the base stations is varied from 0.5 (even division of load) to 0.1 (highly asymmetric load division). The results show that the optimal load division depends on the intensity of the offered load, and is not always balanced but can be significantly asymmetric. As exhibited in the figure, where mean delays are plotted on a logarithmic

scale, the performance implications can be substantial; load balancing may achieve mean delays 15 times higher versus an optimal asymmetric split. This result is surprising, and reveals the complexity and substantial impact that dynamic coupling can have in the context of wireless networks. This motivates the need for careful analysis as well as comparisons with more complex user and system greedy dynamic policies.

6.3 Contributions

My contributions in this context include:

1. I propose a methodology to optimize the performance of wireless systems coupled through dynamic interference and apply it to study one and two dimensional networks. To our knowledge, prior to this work, no closed-form or good approximations were available for wireless systems with coupled queues.
2. For a dynamic model of the user association problem in one dimension, I show that optimal static policies are threshold based. Surprisingly, even for a symmetric network, a policy which balances load can be highly suboptimal. Moreover, I show that asymmetric policies can improve average delays seen by users at *all* spatial locations.
3. I show that an optimized static policy (asymmetric) can substantially outperform dynamic policies which are greedy from the user's or system's points of view and achieves performance close to that of a 'repacking' policy. This suggests that an important objective for protocol and network design will be to achieve such load asymmetries.

4. I present ISAP (Interference-Sensitive, Adaptive Policy), a novel load association policy that uses measurements to infer the degree of performance coupling due to inter-cell interference, and adapts to it.

5. I demonstrate through extensive simulations that the proposed policy consistently outperforms conventional, load balancing based approaches under both spatially homogeneous and heterogeneous loads. These results also show that the performance of conventional dynamic schemes is highly dependent on the spatial load, and no single best scheme can be identified.

In Chapter 7, a generalized system of coupled queues is analyzed and a method to find upper/lower bounds on performance measures is developed. This method is used in Chapter 8 to optimize static association policies based on the long term spatial loads being served. Finally, Chapter 9 presents an adaptive user association policy that infers the pertinent characteristics of the spatial load distribution through simple measurements.

Chapter 7

Analyzing a Coupled Queuing System

At the heart of the interference limited wireless network serving a dynamic user population lies a system of coupled queues, and in order to understand the performance users perceive in the network we need to analyze this underlying system. The focus in this chapter is on a queuing system with coupled processors, where the rate at which users in a queue are served depends on the lengths of the other queues in the system. In particular, we will consider those systems where the service rate at each queue varies depending on the set of queues in the system with non-zero queue lengths. Such a coupled-processors model arises naturally in the study of systems where a resource is shared by several classes of customers

7.1 Prior Work

A queue with multiple user classes using the generalized processor sharing (GPS) discipline [65] is an example of a coupled queuing system, where the performance experienced by users in different classes are coupled. Large deviations asymptotics of the workload in GPS systems have been derived in [6, 30, 94, 95], and the effect of serving customers with long-tailed service times

has been analyzed in [14, 15]. Approximate asymptotic formulas for the stationary distribution of the number of customers in a system of two coupled queues are obtained in [53]. However, steady state queue length distributions are known only in some special cases.

The uniformization technique was used in [54] to study a special case of the GPS model with two classes where the classes are identical. The coupled-processors model with exponentially distributed service requests and the number of classes/queues restricted to two was analyzed in [31, 36, 54]. The generating function of the steady-state queue length distribution was obtained in [31] by solving a Riemann-Hilbert boundary value problem. Closed-form expressions were obtained in [36] for the special case where the queueing system is work conserving, without resorting to the formulation of a Riemann-Hilbert problem. However, the techniques used in the above papers cannot be extended to systems with more than two queues. The stability of queueing systems with coupled processor systems was examined in [20], and even in simple cases, closed-form stability conditions cannot be obtained and one has to resort to numerical techniques.

In this chapter, bounds on the moments of the steady-state queue-length in a system consisting of N coupled queues are obtained by studying the stochastic recursive equations that govern the system. Lower and upper bounds on the moments of the queue length are obtained by formulating and solving a semidefinite relaxation of the original problem. These bounds can be made progressively tighter at the expense of increasing the complexity of

the associated semidefinite program. Our model is motivated by the one used in [5] to analyze GI/GI/1 queues, and draws on results obtained in [7, 56].

7.2 System Model and Notation

We consider a system of N queues. Users arrive at queue n as a Poisson process, with mean arrival rate λ_n . We assume that the users require exponential service times with mean 1. We denote the queue length at each queue at time t by $\vec{Q}(t) = (Q_n(t), n = 1, 2, \dots, N)$. The queue service rates depend on the subset of queues in the system with non-zero queue length, where the number of possible subsets is 2^N . The status of the system at time t is captured by a vector $\vec{\Delta}(t)$ of length n that takes values $\vec{\delta}_i, i = 1, \dots, 2^N$. Suppose $\vec{Q}(t) = \vec{q}$ has associated system status vector $\vec{\delta}^i$, then $\delta_n^i = \mathbf{1}(q_n > 0)$. Under $\vec{\delta}^i$, the rate at which queue n serves users is denoted by $\mu_n^{\delta^i}$. We assume that there is a well defined maximum rate μ^* that bounds the rate at which any queue can be served, irrespective of the state of the system.

The queue length process evolves as a continuous time Markov chain, parametrized by the rates defined above. The maximum rate of transitions is bounded by $\eta = \sum_{n=1}^N \lambda_n + N\mu^*$, thus the continuous time Markov chain can be uniformized by introducing fictitious events that cause no change in the state of the Markov chain. With a slight abuse of notation, let $\vec{Q}(k) = (Q_n(k), n = 1, 2, \dots, N)$ denote the state of the uniformized discrete time Markov chain at discrete time step k , and $\vec{\Delta}(k)$ the associated system state vector. The transition probabilities for the uniformized Markov chain when

$\vec{Q}(k) = \vec{q}$, which corresponds to state vector $\vec{\delta}$ are as follows:

$$\begin{aligned} \mathbf{P}(\text{arrival to queue } n | \vec{Q}(k) = \vec{q}) &= \frac{\lambda_n}{\eta}, \\ \mathbf{P}(\text{departure from queue } n | \vec{Q}(k) = \vec{q}) &= \frac{\mu_n^{\delta_i}}{\eta}, \\ \mathbf{P}(\text{no change} | \vec{Q}(k) = \vec{q}) &= 1 - \frac{\sum_{n=1}^N \lambda_n + \mu_n^{\delta_i}}{\eta}. \end{aligned}$$

Note that, if it exists, the uniformized chain's stationary distribution is identical to that of the original. Also, its evolution can be represented as a stochastic recursion

$$\vec{Q}(k+1) = \vec{Q}(k) + \vec{X}(k), \quad k = 0, 1, \dots,$$

where $\vec{X}(k) = (X_n(k), n = 1, 2, \dots, N)$ denotes increments in the queues. An arrival into queue n at iteration k is represented by $X_n(k) = 1$, a departure by $X_n(k) = -1$ and if the transition corresponds to the self-loop, $\vec{X}(k) = \vec{0}$. Note that $\vec{X}(k)$ and $\vec{Q}(k)$ are not independent, e.g., one can not have a departure from an empty queue. When the system is stable, there is a stationary distribution for (\vec{Q}, \vec{X}) [12, 62] such that

$$\vec{Q} \stackrel{d}{=} g(\vec{Q}, \vec{X}) := \vec{Q} + \vec{X} \tag{7.1}$$

where $\stackrel{d}{=}$ denotes equality in distribution. Our goal is to characterize the behavior of the queuing system by formulating a moments based approach to bound functions of the moments of \vec{Q} . As in [5], we only use information on the moments of \vec{X} derived from the uniformized Markov chain in the formulation.

7.3 Performance Bounds

Let ψ denote a joint distribution for (\vec{Q}, \vec{X}) on $\mathcal{S} = \mathcal{S}_{\vec{Q}} \times \mathcal{S}_{\vec{X}} \subseteq \mathbb{Z}_+^N \times \mathbb{Z}_+^N$ satisfying (7.1) and with marginals $\psi_{\vec{Q}}$ and $\psi_{\vec{X}}$. Note that, $\psi_{\vec{Q}}$ and $\psi_{\vec{X}}$ are not necessarily independent and the joint distribution ψ cannot be expressed as a product form. Eq. (7.1) can in this case be rewritten as

$$\psi_{\vec{Q}} = \psi g^{-1}.$$

We partition $\mathcal{S}_{\vec{Q}}$ into 2^N regions where the *same* set of queue lengths are non-zero, i.e., $\mathcal{S}_{\vec{\delta}^i} := \{\vec{q} : \delta_n^i = \mathbf{1}(q_n > 0), n = 1, \dots, N\}$ for $i = 1, \dots, 2^N$. Let $\psi^{\vec{\delta}^i}$ and $\psi_{\vec{Q}}^{\vec{\delta}^i}, \psi_{\vec{X}}^{\vec{\delta}^i}$ be the conditional distributions for (\vec{Q}, \vec{X}) and its marginals given $\vec{Q} \in \mathcal{S}_{\vec{\delta}^i}$. Note that for all states in $\mathcal{S}_{\vec{\delta}^i}$ the queues share the same service rates, so follows \vec{Q} is conditionally independent of \vec{X} given $\vec{Q} \in \mathcal{S}_{\vec{\delta}^i}$, i.e.,

$$\psi^{\vec{\delta}^i} = \psi_{\vec{Q}}^{\vec{\delta}^i} \psi_{\vec{X}}^{\vec{\delta}^i}, \quad i = 1, \dots, 2^N.$$

We shall use multi-index notation in formulating our bounds. For $\vec{\alpha} \in \mathbb{Z}_+^N$ and $\vec{Y} \in \mathbb{R}^N$, we let $\vec{Y}^{\vec{\alpha}}$ denote the term $Y_1^{\alpha_1} \dots Y_N^{\alpha_N}$, and let $|\vec{\alpha}| = \sum_{n=1}^N \alpha_n$. For $r \in \mathbb{N}$, we define

$$m_{\vec{\delta}^i}^{\vec{\beta}} = \mathbf{E}_{\psi_{\vec{X}}^{\vec{\delta}^i}} \left[\vec{X}^{\vec{\beta}} \right], \quad |\vec{\beta}| \leq 2r \text{ and } i = 1, \dots, 2^N.$$

Given the transition probabilities on each region $\mathcal{S}_{\vec{\delta}^i}$, these can be easily computed. Bounds on functions of the form $\mathbf{E}_{\psi_{\vec{Q}}} \left[\sum_{|\vec{\gamma}| \leq 2r} w_{\vec{\gamma}} \vec{Q}^{\vec{\gamma}} \right]$, where $w_{\vec{\gamma}}$ are constant weights can be obtained by optimizing over distributions ψ satisfying the following constraints:

Problem 7.3.1. Given $\mathcal{S}, \mathcal{S}_{\vec{Q}}$ and $\mathcal{S}_{\vec{X}}$ solve:

$$\sup / \inf_{\psi} \quad \mathbf{E}_{\psi_{\vec{Q}}} \left[\sum_{|\vec{\gamma}| \leq 2r} w_{\vec{\gamma}} \vec{Q}^{\vec{\gamma}} \right]$$

s.t.

$$\psi_{\vec{Q}} = \psi g^{-1} \tag{7.2}$$

$$\psi^{\vec{\delta}^i} = \psi_{\vec{Q}}^{\vec{\delta}^i} \psi_{\vec{X}}^{\vec{\delta}^i}, \quad i = 1, \dots, 2^N, \tag{7.3}$$

$$\mathbf{E}_{\psi_{\vec{X}}^{\delta_i}} \left[\vec{X}^{\vec{\beta}} \right] = m_{\delta_i}^{\vec{\beta}}, \quad i = 1, \dots, 2^N \text{ and } |\vec{\beta}| \leq 2r$$

$$\mathbf{E}_{\psi}[1] = \mathbf{E}_{\psi_{\vec{Q}}}[1] = \mathbf{E}_{\psi_{\vec{X}}}[1] = 1, \tag{7.4}$$

$$\psi \in \mathbb{M}(\mathcal{S}), \psi_{\vec{Q}} \in \mathbb{M}(\mathcal{S}_{\vec{Q}}), \psi_{\vec{X}} \in \mathbb{M}(\mathcal{S}_{\vec{X}}). \tag{7.5}$$

Here, $\mathbb{M}(\mathcal{S})$, $\mathbb{M}(\mathcal{S}_{\vec{Q}})$, and $\mathbb{M}(\mathcal{S}_{\vec{X}})$ are sets of positive Borel measures supported on \mathcal{S} , $\mathcal{S}_{\vec{Q}}$, and $\mathcal{S}_{\vec{X}}$ respectively, and (7.4) ensures they are probability measures. The parameter r controls the degree of accuracy of such bounds [56]. As $r \rightarrow \infty$ the distribution of \vec{X} is specified exactly, in turn uniquely determining the distributions of \vec{Q} and (\vec{Q}, \vec{X}) . To allow numerical computation, we further relax Problem 7.3.1 based on joint moments of degree no higher than $2r$.

7.3.1 The Moments Based Approach

For all $\vec{\alpha}, \vec{\beta}$ such that $|\vec{\alpha}| + |\vec{\beta}| \leq 2r$ and $k = 1, \dots, 2^N = K$, we define the decision variables:

$$x_k^{\vec{\alpha}\vec{\beta}} := \mathbf{E}[\vec{Q}^{\vec{\alpha}} \vec{X}^{\vec{\beta}} | \vec{Q} \in \mathcal{S}_{\vec{\delta}^k}] \mathbf{P}(\vec{Q} \in \mathcal{S}_{\vec{\delta}^k}).$$

We express the objective function and the relaxed versions of the constraints in Problem 7.3.1 in terms of the above decision variables as follows:

Objective function: The objective function in Eq. (7.2) can be written in terms of the conditional expectations as follows using the theorem of total expectation. This expression can then be expressed as a linear function of the decision variables:

$$\begin{aligned} \mathbf{E}_{\psi_{\vec{Q}}} \left[\sum_{|\vec{\gamma}| \leq 2r} w_{\vec{\gamma}} \vec{Q}^{\vec{\gamma}} \right] &= \sum_{k=1}^K \mathbf{E} \left[\sum_{|\vec{\gamma}| \leq 2r} w_{\vec{\gamma}} \vec{Q}^{\vec{\gamma}} \mid \vec{Q} \in \mathcal{S}_{\delta_k} \right] \mathbf{P}[\vec{Q} \in \mathcal{S}_{\delta_k}] \\ &= \sum_{k=1}^K \sum_{|\vec{\gamma}| \leq 2r} w_{\vec{\gamma}} x_k^{\vec{\gamma} \vec{0}} \end{aligned}$$

Constraints: The equality in distribution in the steady state constraint from Eq. (7.2) implies that the equality also holds for moments of all orders. So, we can relax the distributional constraint (7.2) to a constraint on moments of order no higher than $2r$ giving

$$\begin{aligned} \vec{Q} &\stackrel{d}{=} g(\vec{Q}, \vec{X}) = \vec{Q} + \vec{X} \\ \Rightarrow \mathbf{E}[\vec{Q}^{\vec{\alpha}}] &= \mathbf{E}[(\vec{Q} + \vec{X})^{\vec{\alpha}}], \quad \forall |\vec{\alpha}| \leq 2r. \end{aligned}$$

Using the theorem of total expectation, we break down the above equation in terms of the conditional expectations to get

$$\sum_{k=1}^K \mathbf{E}[\vec{Q}^{\vec{\alpha}} \mid \vec{Q} \in \mathcal{S}_{\delta_k}] \mathbf{P}[\vec{Q} \in \mathcal{S}_{\delta_k}] = \sum_{k=1}^K \mathbf{E}[g(\vec{Q}, \vec{X})^{\vec{\alpha}} \mid \vec{Q} \in \mathcal{S}_{\delta_k}] \mathbf{P}[\vec{Q} \in \mathcal{S}_{\delta_k}], \quad \forall |\vec{\alpha}| \leq 2r.$$

Note that the term $g(\vec{Q}, \vec{X})^{\vec{\alpha}}$ can be expanded using the binomial theorem as

$$g(\vec{Q}, \vec{X})^{\vec{\alpha}} = (\vec{Q} + \vec{X})^{\vec{\alpha}} = \sum_{|\vec{\gamma}_1| + |\vec{\gamma}_2| \leq \alpha} g_{\vec{\alpha}}^{(\vec{\gamma}_1, \vec{\gamma}_2)} \vec{Q}^{\vec{\gamma}_1} \vec{X}^{\vec{\gamma}_2},$$

where $\{g_{\vec{\alpha}}^{(\gamma_1, \gamma_2)}\}$ are the coefficients resulting from the expansion. The constraint can now be written as a linear function of the decision variables as follows:

$$\sum_{k=1}^K x_k^{\vec{\alpha}\vec{0}} = \sum_{k=1}^K \sum_{|\gamma_1|+|\gamma_2|\leq\alpha} g_{\vec{\alpha}}^{(\gamma_1, \gamma_2)} x_k^{\gamma_1 \gamma_2}, \quad \forall |\vec{\alpha}| \leq 2r. \quad (7.6)$$

Constraint (7.3) is also relaxed by equating the moments of the product distribution to the products of the moments

$$\begin{aligned} \psi^{\vec{\delta}_i} = \psi_{\vec{Q}}^{\vec{\delta}_i} \psi_{\vec{X}}^{\vec{\delta}_i} &\Rightarrow \mathbf{E}[\vec{Q}^{\vec{\alpha}} \vec{X}^{\vec{\beta}} | \vec{Q} \in \mathcal{S}_{\vec{\delta}_k}] = \mathbf{E}[\vec{Q}^{\vec{\alpha}} | \vec{Q} \in \mathcal{S}_{\vec{\delta}_k}] \mathbf{E}[\vec{X}^{\vec{\beta}} | \vec{Q} \in \mathcal{S}_{\vec{\delta}_k}], \\ &\quad \forall |\vec{\alpha}| + |\vec{\beta}| \leq 2r, \quad k = 1, \dots, K. \end{aligned}$$

The above equation can be used in combination with the given moments of \vec{X} from constraint (7.4) to constrain the moments of the joint conditional distribution for $|\vec{\alpha}| + |\vec{\beta}| \leq 2r$ and $k = 1, \dots, K$:

$$\begin{aligned} \mathbf{E}[\vec{Q}^{\vec{\alpha}} \vec{X}^{\vec{\beta}} | \vec{Q} \in \mathcal{S}_{\vec{\delta}_k}] \mathbf{P}[\vec{Q} \in \mathcal{S}_{\vec{\delta}_k}] &= \mathbf{E}[\vec{Q}^{\vec{\alpha}} | \vec{Q} \in \mathcal{S}_{\vec{\delta}_k}] m_{\vec{\delta}_k}^{\vec{\beta}} \mathbf{P}[\vec{Q} \in \mathcal{S}_{\vec{\delta}_k}] \\ x_k^{\vec{\alpha}\vec{\beta}} &= m_{\vec{\delta}_k}^{\vec{\beta}} x_k^{\vec{\alpha}\vec{0}} \end{aligned} \quad (7.7)$$

Constraint (7.4) can be directly expressed as:

$$\sum_{k=1}^K x_k^{\vec{0}\vec{0}} = 1. \quad (7.8)$$

Finally, we need to ensure that constraint (7.5) is satisfied and $\{x_k^{\vec{\alpha}\vec{\beta}}, |\vec{\alpha}| + |\vec{\beta}| \leq 2r\}$ represents a valid moment sequence for any $k = 1, \dots, K$. We denote the cone of moments supported on $\mathcal{S}_k = \mathcal{S}_{\vec{\delta}_k} \times \mathcal{S}_{\vec{X}}$ by

$$\mathcal{M}_{2r}(\mathcal{S}_k) = \left\{ x_k | x_k^{\vec{\alpha}\vec{\beta}} = \mathbf{E}_{\psi_k}[\vec{Q}^{\vec{\alpha}} \vec{X}^{\vec{\beta}}], \forall |\vec{\alpha}| + |\vec{\beta}| \leq 2r \text{ and for some } \psi_k \in \mathbb{M}(\mathcal{S}_k) \right\}.$$

Denoting the closure of $\mathcal{M}_{2r}(\mathcal{S}_k)$ by $\overline{\mathcal{M}_{2r}(\mathcal{S}_k)}$, constraint (7.5) can be translated to the moment constraint

$$\{x_k\} \in \overline{\mathcal{M}_{2r}(\mathcal{S}_k)}, \quad \forall k = 1, \dots, K. \quad (7.9)$$

The relaxed version of Problem 7.3.1 can then be expressed in the form of the following conic optimization problem

Problem 7.3.2.

$$\begin{aligned} & \sup / \inf_{x_k} \quad \sum_{k=1}^K \sum_{|\vec{\gamma}| \leq 2r} w_{\vec{\gamma}} x_k^{\vec{\gamma}\vec{0}} \\ & \text{s.t.} \\ & \quad \sum_{k=1}^K x_k^{\vec{\alpha}\vec{0}} = \sum_{k=1}^K \sum_{|\vec{\gamma}_1| + |\vec{\gamma}_2| \leq |\vec{\alpha}|} g_{\vec{\alpha}}^{(\vec{\gamma}_1, \vec{\gamma}_2)} x_k^{\vec{\gamma}_1 \vec{\gamma}_2}, \quad \forall |\vec{\alpha}| \leq 2r \\ & \quad x_k^{\vec{\alpha}\vec{\beta}} = m_{\delta_k}^{\vec{\beta}} x_k^{\vec{\alpha}\vec{0}}, \quad \forall |\vec{\alpha}| + |\vec{\beta}| \leq 2r, \quad k = 1, \dots, K \\ & \quad \sum_{k=1}^K x_k^{\vec{0}\vec{0}} = 1 \\ & \quad \{x_k\} \in \overline{\mathcal{M}_{2r}(\mathcal{S}_k)}, \quad \forall k = 1, \dots, K. \end{aligned}$$

7.3.2 A Semidefinite Relaxation

The moment cone can in turn be characterized using positive semidefinite matrices as in [5, 56, 96]. A necessary condition for constraint (7.9) to hold is that the moment matrix associated with $\{x_k\}$ be positive semidefinite [56]. The moment matrix corresponding to the sequence $y = \{y^{\vec{\alpha}\vec{\beta}}, |\vec{\alpha}| + |\vec{\beta}| \leq 2r\}$ is denoted $\mathbf{M}_r(y)$, and is given by the block matrix $\{M_r^{i,j}(y), 0 \leq i, j \leq r\}$ with rows and columns indexed in the basis of polynomials of degree less than or

equal to r . The entries of the moment matrix satisfy the following condition:

$$\text{if } M_r^{1,j}(y) = y^{\vec{\alpha}_1 \vec{\beta}_1} \text{ and } M_r^{i,1}(y) = y^{\vec{\alpha}_2 \vec{\beta}_2}, \text{ then } M_r^{i,j}(y) = y^{(\vec{\alpha}_1 + \vec{\alpha}_2)(\vec{\beta}_1 + \vec{\beta}_2)}.$$

In our formulation, this condition has to be satisfied by the decision variables $\{x_k\}$, i.e.,

$$\mathbf{M}_r(x_k) \succeq 0, \quad k = 1, \dots, 2^N. \quad (7.10)$$

Additionally, note that $\mathcal{S}_{\vec{\delta}_k}$ can be specified by an intersection of linear inequalities and so can \mathcal{S}_k , i.e.,

$$\mathcal{S}_k = \cap_{h \in H_k} \{h(\vec{q}, \vec{x}) \geq 0\},$$

where H_k denotes a set of linear functions defining \mathcal{S}_k . Clearly $\{x_k\}$ should also be a valid truncated moment sequence when restricted to each half plane $\{h(\vec{q}, \vec{x}) \geq 0\}$. Again it can be shown that a necessary condition for this to be the case is that an associated (localizing) moment matrix, denoted $\mathbf{M}_{r-1}(h, x_k)$, depending on the coefficients of the hyperplane h and x_k be positive semidefinite, see [5, 56]. Let the coefficients of h be denoted by $\{h^{\vec{\alpha}}, |\vec{\alpha}| \leq 1\}$. Consider any region \mathcal{S}_k , a localizing matrix $\mathbf{M}_{r-1}(h, x_k)$ associated with one of the polynomials $h \in H_k$ is defined as:

$$M_{r-1}^{(i,j)}(h, x_k) = \sum_{|\vec{\alpha}| \leq 1} h^{\vec{\alpha}} x_k^{\theta(i,j) + \vec{\alpha}},$$

where $\theta(i, j)$ is the subscript of the entry $M_{r-1}^{i,j}(x_k)$ in matrix $\mathbf{M}_{r-1}(x_k)$. The corresponding set constraints for $\{x_k\}$ to be a valid truncated moment sequence is then given by

$$\mathbf{M}_{r-1}(h, x_k) \succeq 0, \quad \forall h \in H_k, k = 1, \dots, 2^N. \quad (7.11)$$

Substituting our semidefinite constraints into Problem (7.8) we obtain :

Problem 7.3.3.

$$\begin{aligned}
& \sup / \inf_{x_k} \sum_{k=1}^K \sum_{|\vec{\gamma}| \leq 2r} w_{\vec{\gamma}} x_k^{\vec{\gamma}\vec{0}} \\
& \text{s.t.} \\
& \sum_{k=1}^K x_k^{\vec{\alpha}\vec{0}} = \sum_{k=1}^K \sum_{|\vec{\gamma}_1| + |\vec{\gamma}_2| \leq \alpha} g_{\vec{\alpha}}^{(\vec{\gamma}_1, \vec{\gamma}_2)} x_k^{\vec{\gamma}_1 \vec{\gamma}_2}, \quad \forall |\vec{\alpha}| \leq 2r \\
& x_k^{\vec{\alpha}\vec{\beta}} = m_{\delta_k}^{\vec{\beta}} x_k^{\vec{\alpha}\vec{0}}, \quad \forall |\vec{\alpha}| + |\vec{\beta}| \leq 2r, \quad k = 1, \dots, K \\
& \sum_{k=1}^K x_k^{\vec{0}\vec{0}} = 1 \\
& \mathbf{M}_r(x_k) \succeq 0, \quad \forall k = 1, \dots, K \\
& \mathbf{M}_{r-1}(h, x_k) \succeq 0, \quad \forall h \in H_k, \quad k = 1, \dots, K
\end{aligned}$$

This semidefinite problem can be solved to obtain the desired upper and lower bounds on the objective function. As r is increased, and information about more moments of \vec{X} are used in the semidefinite program, tighter bounds can be obtained at the cost of increased complexity of the optimization problem. In all the computational results presented in the sequel, Gloptipoly [40] and Sedumi [81] were the tools used to solve the semidefinite program.

7.4 Fidelity of the Bounds

In this section, two examples of queuing systems with coupled processors are examined. Bounds on mean delay are derived by using Little's law in conjunction with the bounds on sum queue length obtained from the semidef-

inite optimization. These bounds are then compared to the simulated delay performance.

Two queue system: A system of two queues with users arriving according to a Poisson process with rate λ . The service time of the users is assumed to be exponentially distributed with mean 1. The service rates of the queues are summarized in Table 7.1. Fig. 7.2 depicts the bounds on the mean

Queue status		Service rates	
Queue 1	Queue 2	Queue 1	Queue 2
Idle	Idle	0	0
Idle	Busy	0	4
Busy	Idle	4	0
Busy	Busy	4	2.5

Figure 7.1: Service Rates in the two queue system.

delay obtained by using progressively higher order relaxations, along with the simulated mean delay. The bounds obtained from assuming that the service rates correspond to the queues being *always* busy/idle are also plotted for comparison. Even when $r = 1$, the bounds computed through the optimization proposed above are close to the actual steady state mean delay. Further, as the relaxation order is increased, the upper and lower bounds further improve, matching the simulated results very well. The discrepancy between the upper bound that assumes the system is saturated and the upper bounds obtained by solving the SDP clearly indicates the impact that coupling has even under heavy loads. In this case, closed form expressions for the steady state delay are known, however, the degree of accuracy of our results is highly encouraging.


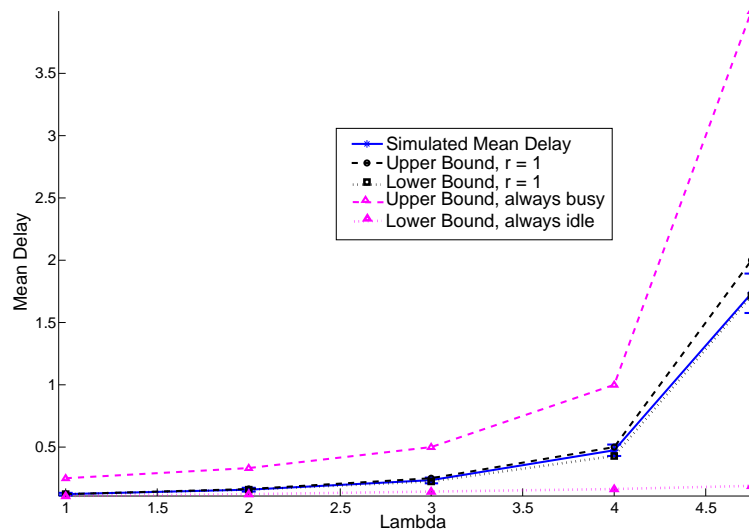


Figure 7.2: Mean delay in the two queue system.

Three queue system: For this queuing system, no closed form expressions or tight bounds for the steady state delay are available to the best of our knowledge. The state-dependent service rates are summarized in Table 7.3a. The mean rate of arrival at each base station is again λ and the service times are exponentially distributed with mean 1. As shown in Fig. 7.3b, the upper and lower bounds are not far apart, and closely trail the simulated mean delay, and in particular capture the trends with increasing load. The bounds assuming the queues are saturated are much worse, demonstrating that coupling plays a significant role in the performance experienced by the users in the system.

Queue status			Service rates		
Queue 1	Queue 2	Queue 3	Queue 1	Queue2	Queue3
Idle	Idle	Idle	0	0	0
Idle	Idle	Busy	0	0	10
Idle	Busy	Idle	0	10	0
Idle	Busy	Busy	0	7	7
Busy	Idle	Idle	10	0	0
Busy	Idle	Busy	7	0	7
Busy	Busy	Idle	7	7	0
Busy	Busy	Busy	5	5	5

(a) Service Rates



(b) Mean delay experienced by users.

Figure 7.3: Three Queue system

7.5 Some Other Instances of Coupled Queuing Systems

In the following chapters, the bounding methodology developed above will be used to study and optimize the user association policy in a wireless network. Here, we consider a few other scenarios where the dynamics of

the system are properly modeled by the coupled processors are examined to demonstrate the applicability of the bounding methodology.

7.5.1 A Generalized Processor Sharing System

Generalized Processor Sharing (GPS) [65] is a service discipline that was developed to allow capacity to be shared in a fair and flexible manner, and is an important mechanism for achieving differentiated quality of service. Unlike a strict priority scheme, the GPS discipline allows for service differentiation while preventing starvation. The following is a brief description of the GPS scheduler. Consider N classes of customers, each with some associated arrival rate and service requirements, sharing a server. Each class, n , also has an associated weight ϕ_n , with $\sum_{n=1}^N \phi_n = 1$ and can be served at rate R_n . If all the queues are non-empty, the fraction of time queue n is served is given by ϕ_n and the corresponding service rate is $\phi_n R_n$. However, if some of the queues are empty, their allotted time-fractions are distributed among the non-empty queues in proportion to their respective weights. Thus, the queues are coupled by the mechanism used to share excess capacity. While the GPS policy is not itself realistic to implement, disciplines that closely track the GPS mechanism such as Weighted Fair queuing (WFQ) [29] are used in practice. Understanding the performance of the system, and of the user classes is very important, for example, in order to optimize the choice of the class weights.

Two Queue System: Consider a server with a service rate of 10 that is shared by two classes of packets with equal weights, i.e., $R_1 = R_2 = 10$

and $\phi_1 = \phi_2$. packets arrive at each class as a Poisson process of rate λ , with exponentially distributed service requirements with mean 1. In this case, since the sum service rate is a constant, independent of the system state, the sum queue length process is a Markov process, and the distribution of the sum queue length (but not the individual queue lengths) can be analytically determined as the sum queue length reduces to a M/M/1 queuing model. The bounds computed by solving the SDP with $r = 1, 2$ are plotted in Fig. 7.4, along with the simulated results. The lower and upper bounds match exactly, for all the values of λ . In addition, they precisely match the analytically computed mean delay (not shown). This is confirmed by the simulated mean delay matching the bound.

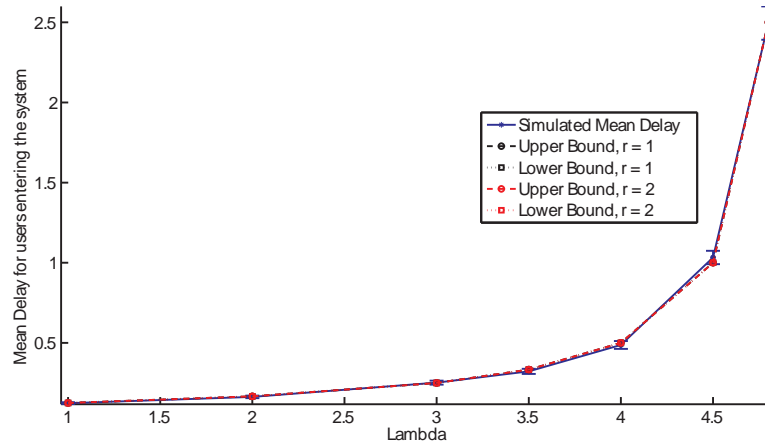


Figure 7.4: Mean overall system queue length for a 2 Queue GPS system.

Fig. 7.5 exhibits the mean queue length at the queue associated with one of the classes as λ increases. While the system as a whole is work-conserving, the individual queues are not. The bounds computed by solving the SDP

with $r = 1, 2$ are plotted along with the bounds computed assuming that the neighboring queue is always busy/idle. Note that, in this case closed-form expressions for the mean delay are known [31, 36]. However, for larger systems, analytic expressions for the mean delay are not known.

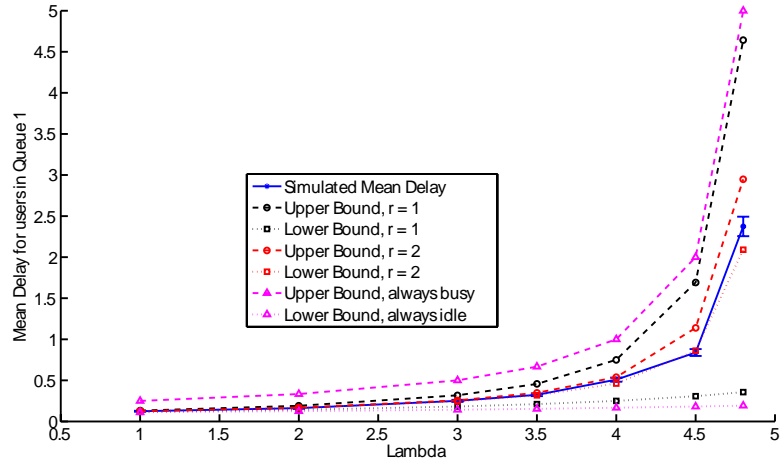


Figure 7.5: Mean queue length at one of the queues in the two queue GPS system.

Three Queue System: Consider the case now, where the GPS server is shared by three classes (queues) of packets with weights $\phi_1 = 3, \phi_2 = 1, \phi_3 = 1$. packets arrive at each class as a Poisson process of rate λ , with exponentially distributed service requirements with mean 1. The sum queue length process is again a Markov process, and the mean sum queue length is plotted in Fig 7.6, along with the computed bounds.

Figs. 7.7 and 7.8 exhibit the mean queue lengths at Queue 1 and at Queue 2 respectively. The mean delay observed at Queue 3 is identical to that observed at Queue 2. In this case, note that trying to bound the delay




Figure 7.6: Mean system queue length

at Queues 2 and 3, assuming that the other queues are saturated and Queue 1 is always busy leads to a predicted mean delay of ∞ . This is because, the system is only stable because Queues 2 and 3 can borrow the excess capacity of the server when queue 1 is idle. The simulations demonstrate that the SDP formulation provides reasonably tight bounds on the mean queue lengths of the individual queues when $r \geq 2$.

Choosing Queue Weights: The choice of class weights is a crucial factor affecting the performance experienced by packets in the system. In systems, such as wireless networks, where the maximum rates at which different classes can be served varies, a systematic method to choose the class weights is not apparent. Note that, in this case, even the sum queue length process is not a Markov process, as the system is not work conserving. The close bounds provided by our formulation can be exploited to determine the weights that should be associated with each class if the objective function is a weighted

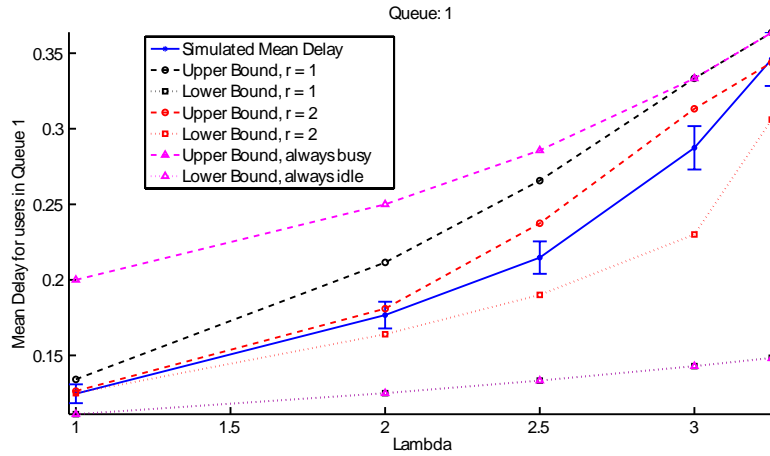


Figure 7.7: Mean queue length at queue 1

combination of the queue length moments.

Consider a system with two classes, with $R_1 = 10$, $R_2 = 5$, where the arrival rate to each class is equal to λ . We consider an objective function of the form $\mathbf{E}[Q_1^z + Q_2^z]$, and use the lower bound from the SDP formulation to approximate the value of the objective function for various class weights. The weight associated with class 1 is assumed to be $\phi_1 = \phi$, the weight associated with class 2 is then $\phi_2 = 1 - \phi$. We do a simple line search to determine the weight to be associated with class 1 to minimize the objective function. As z is increased, the penalty associated with high queue lengths increases rapidly, and the objective function tends to balance the performance experienced by the queues.

Fig. 7.9 exhibits the mean queue lengths of the two queues against the arrival rate λ , for $z = 1, 2, 4$, and Fig. 7.10 plots the weight determined using




Figure 7.8: Mean queue length at queue 2




Figure 7.9: Mean queue lengths resulting from the different objective functions.

the SDP formulation. When $z = 1$, the objective is to minimize the system queue length. The optimal policy here is to always serve queue 1, whenever both queues are busy and $\phi = 1$ irrespective of the load, as shown in Fig. 7.10. However, this leads to very large queue lengths at queue 2. Increasing the value of z leads to larger weights being associated with queue 2, in order to

balance the queue lengths at the two queues. Fig. 7.9 clearly shows the impact of the larger weights that are picked using the SDP formulation on the mean queue lengths of the two queues.




Figure 7.10: Optimal weights for the different objective functions

7.5.2 A Multiple Access System

The final example is an ALOHA-like [2, 72] slotted multiple access system, where N users contend for access to a shared channel in order to communicate with a central server. A collision model is used to capture the effect of simultaneous transmissions, i.e., if a user is the only one contending for the channel in a slot, the transmission is successful and a 1Kb portion of the file is transmitted to the sever. If more than one user transmits at the same time, there is a collision and all transmissions are unsuccessful. Files with exponentially distributed sizes and mean 2Mb arrive for transmission at each user as a Poisson process with mean arrival rate λ . All users with files waiting for

transmission contend at a time slot with probability P_C . The aim is to find a contention probability that results in low file transfer delays and high user throughputs.

Clearly, the contention probability has to be chosen in a load-dependent manner. Intuitively, when the system load is low, the number of users with non-empty queues at any given time will be low and those users should contend with higher probability. However, at high loads, when a larger number of users are likely to have non-empty queues users should contend with a lower probability in order to avoid collisions. If all N users are saturated, the contention probability resulting in the highest throughput is $P_C = 1/N$. This can be seen easily by differentiating the probability of there being exactly one contending user,

$$\mathbf{P}(\text{exactly one contending user}) = \mathbf{P}(\text{successful transmission}) = P_C(1-P_C)^{N-1},$$

with respect to the contention probability P_C and equating to zero. However, in the dynamic system, the number of busy users varies as files arrive and are transmitted. In this case, it is not clear how the contention probability should be chosen.

This system is modeled using a fluid coupled processors model, such that each user transmits data at a rate of $1000P_C(1 - P_C)^{n-1}$ when n of the N users have data to send. Using the lower bound from the semidefinite formulation as an approximation, the contention probability that minimizes the mean sum system queue length is determined. A system with three users is

considered, and as before use a simple line search is used to find the common contention probability that minimizes the mean sum queue length of the system.




Figure 7.11: Optimal probability of contention.

Fig. 7.11 shows how the contention probability determined using the SDP formulation varies with the file arrival rate λ . As expected, the contention probability is high when load is low and collisions are rare, and decreases with increasing loads in order to avoid collisions. At very high loads, the contention probability approaches $\frac{1}{3}$, the optimal contention probability in the saturated system where all users have backlogged queues.

Figs. 7.12 and 7.13 exhibit the mean file transfer delay and mean user throughput for the system using the contention probability determined using our SDP formulation and the system with a contention probability of $\frac{1}{3}$. At very high loads, performance of the two schemes are close to identical. This is to be expected as the probability of contention chosen using the semidefinite




Figure 7.12: Mean delay comparison.




Figure 7.13: Mean throughput comparison.

formulation is very close to $\frac{1}{3}$. At lower loads, the mean file transfer delay and the mean user throughputs are much improved when using the contention probability from the SDP formulation. This is particularly clear from the plot of the mean user throughputs against the user arrival rate. Thus, we can clearly

see that the coupled processor model accurately reflects the dynamics of the system, and shows the potential of the semidefinite formulation in modeling and optimizing such systems.

In the following chapter, the semidefinite programming based methodology developed above is used to optimize user association policies in wireless networks where users see location dependent service rates that are further coupled across base stations through interference. The results further demonstrate the significant impact that coupling can have on user performance and the accuracy of the proposed method.

Chapter 8

Static User Association Policies

User association policies are classified into static and dynamic policies. *Dynamic* policies use information about the current loads being served at the candidate base stations when deciding the base station to which a new user is assigned. A *static* user association policy is one that does not take into account the current state of the system when making this decision. In this chapter, the focus is on understanding the performance of, and optimizing static user association policies. The methodology developed in Chapter 7 will be crucial to optimizing the static user association policy.

8.1 System Model

In Secs. 8.3-8.4, the scenario with two base stations, BS1 and BS2, located a distance d apart on a line, as shown in Fig. 6.1a is considered. User requests are distributed on the line segment joining the two base stations. A user request is identified by the distance between the user and BS1, denoted by $x \in [0, d]$. The distance between the user and BS2 is then given by $d - x$. User requests arrive according to a spatial Poisson process with mean measure $\lambda(\cdot)$ which is absolutely continuous with respect to the Lebesgue measure, i.e.,

the rate at which user requests arrive into a set \mathcal{X} is $\lambda(\mathcal{X})$. Each user request is assumed to correspond to a downlink file transfer which is assumed to be exponentially distributed with mean 1, and the position of the user remains fixed for the duration of the transfer. Once the file transfer is completed, the user leaves the system.

The capacity to users from their serving base station depends on the received signal strength and the strength of the received interference, and is assumed to be monotonically increasing in the perceived signal to interference plus noise ratio (SINR). The base stations transmit, and thereby cause interference only when they are serving users. The base stations use the processor sharing mechanism to serve active users, i.e., the base station splits time evenly among all users currently being served. Thus, a degree of temporal ‘fairness’ is imposed.

A static load allocation policy π partitions the line segment into regions \mathcal{X}_1^π and \mathcal{X}_2^π , served by BS1 and BS2 respectively. The base station that serves a user at location x under policy π is denoted by $\beta^\pi(x)$. Thus, if $x \in \mathcal{X}_1^\pi$ then $\beta^\pi(x) = 1$, otherwise $\beta^\pi(x) = 2$. Base stations transmit at maximum power when there are active associated users, and turn off otherwise. The signal strengths received by a user at location x from BS1 and BS2 are denoted by $s_1(x)$ and $s_2(x)$ respectively. For $i = 1, 2$, the worst and best received signals in $A \subset [0, d]$ are denoted by $\underline{s}_i(A) = \inf_{x \in A} s_i(x)$ and $\overline{s}_i(A) = \sup_{x \in A} s_i(x)$. Let N_0 denote the average power of the additive Gaussian noise.

Under a given policy π , let $\mathbf{U}^\pi(t) = (\mathcal{U}_1^\pi(t), \mathcal{U}_2^\pi(t))$ where $\mathcal{U}_i^\pi(t)$ is the

set of locations for users being served at base stations $i = 1, 2$ at time t . Note that since $\lambda(\cdot)$ is non-atomic, users' locations will be distinct with probability 1. Given the assumptions on arrivals and file sizes, $\mathbf{U}^\pi(t)$ is a Markov process since, given all the users locations, one can determine their service capacities and thus departure rates. Note however that its state space is uncountable. By contrast, the process $\mathbf{Q}^\pi(t) = (Q_1^\pi(t), Q_2^\pi(t))$ where $Q_i^\pi(t) = |\mathcal{U}_i^\pi(t)|$ for $i = 1, 2$ is on a countable state space, but not Markovian.

This model is similar to that of optimally routing n classes of users to m non-identical queues studied in [19], with an infinite number of classes. However, in this case the problem is further complicated by the fact that the queues at the base stations are coupled (through interference) and the system is non-work conserving.

8.2 Simulation Model

In the bulk of the simulation results, we consider two base stations located 500m apart with users arriving according to a Poisson process. The three base station network studied in Sec. 8.6.3 consists of three facing sectors in a hexagonal layout of base stations with cell radius 250m, with users again arriving according to a Poisson process. In Secs. 8.3-8.6, where I develop and study the semidefinite programming based methodology, I assume that the user distribution is spatially homogeneous. In the two base station case, users are assumed to be distributed uniformly on the line joining the two base stations, and in the three base station network, users are assumed to

be distributed uniformly within the hexagon formed by the three interfering sectors. Non-homogeneous spatial load distributions are considered in the simulation results presented in Sec. 9. and the exact load profiles simulated are described in Sec. 9.3.

A carrier frequency of 1GHz, and a bandwidth of 10MHz are assumed. The maximum transmit power is restricted to 10W. Additive white Gaussian noise with power -55dBm is assumed. A log distance path loss model[69] is considered, with path loss exponent 2. Shadowing, and fading are not considered in these results. File sizes are assumed to be exponentially distributed, with mean 5MB. The data rate at which users are served is calculated based on the perceived SINR using Shannon’s capacity formula. The maximum rate at which a user can be served is capped at 54 Mbps. The base stations transmit at maximum power when they have active users, share capacity across users using a processor sharing mechanism, and turn off otherwise. The mean user perceived delay is estimated within a relative error of 2%, at a confidence level of 95%. Note that the sensitivity of the delay performance to the channel and system model is examined in Sec. 9.3.3 where a system with a higher path loss exponent, and cell-edge SNR of 10 dB is simulated.

8.3 Optimal Static Policies

I begin by considering static association policies in the one dimensional, two base station system. Such policies are defined by the service regions corresponding to each base station, which in turn may depend on the long term

offered load $\lambda(\cdot)$. The key result is that under our system model, the service regions are contiguous and thus are defined by a single threshold between the two base stations. The following lemma provides a partial characterization of optimal static policies. Note at the outset that, while this result appears straightforward, the challenge lies in the dynamic nature of the model; specifically, in dealing with the spatial arrivals and departures, the dynamic (on/off) nature of the interference from the neighboring base station, and thus the coupling of delay performance between the two base stations.

Lemma 8.3.1. *Consider the two base station model defined in Sec. 8.1. For any static load allocation policy π_a with $\mathcal{R}_1 \subseteq \mathcal{X}_1^{\pi_a}$, $\mathcal{R}_2 \subseteq \mathcal{X}_2^{\pi_a}$ with $\lambda(\mathcal{R}_1) = \lambda(\mathcal{R}_2)$, and such that $\underline{s}_1(\mathcal{R}_2) \geq \overline{s}_1(\mathcal{R}_1)$ and $\overline{s}_2(\mathcal{R}_2) \leq \underline{s}_2(\mathcal{R}_1)$, the policy π_b with $\mathcal{X}_1^{\pi_b} = (\mathcal{X}_1^{\pi_a} \cup \mathcal{R}_2) \setminus \mathcal{R}_1$, $\mathcal{X}_2^{\pi_b} = (\mathcal{X}_2^{\pi_a} \cup \mathcal{R}_1) \setminus \mathcal{R}_2$ achieves lower (or equal) average user delay.*

The insight underlying this lemma can be grasped by considering Fig. 8.1. It illustrates a policy π_a which satisfies the lemma's conditions if signal strength decays monotonically with distance from the serving base station – although part of our system model, this is not required to prove the lemma. Policy π_b is constructed by merely exchanging service regions $\mathcal{R}_1, \mathcal{R}_2$ between the two base stations. The constraints on the best and worst case signal strengths ensure that this exchange is favorable for both base stations at all the associated user locations, which implies the following straightforward fact.




Figure 8.1: A sub-optimal load allocation policy.

Fact 8.3.1. *Under the assumptions on \mathcal{R}_1 and \mathcal{R}_2 in Lemma 8.3.1, and the assumption that capacity is monotonically increasing in SINR, the capacity from BS1 to any user in \mathcal{R}_2 is greater than that to any user in \mathcal{R}_1 under the same interference regime, i.e., BS2 is transmitting or not. Similarly, the capacity from BS2 to any user in \mathcal{R}_1 is greater than that to any user in \mathcal{R}_2 , whether BS1 is transmitting or not.*

So, the exchange leaves the intensity of arrivals to BS1 and BS2 unchanged, and associates users to them which then can be served at higher capacity under the same interference regime. This allows us to construct a spatial coupling (i.e., by associating users in different regions) for networks under the two policies, showing that the average queue lengths are not increased.

The following definitions provide a characterization of the stochastic ordering relationship between two process, and will be used in the proof of Lemma 8.3.1.

Definition 8.3.1 ([63]). *Let $\mathbf{l}, \mathbf{m} \in \mathbb{R}^n$, and let $l_{[1]} \geq \dots \geq l_{[n]}$ denote the components of \mathbf{l} arranged in descending order.*

$$\mathbf{l} \prec_w \mathbf{m} \text{ if } \sum_{i=1}^k l_{[i]} \leq \sum_{i=1}^k m_{[i]}, \quad k = 1, \dots, n$$

The vector \mathbf{l} is then said to be weakly majorized by \mathbf{m} .

Definition 8.3.2 ([79]). Let \mathbf{L}, \mathbf{M} be random vectors taking values in \mathbb{R}^n . \mathbf{L} is stochastically weak-majorized by \mathbf{M} , written $\mathbf{L} \prec_w^{st} \mathbf{M}$, if there exist random vectors $\tilde{\mathbf{L}}$ and $\tilde{\mathbf{M}}$ taking values in \mathbb{R}^n with the same probability laws as \mathbf{L} and \mathbf{M} respectively, with $\tilde{\mathbf{L}} \prec_w \tilde{\mathbf{M}}$ a.s.

Proof of Lemma 8.3.1. The proof will demonstrate that the policy π_b , which is obtained from π_a by exchanging service regions \mathcal{R}_1 and \mathcal{R}_2 between the base stations, obtains a lower (or equal) mean delay, see Section 8.3. This is shown by constructing a pair of coupled processes $\tilde{\mathbf{U}}^{\pi_a}(t)$ and $\tilde{\mathbf{U}}^{\pi_b}(t)$, such that

$$\tilde{\mathcal{U}}_1^{\pi_b}(t) \subseteq \tilde{\mathcal{U}}_1^{\pi_a}(t) \text{ and } \tilde{\mathcal{U}}_2^{\pi_b}(t) \subseteq \tilde{\mathcal{U}}_2^{\pi_a}(t), \quad (8.1)$$

and such that $\tilde{\mathbf{U}}^{\pi_a}(t) \sim \mathbf{U}^{\pi_a}(t)$ and $\tilde{\mathbf{U}}^{\pi_b}(t) \sim \mathbf{U}^{\pi_b}(t)$. It follows that associated queue length processes $\tilde{\mathbf{Q}}^{\pi_a}(\mathbf{t})$ and $\tilde{\mathbf{Q}}^{\pi_b}(\mathbf{t})$ satisfy similar properties with containment replaced with an inequality. By standard arguments, see [79], this construction suffices to show that $\mathbf{Q}^{\pi_b}(t)$ is *stochastically weak-majorized* by $\mathbf{Q}^{\pi_a}(t)$. As $t \rightarrow \infty$ this implies π_b achieves a lower (or equal) mean queue length, and thus, by Little's Law, a lower (or equal) mean delay.

Note that the arrival rates associated with the exchanged service regions are equal so the arrival rate to each base station under the two policies are the same, i.e., $\lambda_1 = \lambda(\mathcal{X}_1^{\pi_a}) = \lambda(\mathcal{X}_1^{\pi_b})$ and $\lambda_2 = \lambda(\mathcal{X}_2^{\pi_a}) = \lambda(\mathcal{X}_2^{\pi_b})$. Arrivals of the two processes $\tilde{\mathbf{U}}^{\pi_a}(t)$ and $\tilde{\mathbf{U}}^{\pi_b}(t)$ are coupled, as generated by a common Poisson process with intensity $\lambda_1 + \lambda_2$. For convenience, user requests are

indexed based on arrival times (including those in the system at $t = 0$), i.e., $1, 2, \dots$. While arrival times for users to the two systems are identical, their locations may not be, whence $x_i^{\pi_a}$ and $x_i^{\pi_b}$ denotes the locations of the i^{th} request under policy π_a and π_b respectively.

Suppose $x \in \tilde{\mathcal{U}}_1^{\pi_a}(t)$ then let $c_x^{\pi_a}(t)$ be the capacity to the user under policy π_a at time t taking into account the state of the neighboring base station. Since users share capacity via processor sharing, effective service rate to users at locations x and y under the two policies is given by $\mu^{\pi_a}(t, x) = \frac{c_x^{\pi_a}(t)}{Q_{\beta^{\pi_a}(x)}^{\pi_a}(t)}$ and $\mu^{\pi_b}(t, y) = \frac{c_y^{\pi_b}(t)}{Q_{\beta^{\pi_b}(y)}^{\pi_b}(t)}$. So the departure rate of users from BS1 under policy π_a is given by

$$\mu_1^{\pi_a}(t) = \sum_{x \in \tilde{\mathcal{U}}_1^{\pi_a}(t)} \mu^{\pi_a}(t, x).$$

The overall departure rates $\mu_2^{\pi_a}(t)$, $\mu_1^{\pi_b}(t)$, and $\mu_2^{\pi_b}(t)$ are defined analogously.

Let $\tilde{\mathbf{U}}^{\pi_a}(0) = \tilde{\mathbf{U}}^{\pi_b}(0)$ so (8.1) holds at time $t = 0$. The construction will be such that if (8.1) holds at some time t then it is satisfied after the next arrival/departure, while maintaining marginal dynamics that are consistent with systems associated with policies π_a and π_b . Although the two systems see the same overall arrival rates they may see different overall departure rates. In our construction we let

$$\nu(t) = \lambda_1 + \lambda_2 + \max(\mu_1^{\pi_a}(t), \mu_1^{\pi_b}(t)) + \max(\mu_2^{\pi_a}(t), \mu_2^{\pi_b}(t))$$

denote the current rate of events for the *coupled processes* and allow fictitious events to ensure the marginal system processes have the correct dynamics. Let the time at which the next event occurs be t' and z be a realization of

a random variable Z , which is uniformly distributed on $[0, \nu(t)]$. The coupled process events are constructed as follows:




Figure 8.2: Example coupling construction for arrivals/departures based on realization of Z .

Arrivals: If $0 \leq z \leq \lambda_1$, the next event is an arrival, say of user n , to BS1 under both policies. Let random variables $X_n^{\pi_a}$ and $X_n^{\pi_b}$ denote the position of this user under policies π_a and π_b respectively. The distribution $X_n^{\pi_a}$ is given by $\mathbf{P}(X_n^{\pi_a} \in A) = \frac{\lambda(A)}{\lambda_1}$, for a measurable set $A \subseteq \mathcal{X}^{\pi_a}$. The position of the user under policy π_b is identical, except if $X_n^{\pi_a} \in \mathcal{R}_1$. In this case, the user's location falls within \mathcal{R}_2 with a distribution $\mathbf{P}(X_n^{\pi_b} \in B | X_n^{\pi_a} \in \mathcal{R}_1) = \frac{\lambda(B)}{\lambda(\mathcal{R}_2)}$, where $B \subseteq \mathcal{R}_2$. The states of the processes are updated accordingly. If $\lambda_1 \leq z \leq \lambda_1 + \lambda_2$, the next event is an arrival to BS2 under both policies, with the user's location generated analogously to the above. In either case, arrivals to BS1 or BS2 occurs simultaneously for both policies, so (8.1) holds at time t' . Also under the above construction the spatial distribution of Poisson arrivals is maintained.

Departures: If $\lambda_1 + \lambda_2 \leq z \leq \lambda_1 + \lambda_2 + \max(\mu_1^{\pi_a}(t), \mu_1^{\pi_b}(t))$, the event is a potential departure from BS1. Consider any user k such that $x_k^{\pi_b} \in \tilde{\mathcal{U}}_1^{\pi_b}(t)$. Since (8.1) holds, user k is also in the system under policy π_a , i.e., $x_k^{\pi_a} \in \tilde{\mathcal{U}}_1^{\pi_a}(t)$. Since (8.1) holds there are only three cases to consider:

1. $\tilde{\mathcal{U}}_2^{\pi_b}(t) = \tilde{\mathcal{U}}_2^{\pi_a}(t) = \emptyset$: BS2 is idle under both policies. If $x_k^{\pi_a} = x_k^{\pi_b}$, $c_{x_k}^{\pi_b}(t) = c_{x_k}^{\pi_a}(t)$. Otherwise, $x_k^{\pi_a} \in R_1$ and $x_k^{\pi_b} \in R_2$, so Fact 8.3.1 implies $c_{x_k}^{\pi_b}(t) \geq c_{x_k}^{\pi_a}(t)$.
2. $\tilde{\mathcal{U}}_2^{\pi_b}(t) \neq \emptyset, \tilde{\mathcal{U}}_2^{\pi_a}(t) \neq \emptyset$: BS2 is transmitting under both policies, and, as in the previous case, we can argue that $c_{x_k}^{\pi_b}(t) \geq c_{x_k}^{\pi_a}(t)$.
3. $\tilde{\mathcal{U}}_2^{\pi_b}(t) = \emptyset, \tilde{\mathcal{U}}_2^{\pi_a}(t) \neq \emptyset$: In this case, users in BS1 see no interference under policy π_b while they see interference from BS2 under policy π_a . Combining our conclusion in case 1 with the fact that the data rate at which users can be served is an increasing function of the received signal to interference plus noise ratio, we see that $c_{x_k}^{\pi_b}(t) \geq c_{x_k}^{\pi_a}(t)$.

Also, by assumption $\tilde{Q}_1^{\pi_b}(t) \leq \tilde{Q}_1^{\pi_a}(t)$, thus $\mu^{\pi_b}(t, x_k^{\pi_b}) \geq \mu^{\pi_a}(t, x_k^{\pi_a})$. This permits us to couple User k 's departure such that if it leaves under policy π_a , it also leaves under policy π_b . To see this, consider Fig. 8.2 where $[0, \nu(t)]$ has been subdivided based on the arrival rates and service rates of the users in the system under the two policies. If a user is present in both systems then a set of length $\mu^{\pi_a}(t, x_k^{\pi_a})$ for policy π_a is contained within one of length $\mu^{\pi_b}(t, x_k^{\pi_b})$ for policy π_b . If the user has already left the system under policy π_a , the corresponding set for policy π_b can be arranged arbitrarily (need not be contiguous) within $[0, \nu(t)]$. Unused intervals correspond to dummy events. Which departures (if any) occur for the two systems depend on which sets contain z . However, clearly a departure of User k from BS1 under policy π_a results in the same under policy π_b unless it has already left the system, and (8.1) still hold at time t' . If $(\lambda_1 + \lambda_2 + \max(\mu_1^{\pi_a}(t), \mu_1^{\pi_b}(t))) \leq z$, the event is a

potential departure from BS2, and is treated analogously to departures from BS1.

Since relationship (8.1) holds after any future event, by induction the relationship holds for all times in the future. We show that the following relationship holds at any given time

1. $\tilde{Q}_1^P(t) \geq \tilde{Q}_1^{P_E}(t)$ and $\tilde{Q}_2^P(t) \geq \tilde{Q}_2^{P_E}(t)$
2. Corresponding to every user attached to BS1 under policy P_E , there exists a user attached to BS1 under policy P , that is served at lower rates both when BS2 is idle and active.
3. Corresponding to every user attached to BS2 under policy P_E , there exists a user attached to BS2 under policy P , that is served at lower rates both when BS1 is idle and active.

Now, consider a sequence of user arrivals and departures resulting from a static load allocation policy that associates users in region r_1 with base station 1, and users in region r_2 with base station 2. An alternate sample path is constructed based on this sequence of user arrivals. User arrivals in region r_1 are moved to instead arrive in region r_2 while still being served by base station 1, and user arrivals in region r_2 are moved to r_1 and served by base station 2. All other user arrivals are unchanged. Since the probability of user arrivals in region r_1 is equal to the probability of user arrivals in r_2 , and there are no correlations between user arrivals, the constructed sequence of user arrivals is also representative of the arrival process.

The user queues associated with the two base stations evolve identically until a user arrives in either region r_1 or r_2 . As proved previously, this user is served at a higher rate in the alternate sample path. All other users currently in the queues are served at exactly the same rates as in the original sample path as the queue lengths of the two queues are identical. As a result the user that arrived to one of the regions r_1 or r_2 is served and leaves the system earlier in the alternate sample path. This results in all the other remaining users in the system being served at the same or greater rate. Thus, all users in the alternate sample path perceive delays that are less than or equal to those in the original sample path.

Since both systems are ergodic, the sample mean of the user delays in both sample paths converge eventually to the expected values, and the expected user delay in the alternate sample path has to be lower or equal to that in the original. Note that this alternate sample path corresponds to a policy which associates users in region r_2 with base station 1, and users in r_1 to base station 2. \square

The argument can also be extended to other service disciplines, e.g., FCFS and LCFS.

Theorem 8.3.1. *For the two base station model defined in Sec. 7.2, there exists a static load allocation policy minimizing mean delay corresponding to a spatial threshold $x^* \in [0, d]$ such that a user at location x is served by BS1 if $x \leq x^*$ and by BS2 otherwise. This can also be expressed as a threshold on*

the ratio of received signal strengths from the two base stations.

Proof Sketch: Since traffic intensity measure $\lambda(\cdot)$ is non-atomic, if the service regions associated with the BS1 and BS2 are not contiguous, one can construct regions \mathcal{R}_1 and \mathcal{R}_2 satisfying Lemma 8.3.1 which can then be exchanged without increasing the mean delay. Thus an optimal policy must be defined by contiguous regions, i.e., specified by a spatial threshold. Since the ratio of the received signal strengths is strictly decreasing or increasing with the received signal strength (or distance) from a base station, the policy can also be implemented as a threshold on this ratio.

Note that optimal static load allocation policies need not necessarily be unique. For example, consider the case when user requests are distributed homogeneously on the line segment joining the two base stations. If the optimal threshold does not correspond to the midpoint, then by symmetry, the policies that divide the service areas using thresholds at a distance d^* from BS1 and d^* from BS2 will result in identical mean user delays.

8.4 Optimal Threshold Trends

As a consequence of Theorem 8.3.1, we need only consider threshold-based static allocation policies. Fig. 8.3 exhibits again the simulated mean user delay for varying thresholding policies as the (spatially homogeneous) arrival rate between the base stations increases. The policies are characterized by the fraction of load served by BS1 with 0.5 corresponding to load balancing




Figure 8.3: Impact of load split threshold on delay performance.

and 0.1 to only 10% of the load. As noted earlier, due to symmetry, the delay performance would be identical if the threshold were moved closer to BS2. For each arrival rate, the optimal load split, i.e., roughly achieving the minimum mean delay, is highlighted. One can make the following observations:

1. The location of the optimal threshold is a function of the load on the system.
2. Except at very low loads, delay performance is improved by moving the threshold away from the mid-point, thus inducing asymmetrical loads on the two base stations.

Why does this happen? Load balancing increases parallelism, i.e., base stations are more likely to be simultaneously active. In our model, load balancing associates users with close by base stations providing them a stronger

signal. Finally, it would appear that load balancing might be beneficial in terms of statistical multiplexing at the two base stations. If capacity users were fixed, these points would provide the right insight. Yet, when dynamic interference is present, the capacity users see (particularly those far from either base station) can be substantially reduced by interference, and the fraction of time that base stations interfere with each other depends on the traffic and the load allocation policy. Thus, when arrival rate is low, the probability of the base stations being simultaneously active is low; the base stations operate in an interference-free environment, and load balancing is roughly optimal. For higher arrival rates, performance is strongly impacted by interference, and skewing the load is beneficial. Intuitively, this skew reduces the utilization of one of the base stations, say BS1, and thus the interference it causes on BS2's users, which reduces BS2's utilization, in turn benefiting BS1. However, one cannot overdo this skew as serving users that are far away, and thus have poor received signal, is also detrimental. Finally, it is tempting to assume that as load increases, base stations are always busy and the role of dynamic coupling reduces. Yet, as can be seen, at high loads performance sensitivity is also high, and the gains of an optimal asymmetric split increase further. The optimal threshold reflects a complex tradeoff among dynamic interference, statistical multiplexing, and users' signal strengths.

8.5 Optimizing the Threshold

In this section, an approximation methodology is developed for optimizing static load allocation policies for the wireless network model in Sec. 7.2, naturally extended to N base stations serving a possibly higher dimensional region. A policy π partitions the service area such that base station n has service area \mathcal{X}_n^π and overall arrival rate $\lambda_n = \lambda(\mathcal{X}_n^\pi)$. First, we approximate the Markovian model with uncountable state space by one with a countable state space, i.e., we will no longer keep track of the locations of users associated with each base station. This involves introducing an ‘effective’ rate for *all* users associated with a base station which depends on the busy state of the remaining base stations. Thus, the model preserves the dynamic interference characteristics. Second, the methodology developed in Chapter 7.3 is used to upper/lower bound the performance for the approximated model. Finally, performance is optimized over families of static policies that can be easily parametrized, e.g., for our one dimensional example, one need only determine the threshold.

Countable state-space approximation: Let $\vec{Q}(t) = (Q_n(t), n = 1, \dots, N)$ denote the number of active users at each base station at time t for our approximated process. For notational simplicity, we have suppressed its dependency on π . As mentioned earlier, the capacity to a user depends on *both* its current location and the interference profile it sees from neighboring base stations. Let $\vec{\Delta}(t) = (\Delta_n(t), n = 1, \dots, N)$ where $\Delta_n(t) = \mathbf{1}(Q_n(t) > 0)$ denote the status (idle or busy) or the ‘interference profile’ of the base stations.

Note that $\vec{\Delta}(t)$ can take 2^N possible values which we denote $\vec{\delta}^i, i = 1, \dots, 2^N$. Let $c_n(x, \vec{\delta}^i)$ denote the actual capacity at which base station n can serve a user at location $x \in \mathcal{X}_n^\pi$ under interference profile $\vec{\delta}^i$. In our approximate model, the effective capacity for such a user depends *only* on $\vec{\delta}^i$ and is given by

$$c_n^{\vec{\delta}^i} = \left(\int_{\mathcal{X}_n^\pi} \frac{1}{c_n(x, \vec{\delta}^i)} \frac{\lambda(dx)}{\lambda_n} \right)^{-1}.$$

In other words, $c_n^{\vec{\delta}^i}$ is the *harmonic mean* of the users actual service capacities weighted by the spatial distribution of arrivals to the base station, i.e., $\frac{\lambda(dx)}{\lambda_n}$. The harmonic mean is the appropriate average since the incremental time users spend in the system is inversely proportional to their service capacity, as shown in Sec. 3.2.1.1. Since files have mean size of 1, the total service rate $\mu_n^{\vec{\delta}^i}$ at base station n under interference profile $\vec{\delta}^i$ is given by $\mu_n^{\vec{\delta}^i} = c_n^{\vec{\delta}^i}$. This model fits precisely the one in Sec. 7.2, and we can use the bounding methodology developed in Chapter 7 to obtain lower and upper bounds on the sum queue length. Then, using Little's law, one can obtain bounds on the mean delay experienced by users.

8.5.1 Determining Optimal Thresholds

As mentioned earlier, when policies can be easily parameterized, one can use these bounds to optimize performance. For the two base-station scenario, Theorem 8.3.1 shows the optimal static load allocation policy is a simple threshold. So for any threshold, Problem 7.3.3 can be solved with the sum queue length as the objective function to determine bounds on the mean de-

lay, and a simple line search can be used to determine the threshold giving the smallest lower bound on the mean delay. In the case of the three base station network considered in the sequel, we parametrize policies based on weights associated with the base stations, as described in Sec. 8.6.3.




Figure 8.4: Goodness of optimized thresholds.

Fig. 8.4 exhibits the computed approximate optimal thresholds versus those obtained via brute force simulation for our two base station model. Semidefinite optimization problems associated with relaxations of order 2 were solved using [40] and [81] to determine the necessary bounds. As can be seen both load splits (thresholds) and resulting mean delay performance are very close, supporting the accuracy of our optimization methodology. The optimization approach however provides the flexibility to address complex traffic loads as well as systems with a larger number of base stations.

8.6 Performance Comparison

8.6.1 Comparing Static Policies

Fig. 8.5-i again illustrates the impact that the choice of threshold location has on delay performance. The user distribution is spatially homogeneous, so locating the threshold at the midpoint between the base stations corresponds to a static load balancing approach. As can be seen, the resultant mean user delays are greatly decreased by choosing an optimal threshold, particularly at moderate to high system loads. Fig. 8.5-ii further exhibits the spatial distribution of user delays under the two schemes when the rate at which user requests arrive in the network is 1.2 per second. Surprisingly, skewing the load towards one base station does not result in a trade off where a subset of the users, e.g., at the heavily loaded base station, experience poor performance. Instead, under the optimal policy, the overall impact of inter-cell interference is reduced such that all users, irrespective of their spatial location or perceived signal strength, see improved performance on average.

8.6.2 Optimized Policy vs. Dynamic Strategies

Next, the performance of the optimal static policy is compared versus the following three dynamic policies:

Greedy User: each new user joins the base station which offers the highest current service rate. This requires knowledge of the new user's capacity to each base station when the neighbor is active/idle and the number of users each is serving.




Figure 8.5: Performance of the optimized policy vs. load balancing.

Greedy System: each new user is assigned to the base station so as to maximize the resulting current sum service rate of the base stations. This policy is more complex than the Greedy User policy as in addition it requires knowledge of the capacity for *all* ongoing users with and without interference.

Repacking: each time a user arrives or leaves, the assignments of *all* users are chosen so as to maximize sum service rate of the base stations via a brute force search – the overheads and complexity of such a scheme would be unrealistically high, yet we hypothesize that it results in the best delay performance among non-anticipative dynamic schemes.

Fig. 8.6 illustrates the mean delay (logarithmic scale) for varying traffic loads under the above-mentioned greedy policies. Surprisingly, the optimal static policy substantially outperforms the two greedy policies at moderate to high loads. Indeed, at high load, the mean delay of the static policy is 6




Figure 8.6: Mean Delay - Static vs. Greedy

times lower than the greedy system policy which itself is orders of magnitude lower than the greedy user policy. As expected, the repacking policy shown in Fig. 8.7 (linear scale) is the best, but indeed very close to the optimal static policy.

Fig. 8.8 exhibits the spatial delay distribution under the system-level greedy scheme vs. the static policy. While the greedy policy exhibits perhaps desirable spatially symmetric performance, it is still the case that the optimal static policy gives better performance to all user locations.

8.6.3 Three Base Station Network

The three base station case can be used as a building block to develop a load allocation policy in a larger network. The number of base stations that can potentially serve a particular user request is unlikely to be very large. A load




Figure 8.7: Mean Delay - Static vs. Repacking

association policy that decides only between the three strongest base stations for each user request seems to be a reasonable tradeoff between complexity and performance. For the 2 dimensional three base station network described in Sec. 8.2, the form of the optimal static association policy is difficult to characterize. The ‘optimal’ static association policy is computed within a family of policies that can be easily parametrized.

8.6.3.1 Weighted Signal Strengths

The first family of policies considered is parametrized by base station weights. Each base station is assigned a weight and a user is associated with the base station that offers the maximum weighted received signal strength. The weight associated with one of the base stations is set to 1, and a simple gradient descent is used to determine weights for the remaining base stations.




Figure 8.8: Spatial delay distribution - Static vs. Greedy system

The bounding methodology described in Chapter 7 is used to approximate the mean delay at each step of the gradient descent algorithm.

If the base stations are part of a larger network, accurately accounting for the activity levels of other neighboring base stations could be important. The proposed methodology can still be used in such a scenario by including queues corresponding to the neighboring base stations when the performance bounds are computed using the semidefinite optimization. The objective function would remain the expected sum queue length at the three sectors under consideration. Note that including additional base stations will increase the complexity of the bounding procedure.




Figure 8.9: Load division after projecting down to the two dimensional hyper-plane

8.6.3.2 Pairwise Optimization

As an alternative to the methodology proposed above, we consider a family of policies where modifying a single parameter while keeping the rest constant allows the load division between two base stations to be modified without affecting the set of users served by the other base station. Note that the policy presented in Sec. 8.6.3.1 does not possess this property as changing the weight associated with any base station potentially changes the load served by all three base stations. This property allows the sequential optimization of the policy parameters, and the optimal policy can be determined using a sequence of iterations where one parameter is adjusted in each iteration.

This is particularly important if additional neighboring base stations

have to be taken into account. When one of the parameters is being optimized, the only sectors that have to be considered in the optimization are the neighbors of the two base stations that are affected. Thus, each parameter can be optimized while accounting for a different set of neighbors. This reduces the complexity of the semidefinite program that has to be solved to obtain the performance bounds. In this paper, we only consider the three base station scenario in isolation. Evaluating the performance of the proposed technique in large networks is a topic for future study.

The vector of received signal strengths from the three base stations, $\vec{s}(x) = (s_1(x), s_2(x), s_3(x))$, is projected down on to the two dimensional hyperplane that passes through the origin and is orthogonal to the vector $(1, 1, 1)$. The family of static policies that is considered divides this hyperplane into regions, and a base station serves all users whose projected signal strength vector falls in its region. The hyperplane is chosen such that users with identical relative received signal strengths from the base stations are mapped to the same point. The projected vector, after an orthogonal transformation is given by $\vec{z} = \{z_1, z_2\}$, where

$$\begin{aligned} z_1 &= \frac{1}{\sqrt{6}}(2s_1(x) - s_2(x) - s_3(x)) \\ z_2 &= \frac{1}{\sqrt{2}}(s_2(x) - s_3(x)). \end{aligned}$$

The hyperplane is divided into three regions by three rays extending from the origin, as shown in Fig. 8.9. Each base station serves the region between two rays as illustrated in the figure. The rays are specified by the angles

α, β , and γ that they subtend with the z_1 axis, and these angles parametrize a policy within the family. Rotating one of the rays only exchanges load between the two base stations whose service regions adjoin the ray. The optimal static policy is determined through a series of iterations. At each iteration, one of the parameters is modified, and a new value that improves the overall delay experienced by the set of users served by the three base stations is chosen. Thus, each iteration lowers the overall mean delay experienced by users in the system, ensuring that the optimization procedure converges.




Figure 8.10: Three Base Station Network: Mean Delay Performance under the weighted signal strength policy

Fig. 8.10 exhibits the mean delay performance in a three base station network. The repacking policy for this case is a hard combinatorial problem to be solved upon each arrival/departure and so was infeasible. The static load balancing and the greedy user policies exhibit similar performance, i.e., overlap, while the optimized static (asymmetric) policy exhibits substantial performance gains. Even the greedy system policy (itself unrealistic in practice)

achieves mean delays up to 20 times higher than the weighted signal strength based policy. The projection based policy performs significantly better than even the weighted signal strength policy, reducing the user perceived mean delay further by 6-10 times at high loads. These results suggest that quasi-static policies adapting to slowly changing, predictable traffic loads would achieve excellent performance.

Chapter 9

The Interference-Sensitive Adaptive Policy

The static policy developed in Chapter 8 requires knowledge of the long-term traffic loads served by the wireless network. Also, several iterations of a semidefinite optimization problem have to be solved in order to determine the optimal thresholding policy. Further, the static policy determined through the optimization procedure may not be robust to quickly changing traffic loads. In this section, I present the Interference-Sensitive, Adaptive Policy (ISAP) that divides load among base stations and induces asymmetry by tracking the impact of performance coupling among base stations resulting from dynamic inter-cell interference. The proposed policy requires no communication among base stations and only requires each base station to track two simple measures of the load being served.

The load on the system depends not only on the rate at which users arrive and the mean file size requirements but also on their location with respect to the base stations. A load allocation policy must be sensitive to both the intensity of the load as well as its distribution in space. The policy must be able to distinguish between scenarios where inter-cell interference is responsible for causing high user delays, and scenarios where user delays are

driven by high traffic loads inherent in the system. As seen in the previous sections, in an interference dominated scenario, an adaptive scheme may need to create an asymmetric division of load among base stations. The scheme should ideally create an asymmetric division of load even from an initial point at which the base stations could be identically loaded. Note that schemes where the desirability of a base station depends solely on the nature of the load supported by just that base station will not possess this property.

For example, when traffic is concentrated near the midpoint between cells, even a low intensity of user arrivals can result in high user perceived delays. In this case, inducing asymmetry in the load carried by the base stations would improve performance by reducing the impact of inter-cell interference. However, if the traffic carried by the base stations is concentrated near them (and away from the other base stations), the users are affected only marginally by inter-cell interference. In this case, there is not much gain possible from inducing asymmetry even if the observed mean delays are high. Ideally, a dynamic load allocation scheme should favor a load balancing approach in such a scenario.

9.1 Measuring the impact of interference

In order to estimate the effect of inter-cell interference, each base station i tracks and maintains estimates for the *system load* and *virtual load* as described below. Let $\hat{\rho}_i^S$ denote an estimate of the true *system load*, the current utilization of the downlink queue of base station i . This will reflect the

effect of interference from neighboring base station transmissions. Also, let $\hat{\rho}_i^V$ denote an estimate of the *virtual load*, the base station utilization that would result if base station i experienced no interference. Measurements are updated in discrete time slots of length δ . In any slot, the base station is either idle or transmitting to exactly one user, say user j . Slots are indexed by $t \in \mathbb{Z}$, corresponding to times δt . The transmission rate to user j in slot t under the current policy, taking into account the current activity state of the neighboring base station(s) is denoted $r_{ij}(t)$, and the rate to the user in the absence of *any* interference is denoted $r_{ij}^0(t)$. Each base station estimates the current system load and the virtual load resulting from the current load allocation policy as follows.

Estimating the system load: The system load can be estimated by periodically checking if there are active users associated with the base station. Each base station updates the estimate for the system load at δ intervals as follows:

$$\begin{aligned} \hat{\rho}_i^S(t+1) &= \beta_S \mathbf{1}(\text{BS } i \text{ is transmitting in slot } t) + \\ &\quad (1 - \beta_S) \hat{\rho}_i^S(t), \end{aligned}$$

where $\beta_S \in (0, 1)$ is a small constant determining the averaging time scale. Note that we assume time slots are small enough that base stations are either on or off for the entire duration of a slot.

Estimating the virtual load: A base station's virtual load is measured as the fraction of time the base station would be actively transmitting

to users if it were to serve the same traffic in the absence of interference. The estimate for the virtual load is updated along with the system load as follows:

$$\hat{\rho}_i^V(t+1) = \beta_V q_i^V(t) + (1 - \beta_V) \hat{\rho}_i^V(t),$$

where $\beta_V \in (0, 1)$ is a small constant, and the function $q_i^V(t)$ is defined as

$$q_i^V(t) = \begin{cases} 0, & \text{BS } i \text{ is idle} \\ \frac{r_{ij}(t)}{r_{ij}^0(t)}, & \text{BS } i \text{ is transmitting to user } j. \end{cases}$$

One can interpret the virtual load as follows: The virtual system serves exactly the same user as the real system in each slot. The virtual system transmits exactly the same number of bits to the user as the real system, by using only a fraction $\frac{r_{ij}(t)}{r_{ij}^0(t)}$ of the slot. Thus, when the user in the real system experiences interference in a slot, the slot is only partially used in the virtual system and the base station is idle for the remainder of the slot. In the case where the channel to the users is time invariant, the fraction of time that the base station is transmitting under the virtual system is identical to the utilization that would result from any work conserving scheduling policy in the absence of interference. Since the rate at which a user can be served is a constant in this case, one could rearrange the on and off periods of the base station to match any work conserving policy. Note that this is not true in general, in the case of time varying channels and arbitrary scheduling disciplines. However, a similar virtual system could hypothetically be constructed for such a case also.

Estimating interference impact: Clearly, $\hat{\rho}_i^S$ will always be greater than $\hat{\rho}_i^V$. The overall impact that inter-cell interference has on base station i

can be measured by a function of both $\hat{\rho}_i^S$ and $\hat{\rho}_i^V$, such as $(\hat{\rho}_i^S - \hat{\rho}_i^V)$ or $\frac{\hat{\rho}_i^S}{\hat{\rho}_i^V}$. For instance, if the load served by the base station is concentrated near the mid point between the base stations, $\hat{\rho}_i^S$ will be high due to the effect of inter-cell interference. However, $\hat{\rho}_i^V$ will be significantly lower even after accounting for path loss, as the major factor contributing to the low rates experienced by cell edge users is inter-cell interference. However, if the bulk of the load is concentrated near the base station, $\hat{\rho}_i^S$ will be close to $\hat{\rho}_i^V$.

9.2 Algorithm to determine the serving BS

9.2.1 Two base station scenario

We begin by considering the two base station case. Suppose a request from user j arrives at slot t , and is to be assigned to one of base stations $i = 1, 2$. The proposed ISAP policy is a simple weighted maximum rate policy, that is, the user connects to the base station $i^* = \arg \max_{i=1,2} w_i(t)r_{ij}^0(t)$ with ties broken arbitrarily. The novelty in the policy lies in specification of time dependent weights $w_i(t)$ which in turn are nonlinear functions of the current estimates for the true system and virtual traffic loads seen at *both* of the base stations.

Specifically, the base stations share their current estimates $\hat{\rho}_i^S, \hat{\rho}_i^V$, $i = 1, 2$ with the user. The user in turn assigns a weight of 1 to the base station which currently has the lowest system load. The other base station is assigned a weight which exceeds 1 and increases with the degree of inter-cell interference experienced by the base stations. There are many possibilities for doing this,

yet in this dissertation I consider the following specific weight assignment:

$$w_1(t) = \begin{cases} 1, & \hat{\rho}_1^S(t) \leq \hat{\rho}_2^S(t) \\ 10^{(\gamma \prod_{i=1}^2 (\hat{\rho}_i^S - \hat{\rho}_i^V))}, & \hat{\rho}_1^S(t) > \hat{\rho}_2^S(t) \end{cases}$$

where γ is an appropriately chosen constant. The weight assigned to base station 2 is similarly computed. The idea underlying this weight assignment is as follows. If $\hat{\rho}_i^S \approx \hat{\rho}_i^V$ for *one or both* of the base stations then the weight's exponent is roughly 0 and the base station with heavier load will have a weight which is only slightly larger than 1. In this case the policy reduces to a greedy max rate policy. However, if *both* base stations are subject to interference, then $\hat{\rho}_i^S - \hat{\rho}_i^V > 0$ for $i = 1, 2$. So the weight associated with the heavier loaded base station is quickly increasing with the degree of interference seen by base stations's users, and larger than 1. In this case, the policy becomes a weighted max rate policy with a bias towards attracting additional load to the heavier loaded of the base stations, the type of load asymmetry that was found to be advantageous earlier. Clearly other suitable functions of the system and virtual loads are possible, yet the above appears to be reasonable and work well.

9.2.2 Multi base station scenario

The proposed policy readily extends to the case of a network with multiple base stations. Suppose a user j arrives at time t and can associated with any one of n base stations. For simplicity assume the possible base stations are indexed $i = 1, \dots, n$ such that they have decreasing system loads. The multiple base station association policy exhibited in Algorithm 1, relies

Algorithm 1 Assigning user j to one of n BSs

```
1: Sort the base stations in decreasing order of  $\hat{\rho}_i^S(t)$ .
2: Let  $i^* = 1$ 
3: for  $i = 2$  to  $n$  do
4:    $w_{i^*} = 10^{(\gamma(\hat{\rho}_{i^*}^S - \hat{\rho}_{i^*}^V)(\hat{\rho}_i^S - \hat{\rho}_i^V))}$ 
5:    $w_i = 1$ 
6:   if  $w_{i^*}(t)r_{i^*j}^0(t) < w_i(t)r_{ij}^0(t)$  then
7:      $i^* = i$ 
8:   end if
9: end for
10: return  $i^*$ 
```

on making pairwise comparisons among base stations starting from the base stations which are seeing the heaviest loads. The idea is once again to favor asymmetries in load towards base stations which are seeing high system loads, but only if they are also strongly coupled through interference with one of the other base stations.

By assigning a weight larger than 1 to the heavily loaded base station, ISAP induces asymmetry in the loads served. However, the asymmetry is controlled as the policy is also sensitive to the rate at which the users can be served by the base stations. Note that if either base station is not affected by interference, the weight associated to the heavily loaded base station will be very close to 1, resulting in a policy that resembles a greedy maximum rate policy.

9.3 Performance Evaluation

The delay performance of ISAP is compared to the other dynamic schemes introduced in Sec. 8.6, and the static policy resulting from our approximate SDP based optimization. The value for the constant multiplicative factor γ is 3, unless noted otherwise. This constant should be chosen to ensure that the dynamic range of the weights is sufficiently large.


9.3.1 Two base station scenario

Thus far, all the simulation results exhibited performance under spatially homogeneous user (load) distributions. In this section, various spatially non-homogeneous load profiles are additionally considered as shown in Fig. 9.1a-9.1e. The line segment joining the two base stations is split into four quarters, and the load distribution is varied by varying the proportion of users in each quarter. Users in a particular quarter are uniformly distributed within that quarter. Load profiles 1, 2 and 4 are symmetric with respect to the midpoint between the base stations. The users are concentrated near the base stations in profile 2, and the impact of inter-cell interference is diminished. The effective load on the network under this profile is lighter at a fixed user arrival rate compared to profile 4, where users are concentrated close to the midpoint and are strongly impacted by inter-cell interference. The load distribution under profiles 3 and 5 is asymmetric.


The optimized static policy and ISAP perform consistently well under all spatial load profiles, and perform as well as or outperform all the dynamic

policies. This demonstrates their robustness to spatially heterogeneous traffic loads. Under load profile 3, for example, they outperform all the other schemes by a wide margin. None of the other schemes perform well under all profiles. The proposed schemes are able to infer the nature of the spatial load and adapt to it. Under load profiles 4 and 5, choosing a higher value for the multiplicative constant γ is necessary for the performance of ISAP to match the optimized static policy. Thus, in order to achieve optimal performance, ISAP has to be parametrized depending on the load distribution. However, even if a nominal value is chosen for γ , ISAP performs very well, even if it is not optimal.


The relative performance of the dynamic schemes can vary dramatically with the distribution of the spatial load. The greedy system scheme performs well under profiles 1 and 4. However, it is the worst among the schemes under load profile 2. Since the greedy system scheme tries to maximize the average throughput realized by all users in the system, it might deviate from a load balancing policy so as to ensure that a base station stays idle. However, since users cannot be reassigned, such decisions adversely affect long-term delay performance. The static max rate scheme performs well under load profile 2, where the effect of interference is minimal and under profile 4, where the spatial load is inherently asymmetric. It performs very poorly under the other spatial profiles.



(a) Spatial Load Profile 1

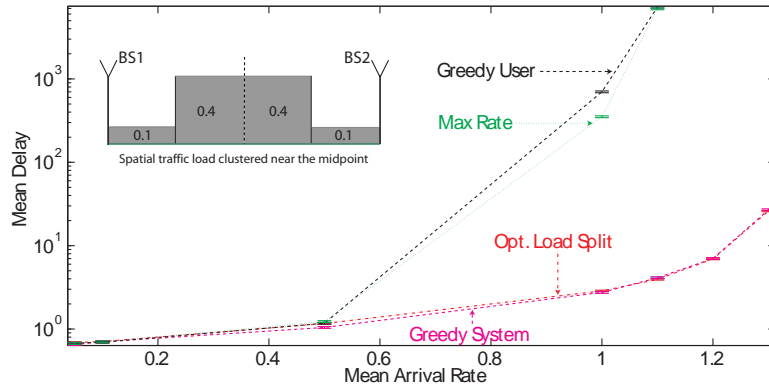


(b) Spatial Load Profile 2




(c) Spatial Load Profile 3

Figure 9.1: ISAP: Delay performance under heterogeneous spatial loads



(d) Spatial Load Profile 4



(e) Spatial Load Profile 5

Figure 9.1: ISAP: Delay performance under heterogeneous spatial loads

9.3.2 Three (or More) Base Station Network:

Fig. 9.2 exhibits the mean delay performance that users perceive in a three base station network under ISAP with γ set to the nominal value of 3, the projection based static policy, and the greedy dynamic policies. We only consider the spatially homogeneous user distribution in this case. Evaluating the performance of the various static and dynamic policies under spatially het-




Figure 9.2: ISAP: Delay performance in a three base station network with homogeneous load.

erogeneous user distributions is a topic for future study. ISAP results in delay performance that is comparable to the projection based static policy. Both policies perform significantly better than the dynamic policies. The results indicate that ISAP performs well even in a multiple base station network.

9.3.3 Performance Sensitivity

Channel Model: The dependence of the performance results on the parameters of the system channel model is examined using parameters that model cellular base stations in an urban environment. A system consisting of two base stations 2800 meters apart is simulated, and the performance of the schemes presented earlier are compared using a path loss exponent 3.5, and a cell-edge signal to noise ratio of 10 dB. The data rate at which users are served is calculated using Shannon’s capacity formula, after a 6dB backoff is applied to the perceived SINR.




Figure 9.3: Sensitivity to channel model: Optimal load split thresholds

Fig. 9.3 shows the simulated delay performance when load split between the base stations is varied from 0.5 (even division of load) to 0.1, similar to Fig. 6.1b. The results again show that the optimal load division depends on the intensity of the offered load, and is not balanced but significantly asymmetric.




Figure 9.4: Sensitivity to channel model: Delay performance

Fig. 9.4 compares the delay performance of the optimized static policy and ISAP to the dynamic schemes described earlier. The proposed schemes significantly outperform the dynamic schemes. The mean delay under the

greedy system scheme, for example, is over 50 times the mean delay under the optimized static scheme at high loads. These results demonstrate that the performance trends observed earlier do not depend on the particular channel model used and are a consequence of the dynamics introduced by inter-cell interference.

Long tailed file size distributions: The effect of long tailed file size distributions on the mean delay experienced by users is examined. In the process of determining the optimized static threshold, we still assume that file sizes are exponentially distributed. We assume that the users' file size requirements are log normally distributed with mean 5 MB, and variance 12.276×10^6 . The performance of the various schemes under a spatially homogeneous user distribution is shown in Fig. 9.5. The relative performance of the different

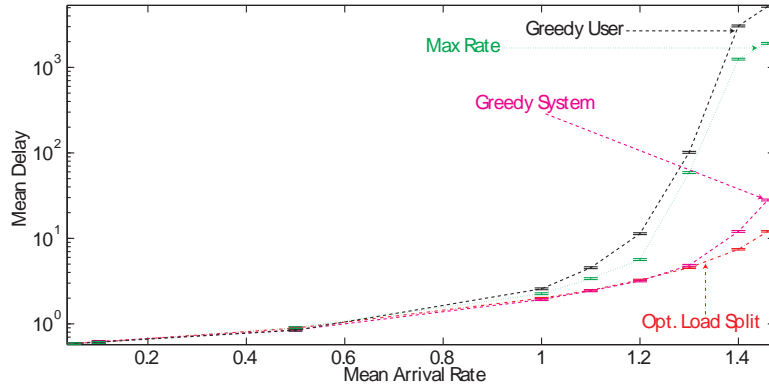


Figure 9.5: Sensitivity to file size distribution: Delay performance

schemes is very similar to the case of exponential file sizes. The optimized static policy and the policy developed above, ISAP, result in the best performance. The optimization procedure and the proposed adaptive policy appear

to be robust to variations in the distribution of users' file size requirements.

9.4 Summing Up

For the one and two dimensional models considered, the performance gain from optimized static policies is substantial, even outperforming natural greedy user and system dynamic policies. The load-balancing static policy was shown to be very poor, showing that the critical aspect is inducing asymmetry in the load, even when the network and loads are symmetric. Our simulation results demonstrated that our proposed policies perform consistently well under all spatial loads and are robust to variations in file size distributions and channel parameters. The performance of the conventional dynamic policies was found to vary dramatically with the load distribution, and no one policy performed consistently well. This work suggests the possibility that substantial gains might be achieved if network functions coupled through interference (or otherwise) are optimized for dynamic loads.

Conclusion

This dissertation focused on the analysis and optimization of wireless networks serving dynamic loads and thus coupled through interference. I addressed two key questions regarding the operation of interference dominated systems: cooperative scheduling and user association. A low complexity, system-level approach that substantially improves performance perceived by best-effort users without requiring high channel measurement and estimation, communication, and computational overheads was developed. The proposed approach simultaneously achieved spatially homogeneous performance while also reducing the transmit power requirements. I also studied the user association problem in such networks, and proposed a methodology to bound and optimize performance in the underlying coupled queuing systems. The proposed user association policies were shown to outperform a load balancing policy as well as several dynamic policies under both spatially homogeneous and heterogeneous loads.

The results in this dissertation demonstrated the impact that interference induced coupling can have on user performance, and the sometimes counter-intuitive implications for network design. Performance and even the stability of wireless networks in this dynamic context is highly dependent on the spatial traffic loads being served. Policies that are tailored to the spatial load and the resulting interference characteristics can make a great difference. There are also subtle differences in the interference characteristics between the

uplink and the downlink resulting from the spatial origin of interference, and algorithms have to be adjusted depending on the context in which they are applied.

The two problems studied as part of this dissertation are, in reality, intertwined. Clearly, the user association policy determines the spatial load that each base station has to serve, while coordination across base stations alters the dynamics of interference and inter base station coupling. Ideally, these policies would be optimized jointly. However, as we have seen, optimizing interference while accounting for interference coupling is not an easy task, and any joint optimization is likely to be highly complex. Therefore, a division like the one proposed in this dissertation might be the most practical way around the problem. For instance, the coordination mechanism and the user association policy developed in this paper could operate together with the coordination scheme inferring the spatial distribution induced by the user association policy. One of the topics I intend to pursue in the future is the development of an adaptive scheme (like ISAP) for base station coordination that could work in tandem with user association and other policies impacted by the spatial load distribution. Opportunistic scheduling in the presence of fading channels offers similar challenges, with performance coupling among users. I believe that understanding the interactions between these various layers, and designing interference and load aware policies that operate together has the potential to significantly improve user performance in wireless networks. In the future, I intend to pursue research in this direction using the framework

established in this dissertation.

Bibliography

- [1] *Telecommunication Industry Association / Electronics Industry Association Interim Standard (TIA/EIA) IS95 Standard*. 1999.
- [2] N. Abramson. The aloha system-another alternative for computer communications. In *Proceedings of Fall Joint Computer Conference, AFIPS Conference*, volume 36, pages 295–298, 1970.
- [3] M. Anderberg. *Cluster Analysis for Applications*. Academic Press, 1973.
- [4] J. G. Andrews. Interference cancellation for cellular systems: A contemporary overview. *IEEE Wireless Communications Magazine*, 12(2):19–29, Apr. 2005.
- [5] D. Bertsimas and K. Natarajan. A semidefinite optimization approach to the steady-state analysis of queueing systems. *Queueing Syst. Theory Appl.*, 56(1):27–39, 2007.
- [6] D. Bertsimas, I. C. Paschalidis, and J. N. Tsitsiklis. Large deviations analysis of the generalized processor sharing policy. *Queueing Syst. Theory Appl.*, 32(4):319–349, 1999.
- [7] D. Bertsimas and I. Popescu. Optimal inequalities in probability theory: A convex optimization approach. *SIAM Journal on Optimization*, 15(3):780–804, 2005.

- [8] G. Bianchi and I. Tinnirello. Improving load balancing mechanisms in wireless packet networks. In *IEEE ICC*, volume 2, pages 891–95, 2002.
- [9] T. Bonald, S. Borst, N. Hegde, and A. Proutiere. Wireless data performance in multicell scenarios. In *ACM SIGMETRICS*, volume 32, June 2004.
- [10] T. Bonald, S. Borst, and A. Proutiere. Inter-cell scheduling in wireless data networks. In *European Wireless Conference*, 2005.
- [11] T. Bonald, S.C. Borst, and A. Proutiere. How mobility impacts the flow-level performance of wireless data systems. In *INFOCOM 2004*, volume 3, pages 1872–1881, 2004.
- [12] A. Borovkov and S. Foss. Stochastically recursive sequences and their generalizations. *Siberian Advances in Mathematics*, 2(1):16–81, 1992.
- [13] S. Borst. User-level performance of channel-aware scheduling in wireless data networks. In *INFOCOM*, pages 321 – 331, 2003.
- [14] S. Borst, O. Boxma, and P. Jelenkovic. Coupled processors with regularly varying service times. In *IEEE INFOCOM*, volume 1, page 157164, 2000.
- [15] S. Borst, O. Boxma, and M. van Uitert. The asymptotic workload behavior of two coupled queues. *Queueing Systems*, 43(1-2):81–102, January 2003.

- [16] S. Borst, A. Buvanewari, L. M. Drabeck, M. J. Flanagan, J. M. Graybeal, G. K. Hampel, M. Haner, W. M. MacDonald, P. A. Polakos, G. Rittenhouse, I. Saniee, A. Weiss, and P. A. Whiting. Dynamic optimization in future cellular networks. *Bell Labs Technical Journal*, 10(2):99–119, 2005.
- [17] S. Borst, N. Hegde, and A. Proutiere. Interacting queues with server selection and coordinated scheduling - application to cellular data networks. *Annals of Operations Research*, 170(1):59–78, September 2009.
- [18] S. Borst, I. Saniee, and A. Whiting. Distributed dynamic load balancing in wireless networks. In *ITC*, pages 1024–37, 2007.
- [19] S. C. Borst. Optimal probabilistic allocation of customer types to servers. In *ACM SIGMETRICS*, pages 116–25, 1995.
- [20] S.C. Borst, M. Jonckheere, and L.S. Leskel. Stability of parallel queueing systems with coupled service rates. *Discrete Event Dynamic Systems*, 18(4):447–472, 2008.
- [21] K. Chawla and Xiaoxin Qiu. Quasi-static resource allocation with interference avoidance for fixed wireless systems. *IEEE J. Select. Areas Commun.*, 17(3):493–504, Mar. 1999.
- [22] R. Choi, R. Murch, and K. Letaief. MIMO CDMA antenna system for SINR enhancement. *IEEE Transactions on Wireless Communications*, 2(2):240–249, March 2003.

- [23] W. Choi and J. Andrews. The capacity gain from intercell scheduling in multi-antenna systems. *IEEE Transactions on Wireless Communications*, 7(1):714–725, February 2008.
- [24] W. Choi and J. G. Andrews. Base station cooperatively scheduled transmission in a cellular MIMO TDMA system. In *Conference on Information Systems and Sciences (CISS)*, March 2006.
- [25] W. Choi and J. G. Andrews. Downlink performance and capacity of distributed antenna systems in a multicell environment. *IEEE Transactions on Wireless Communications*, 6(1), January 2007.
- [26] T. Chu and S. Rappaport. Coverage with reuse partitioning in cellular communication systems. *IEEE Transactions on Vehicular Technology*, 46(1):41–54, February 1997.
- [27] H. Dai, A. Molisch, and H.V. Poor. Downlink capacity of interference-limited MIMO systems with joint detection. *IEEE Trans. Wireless Commun.*, 3(2):442–453, March 2004.
- [28] S. Das, H. Viswanathan, and G. Rittenhouse. Dynamic load balancing through coordinated scheduling in packet data systems. In *IEEE INFOCOM*, volume 1, pages 786–96, 2003.
- [29] A. Demers, S. Keshav, and S. Shenker. Analysis and simulation of a fair queueing algorithm. In *Symposium proceedings on Communications*

architectures & protocols, pages 1–12, Austin, Texas, United States, 1989. ACM.

- [30] P. Dupuis and K. Ramanan. A skorokhod problem formulation and large deviation analysis of a processor sharing model. *Queueing Syst. Theory Appl.*, 28(1-3):109–124, 1998.
- [31] G. Fayolle and R. Lasnogorodski. Two coupled processors: The reduction to a Riemann–Hilbert problem. *Wahrscheinlichkeitstheorie*, (3):1–27, Jan. 1979.
- [32] T. K. Fong, P. S. Henry, K. K. Leung, Xiaoxin Qiu, and N. K. Shankaranarayanan. Radio resource allocation in fixed broadband wireless networks. *IEEE Trans. Commun.*, 46(6):806–818, Jun. 1998.
- [33] G. J. Foschini, K. Karakayali, and R. A. Valenzuela. Coordinating multiple antenna cellular networks to achieve enormous spectral efficiency. *IEEE Trans. Commun.*, 153(4):548–555, Aug. 2006.
- [34] A. Ghasemi and E. S. Sousa. Distributed intercell coordination through time reuse partitioning in downlink CDMA. In *IEEE Wireless Communications and Networking Conference*, volume 4, pages 1992–1997, Mar. 2004.
- [35] R. Graham. An efficient algorithm for determining the convex hull of a finite point set. *Info. Proc. Letters*, 1(4):132–133, 1972.

- [36] F. Guillemin and D. Pinchon. Analysis of generalized processor-sharing systems with two classes of customers and exponential services. *Journal of Applied Probability*, 41(3):832–858, 2004.
- [37] S. V. Hanly. An algorithm for combined cell-site selection and power control to maximize cellular spread spectrum capacity. *IEEE Journal on Selected Areas in Communications*, 13(7):1332–1340, September 1995.
- [38] G. Hardy, J. E. Littlewood, and G. Polya. *Inequalities*. Cambridge University Press, 2 edition, January 1952.
- [39] R. Hasegawa, M. Shirakabe, R. Esmailzadeh, and M. Nakagawa. Down-link performance of a CDMA system with distributed base station. In *IEEE Vehicular Technology Conference*, pages 882–886, October 2003.
- [40] D. Henrion and J. B. Lasserre. Gloptipoly: global optimization over polynomials with matlab and sedumi. *ACM Transactions on Mathematical Software*, 29(2):165–194, 2003.
- [41] P. A. Hoeher, S. Badri-Hoeher, Wen Xu, and C. Krakowski. Single-antenna co-channel interference cancellation for TDMA cellular radio systems. *IEEE Wireless Communications Magazine*, 12(2):30–37, Apr. 2005.
- [42] S. Jafar, G. J. Foschini, and A. J. Goldsmith. PhantomNet: Exploring optimal multicellular multiple antenna systems. In *IEEE Vehicular Technology Conference*, pages 24–28, Sep. 2002.

- [43] A. Jain and R. Dubes. *Algorithms for Clustering Data*. Prentice Hall, 1988.
- [44] W. H. Jean and B. P. Helms. Geometric mean approximations. *The Journal of Financial and Quantitative Analysis*, 18(3):287–293, September 1983.
- [45] S. Jing, D. N. C. Tse, J. Hou, J. B. Soriaga, J. E. Smee, and R. Padovani. Multi-cell downlink capacity with coordinated processing. In *Information Theory and Applications Workshop*, Jan. 2007.
- [46] S. Jing, D. N. C. Tse, J. B. Soriaga, J. Hou, J. E. Smee, and R. Padovani. Downlink macro-diversity in cellular networks. In *IEEE International Symposium on Information Theory*, pages 24–29, June 2007.
- [47] M. Jonckheere. Stability of two interfering processors with load balancing. In *Third International Conference on Performance Evaluation Methodologies and Tools*, 2008.
- [48] N. Kahale and P. E. Wright. Dynamic global packet routing in wireless networks. In *IEEE INFOCOM*, volume 3, pages 1414–1421, Apr. 1997.
- [49] M. Kang and M. Alouini. Quadratic forms in complex gaussian matrices and performance analysis of MIMO system with cochannel interference. *IEEE Transactions on Wireless Communications*, 3(2):418–431, March 2004.

- [50] K. Karakayali, G. J. Foschini, and R. A. Valenzuela. Network coordination for spectrally efficient communications in cellular systems. *IEEE Wireless Communications Magazine*, 13(4):56–61, August 2006.
- [51] J. Karlsson and B. Eklundh. A cellular mobile telephone system with load sharing—an enhancement of directed retry. *Communications, IEEE Transactions on*, 37(5):530–35, 1989.
- [52] H. Kim, Y. Han, and J. Koo. Optimal subchannel allocation scheme in multicell OFDMA systems. In *IEEE Vehicular Technology Conference*, pages 1821–1825, May 2003.
- [53] C. Knessl. On the diffusion approximation to two parallel queues with processor sharing. *IEEE Transactions on Automatic Control*, 36:1356–1367, December 1991.
- [54] A. G. Konheim, I. Meilijson, and A. Melkman. Processor-sharing of two parallel lines. *J. Appl. Probab.*, 18(4):952–956, 1981.
- [55] X. Lagrange and B. Jabbari. Fairness in wireless microcellular networks. *IEEE Transactions on Vehicular Technology*, 47(2):472–479, May 1998.
- [56] J.B. Lasserre. Bounds on measures satisfying moment conditions. *Annals of Applied Probability*, 12:1114–1137, 2002.
- [57] J. W. Lee, R. Mazumdar, and N. B. Shroff. Joint resource allocation and base-station assignment for the downlink in CDMA networks. *IEEE/ACM Trans. Netw.*, 14(1):1–14, 2006.

- [58] K. K. Leung and A. Srivastava. Dynamic allocation of downlink and uplink resource for broadband services in fixed wireless networks. *IEEE J. Select. Areas Commun.*, 17(5):990–1006, May 1999.
- [59] G. Li and H. Liu. Downlink dynamic resource allocation for multi-cell OFDMA system. In *IEEE Vehicular Technology Conference*, pages 1698–1702, October 2003.
- [60] J. Li, N. B. Shroff, and E. K. P. Chong. A static power control scheme for wireless cellular networks. In *IEEE INFOCOM*, volume 2, pages 932–939, 1999.
- [61] X. Liu, E. K. P. Chong, and N. B. Shroff. A framework for opportunistic scheduling in wireless networks. *Computer Networks*, 41:451–474, March 2003.
- [62] R. Loynes. The stability of a queue with non-independent inter-arrival and service times. *Proc. Cambr. Phil. Soc.*, 58:497–520, 1962.
- [63] A. Marshall and I. Olkin. *Inequalities: Theory of Majorization and its Applications*. New York: Academic Press, 1979.
- [64] K. Navaie and H. Yanikomeroglu. Downlink joint base-station assignment and packet scheduling algorithm for cellular CDMA/TDMA networks. In *IEEE ICC*, volume 9, pages 4339–44, 2006.

- [65] Abhay K. Parekh and Robert G. Gallager. A generalized processor sharing approach to flow control in integrated services networks: the single-node case. *IEEE/ACM Trans. Netw.*, 1(3):344–357, 1993.
- [66] S. Pietrzyk and G. J. Janssen. Subcarrier allocation and power control for QoS provision in the presence of CCI for the downlink of cellular OFDMA systems. In *IEEE Vehicular Technology Conference*, pages 2221–2225, April 2003.
- [67] F. P. Preparata and S. J. Hong. Convex hulls of finite sets of points in two and three dimensions. *Commun. ACM*, 20(2):87–93, Feb. 1977.
- [68] X. Qiu and K. Chawla. Resource assignment in a fixed broadband wireless system. *IEEE Communications Letters*, 1(4):108–110, Jul. 1997.
- [69] T. S. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall, 2002.
- [70] F. Rashid-Farrokhi, K. Liu, and L. Tassiulas. Transmit beamforming and power control for cellular wireless systems. *IEEE Journal on Selected Areas in Communications*, 16:1437–1450, October 1998.
- [71] W. Rho and A. J. Paulraj. Performance of the distributed antenna systems in a multicell environment. In *IEEE Vehicular Technology Conference*, pages 587–591, April 2003.
- [72] L. G. Roberts. Aloha packet system with and without slots and capture. *SIGCOMM Comput. Commun. Rev.*, 5(2):28–42, 1975.

- [73] A. Sang, X. Wang, M. Madhian, and R. D. Gitlin. Coordinated load balancing, handoff/cell-site selection, and scheduling in multi-cell packet data systems. *Wireless Networks*, 14(1):103–120, January 2008.
- [74] S. Shakkottai and A. Stolyar. Scheduling for multiple flows sharing a time-varying channel: The exponential rule. *American Mathematical Society Translations*, 207:185–202, 2002.
- [75] S. Shamai and B. Zaidel. Enhancing the cellular downlink capacity via co-processing at the transmitting end. In *IEEE Vehicular Technology Conference*, pages 1745–1749, May 2001.
- [76] A. Srivastava, N. K. Shankaranarayanan, and K. K. Leung. Sector-based resource allocation for broadband fixed wireless networks. In *IEEE Vehicular Technology Conference*, volume 3, pages 1680–1684, May 1998.
- [77] A. L. Stolyar and H. Viswanathan. Self-organizing dynamic fractional frequency reuse in ofdma systems. In *Infocom*, Apr 2008.
- [78] A. L. Stolyar and H. Viswanathan. Self-organizing dynamic fractional frequency reuse for best-effort traffic through distributed inter-cell coordination. In *Infocom*, 2009.
- [79] D. Stoyan. *Comparison Methods for Queues and Other Stochastic Models*. New York: John Wiley, 1983.
- [80] D. Stoyan, W. Kendall, and J. Mecke. *Stochastic Geometry and its Applications*. J. Wiley & Sons, Chichester, 1995.

- [81] J. F. Sturm. Using SeDuMi 1.02, a matlab toolbox for optimization over symmetric cones. *Optimization Methods and Software*, 11(12):625–653, 1999.
- [82] Third Generation Partnership Project. Radio access network work group 1 contributions. <http://www.3gpp.org>, September 2005.
- [83] Third Generation Partnership Project 2. Ultra mobile broadband technical specifications. <http://www.3gpp2.org>, March 2007.
- [84] P. Viswanath, D. N. C. Tse, and R. Laroia. Opportunistic beamforming using dumb antennas. *IEEE Trans. Inform. Theory*, 48(6):1277–1294, Jun. 2002.
- [85] S. W. Wales. Technique for cochannel interference suppression in TDMA mobile radio systems. *IEEE Trans. Commun.*, 142(2):106–114, Apr. 1995.
- [86] J. F. Whitehead. Signal-level-based dynamic power control for co-channel interference management. In *IEEE Vehicular Technology Conference*, pages 499–502, 1993.
- [87] J. H. Winters. Optimum combining in digital mobile radio with cochannel interference. *IEEE Transactions on Vehicular Technology*, 33(3):144–155, 1984.

- [88] X. Wu, A. Das, J. Li, and R. Laroia. Fractional power reuse in cellular networks. In *Proceedings of the 44th Allerton Conference on Communication, Control, and Computing*, September 2006.
- [89] E. Yanmaz, O.K. Tonguz, and R. Rajkumar. Is there an optimum dynamic load balancing scheme? In *IEEE GTC*, volume 1, 2005.
- [90] R. Yates and C. Y. Huang. Integrated power control and base station assignment. *IEEE Transactions on Vehicular Technology*, 44(3):638–644, August 1995.
- [91] William E. Young and Robert H. Trent. Geometric mean approximations of individual security and portfolio performance. *The Journal of Financial and Quantitative Analysis*, 4(2):179–199, June 1969.
- [92] A. Zemlianov and G. de Veciana. Load balancing of best effort traffic in wireless systems supporting end nodes with dual mode capabilities. In *CISS*, pages 1–6, March 2005.
- [93] H. Zhang and H. Dai. Co-channel interference mitigation and cooperative processing in downlink multicell multiuser MIMO networks. *European Journal on Wireless Communications and Networking*, 4th quarter 2004.
- [94] Z. Zhang. Large deviations and the generalized processor sharing scheduling for a two-queue system. *Queueing Syst. Theory Appl.*, 26(3-4):229–254, 1997.

- [95] Z. Zhang. Large deviations and the generalized processor sharing scheduling for a multiple-queue system. *Queueing Systems*, 28(4):349–376, 1998.
- [96] L. Zuluaga and J. F. Pena. A conic programming approach to generalized Tchebycheff inequalities. *Mathematics of Operations Research*, 30(2):369–388, 2005.

Vita

Balaji Rengarajan received the degree of Bachelor of Engineering in Electronics and Communication Engineering from the University of Madras in May 2002, graduating at the top of his class. In 2002, he entered the graduate school at The University of Texas at Austin. He was the recipient of a 2003 Texas Telecommunications Engineering Consortium (TxTEC) graduate fellowship. He received the degree of Master of Science in Electrical Engineering in May 2004. During the summers of 2005 and 2008, he was a student intern at Intel Corporation and Bell Laboratories, Alcatel-Lucent respectively. He was admitted to Ph.D. candidacy in 2007.

Permanent address: 301 West 39th St., Apt. 303
Austin, TX-78751

This dissertation was typeset with L^AT_EX[†] by the author.

[†]L^AT_EX is a document preparation system developed by Leslie Lamport as a special version of Donald Knuth's T_EX Program.