

**Multiplexing, Scheduling, and Multicasting
Strategies for Antenna Arrays in Wireless
Networks**

by

Michael J. Lopez

B.S., Johns Hopkins University (1995)
M.S.E., Johns Hopkins University (1995)

Submitted to the Department of Electrical Engineering and Computer
Science

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

August 2002

© Massachusetts Institute of Technology 2002. All rights reserved.

Author

Department of Electrical Engineering and Computer Science
August 30, 2002

Certified by

Gregory W. Wornell
Professor of Electrical Engineering
Thesis Supervisor

Accepted by

Arthur C. Smith
Chairman, Department Committee on Graduate Students

Multiplexing, Scheduling, and Multicasting Strategies for Antenna Arrays in Wireless Networks

by

Michael J. Lopez

Submitted to the Department of Electrical Engineering and Computer Science
on August 30, 2002, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

A transmitter antenna array has the ability to direct data simultaneously to multiple receivers within a wireless network, creating potential for a more integrated view of algorithmic system components. In this thesis, such a perspective informs the design of two system tasks: the scheduling of packets from a number of data streams into groups; and the subsequent spatial multiplexing and encoding of these groups using array processing. We demonstrate how good system designs can help these two tasks reinforce one another, or alternatively enable tradeoffs in complexity between the two. Moreover, scheduling and array processing each benefit from a further awareness of both the fading channel state and certain properties of the data, providing information about key flexibilities, constraints and goals.

Our development focuses on techniques that lead to high performance even with very low-complexity receivers. We first consider spatial precoding under simple scheduling and propose several extensions for implementation, such as a unified time-domain precoder that compensates for both cross-channel and intersymbol interference. We then show how more sophisticated, channel-aware scheduling can reduce the complexity requirements of the array processing. The scheduling algorithms presented are based on the receivers' fading channel realizations and the delay tolerances of the data streams. Finally, we address the multicasting of common data streams in terms of opportunities for reduced redundancy as well as the conflicting objectives inherent in sending to multiple receivers. Our channel-aware extensions of space-time codes for multicasting gain several dB over traditional versions that do not incorporate channel knowledge.

Thesis Supervisor: Gregory W. Wornell
Title: Professor of Electrical Engineering

Acknowledgments

Any thanks must begin with my adviser, Prof. Greg Wornell, who is an endless source of enthusiasm, patience, and ideas. Most of all, he places his students first, and through his mentorship, I have become a better researcher and a better engineer.

In my developing career path, I have had the good fortune to work under a number of other excellent supervisors. These include Prof. Jerry Prince, who took me under his wing at Johns Hopkins; Per Hamnqvist, who provided lots of real world career advice during my time at Digital Equipment Corporation; and Andy Singer, who helped show me the role of engineering research in industry. I would also like to thank my thesis readers, Prof. Vahid Tarokh, Prof. Hari Balakrishnan, and Kambiz Zangi, for their comments, and especially for helping to make this thesis useful to a larger community.

The Digital Signal Processing Group has provided a wonderful environment for research. It has been an inspiration to work with so many talented and successful people. In particular, I thank Nick Laneman, Albert Chan, and Stark Draper for many helpful discussions during the long process of pulling together the main themes of my work.

Thanks goes out to the National Science Foundation for providing me with a Graduate Research Fellowship. I would also like to acknowledge the generous support of Hewlett-Packard through the HP/MIT Alliance, the NSF again through Grant No. CCR-9979363, and the MARCO/DARPA Center for Circuits, Systems, and Software.

I am grateful for all of my friends, and Sunil Konath, Victor Luchangco, Marissa Schwartz, Christa Beranek, and Edwin Yuen have been friends in the truest sense. I have relied on each of you many times and shared many joyful experiences as well. Among the other friends and good times during these MIT years: the day trip to Salem with Luis, Michelle and others; Keith playing ultimate frisbee with a little too much enthusiasm for his own good; long conversions about theology with Nick and Terri; singing in the choir at Desmond and Beata's wedding; and listening to Ryan expound on various esoteric and scholarly topics. I am especially thankful for the Catholic Fellowship group, which provided a prayerful community, lasting friendships, and helped me become more aware of the role of God in my life.

Special mention goes to Rebecca Martel, who has been at my side throughout all of the stressful and exhilarating moments of the past year. You have supported me in so many ways and helped keep me going through this intense time. My life has grown so much richer since meeting you.

To my brother Dave, sister Kristen, grandparents, and the rest of my family: I treasure the family bond we have, and thank you for all you have done for me.

Finally, I thank my parents for a lifetime of support and love. I truly appreciate how you have encouraged, but not pressured, me in all my efforts. From an early age, you instilled in me the values that are so much a part of the person I am today.

Contents

1	Introduction	19
1.1	Outline of Thesis	21
2	Background on Transmitter Antenna Arrays	25
2.1	Notational Conventions	25
2.2	Communication with Single-Element Antennas	26
2.2.1	Performance Measures	28
2.3	Transmitter Antenna Arrays	32
2.3.1	Array Processing Techniques Under Timesharing	33
2.3.2	Spatial Multiplexing of Multiple Streams	37
2.3.3	Coordinated Versus Uncoordinated Receivers	41
3	Precoding with Simple Scheduling	43
3.1	Precoding for Multiuser Communications	44
3.1.1	Precoding for Triangular Channels	45
3.1.2	Precoding Over Arbitrary Channel Matrices	49
3.1.3	Performance of Precoding	51
3.1.4	Improving Upon Zero-Forcing Precoding	58
3.1.5	Relation to Other Matrix Channel Problems	58
3.2	Implementation Issues	63
3.2.1	Overview of New Implementation Issues	63
3.2.2	Ordering of Streams	64
3.2.3	Constellation Design to Reduce Precoding Power Loss	71
3.2.4	Channel Information	78
3.3	Precoding for Combined Multiuser and Intersymbol Interference	81
3.3.1	Multiple-Receiver Dispersive Channels	82
3.3.2	Canceling Multiuser Interference with Causal Processing	83

3.3.3	Precoding with Noncausal Filtering	85
3.3.4	Comparison with DMT Method	91
3.4	Concluding Remarks	95
4	Informed Data Scheduling	97
4.1	Data Model: Classification by Delay Tolerance	98
4.2	Spatial Multiplexing Performance for Data with Medium Delay Tolerance	100
4.2.1	Diversity Analysis with Random Channel Vectors	100
4.2.2	Diversity With Orthogonal Channel Vectors	102
4.3	Scheduling Algorithms for Data with Medium Delay Tolerance	105
4.3.1	Static Model Example	106
4.3.2	Dynamic Queuing Model	107
4.4	Large Delay Tolerance	113
4.4.1	Relation to Timesharing Strategies	113
4.4.2	Scheduling for Spatial Multiplexing	116
4.5	Tight Delay Constraints	119
4.6	Multiplexing Different Classes of Data	121
5	Multicasting of Common Information	125
5.1	Overview of Multicast	126
5.2	Operating Points for Beamforming	128
5.3	Maximizing Average SNR per Receiver	132
5.3.1	Average Performance Per Receiver	132
5.3.2	Individual Receiver Performance	135
5.4	General Space-Time Multicast Coding	137
5.4.1	Optimal Structures	137
5.4.2	Beamforming Versus Higher-Rank Covariances	142
5.4.3	Implementation Issues	146
5.4.4	Performance of Higher-Rank Covariance Matrices	150
5.5	Multicast Within Larger Systems	152
5.5.1	Integration Among Array Processing Subblocks	152
5.5.2	Integration at the Scheduling Layer	154
6	Conclusions and Future Work	157
A	Ordering of Two Streams to Maximize Sum Capacity	161

List of Figures

1-1	Block diagram of transmitter with an antenna array, illustrating scheduling and array processing system tasks.	20
2-1	Channel model for transmitter antenna array and multiple receivers.	33
2-2	Ergodic capacities for an M -element transmit array and input SNR per link of 5 dB. We show the curve for Rayleigh fading, and for comparison the additive white Gaussian noise channel that does not encounter fading.	35
2-3	Single-receiver outage probabilities for a transmit antenna array with $M = 1, 2, 3,$ and 4 elements, using a Rayleigh fading model at an input SNR per link of 5 dB.	36
2-4	Achievable rate regions for a sample channel realization with orthogonal channel vectors. $\ \mathbf{h}_1\ ^2 = 1, \ \mathbf{h}_2\ ^2 = 1.5,$ and $\mathcal{P}/\mathcal{N}_0 = 1.$	38
3-1	Transmitter for TH precoding system.	46
3-2	TH precoding system using matrix model. The last box is a “slicer,” which implements nearest-neighbor detection on a modulo-extended constellation.	47
3-3	Example of modulo-extended constellation for 4-QAM.	48
3-4	Ergodic sum capacity for spatial precoding compared with other methods, when transmitting from an 8-element array to up to 8 receivers. Results are from simulations assuming independent Rayleigh fading at 5 dB SNR per link. The top curve represents a upper bound based on coordinated receivers.	54
3-5	Embedding of the symbol “00” for the example of Fig. 3-3.	57

3-6	Sum capacities, from simulations, for various stream orderings with $M = 8$ transmit antenna elements and $K = 8$ receivers. Waterfilling is used.	68
3-7	Outage probability, from simulations, for precoding of $K = 8$ streams from an $M = 8$ element array using various ordering methods. We use power control at an input SNR per link of 5 dB.	70
3-8	Outage probability, from simulations, for precoding of $K = 7$ streams from an $M = 8$ element array using various ordering methods. We use power control at an input SNR per link of 5 dB.	71
3-9	Precoding power loss, in terms of SNR gap from a scenario without any interference. The points shown are for uncoded A^2 -QAM input symbols.	73
3-10	Simple embedding for (a) very small or (b) very large interference. Possible interference values are shown as \bullet , and embedding points with \triangleright , $*$, \circ , and \times	75
3-11	Modified versions of the embeddings in Fig. 3-10 for more moderate-power interference.	76
3-12	Precoding power loss for the methods of Fig. 3-11, compared with using a fixed constellation that does not depend on the interference distribution.	77
3-13	Order of processing symbols for precoding	84
3-14	Simulated outage probability for received SNR for the second of two receivers. We compare the methods of of Section 3.3.2 and Section 3.3.3, with a 4-element array, 4 i.i.d. Rayleigh-distributed taps each of variance 0 dB, $\mathcal{P}/\mathcal{N}_0 = 0$ dB, and equal power distributed between the two receivers.	90
3-15	Block diagrams for DMT method with two receivers. Shown are the processing at the transmitter and at a typical receiver.	92
3-16	Ergodic sum capacity across two receivers for the different methods. Parameters are the same as for Fig. 3-14, except the input SNR, $\mathcal{P}/\mathcal{N}_0$, is made variable. The DMT method used 32 tones, and the rate penalty from the cyclic prefix is ignored.	93
4-1	Outage probability for an 8-element array transmitting to 8 receivers, with an SNR per link of 5 dB and power control.	103

4-2	Deterministic received SNR (per receiver) for a large system with $K/M = \beta$ and an input SNR per link of 5 dB.	104
4-3	Simulated outage curves when transmitting from a 4-element array to groups of 3 receivers using zero-forcing beamforming, at an input SNR per link of $\mathcal{P}/\mathcal{N}_0 = 5$ dB. Streams are partitioned into groups using a “greedy” algorithm.	106
4-4	Diagram of queueing model	108
4-5	Outage probabilities for a queueing system with $M = 8$ transmit antenna elements and $K = 8$ simultaneous receivers chosen from a window size varying from 8 to 20. We use zero-forcing beamforming, power control, and an input SNR per link of 5 dB. Shown for comparison is a bound on outage corresponding to orthogonal channel vectors.	109
4-6	Same as Fig. 4-5, but we now add similar curves for precoding, with power control and proposed max min ordering.	110
4-7	Ergodic sum capacity for zero-forcing beamforming and precoding, for the same simulation as in Fig. 4-5 except without power control.	111
4-8	Ergodic sum capacity for channel-aware timesharing strategies with various numbers of transmit antenna elements. The curves were computed using numerical integration over independent Rayleigh fading at an input SNR per link of 5 dB.	115
4-9	Ergodic sum capacity for zero-forcing beamforming and precoding with an 8-element array and large delay constraints. “Max sum proposed” uses the method of Section 3.2.2, while “Orthogonality info only” uses only orthogonality information, as in Section 4.3. Waterfilling across streams was used once the receivers were selected.	117
4-10	Ergodic sum capacity for zero-forcing beamforming and precoding with a 4-element array and large delay constraints. “Max sum proposed” uses the method of Section 3.2.2, while “Channel strength” uses only single-user SNR information. Waterfilling across streams was used once the receivers were selected.	118
4-11	Typical delay regions for certain two-receiver strategies, computed for a sample channel matrix realization. A $(\text{delay}_1, \text{delay}_2)$ point is achievable if it is on or <i>outside</i> (i.e., up and to the right) of the boundaries shown.	120

5-1	The curve shown describes a typical frontier of achievable received SNR pairs when transmitting common information from an 8-element array to two receivers, at an input SNR per link of 5 dB. SNR pairs are achievable if and only if they lie on or inside this transmitter operating characteristic (TOC). Various operating points of interest are also shown.	130
5-2	Expected average SNR per receiver for various values of M/K and an input SNR per link of 0 dB. The solid line shows the deterministic asymptotic values when both M and K go to ∞ with the ratio M/K held fixed. The dashed curves denote representative points corresponding to finite M and K for $K = 4$ (\diamond) and $K = 8$ (∇), from simulations.	134
5-3	Single-user ergodic capacity when multicasting a stream from an 8-element array to a number of receivers, at an input SNR per link of 5 dB. Also shown is a curve for a space-time code that does not make use of channel knowledge and achieves received SNR = $\ \mathbf{h}\ ^2/M \cdot \mathcal{P}/\mathcal{N}_0$.	136
5-4	Single-user outage probabilities when multicasting a stream from an 8-element array to a number of receivers, at an input SNR per link of 5 dB. Also shown are curves for a space-time code that does not make use of channel knowledge.	137
5-5	Possible structure for space-time multicast coding.	138
5-6	Weakest-receiver effective SNR, from simulations, when multicasting a stream from an 8-element array to 8 receivers, at an input SNR per link of 5 dB. The schemes shown are: “No Array”: single transmit antenna element; “No Tx Knowledge”: space-time coding with orthogonal matrix \mathbf{G} , “ST Multicast”: using channel knowledge with the weights chosen by the method of (5.18)	151
B-1	Sample embedding of 4-QAM inside large-order QAM interference. In this example, each “embedding constellation” surrounds four possible interference points.	164
B-2	Precoding power loss for various relative interference ratios.	165

List of Tables

2.1	Maximum achievable coded rates for sending distinct streams to two receivers using different multiplexing methods. All methods use the same symbol duration and bandwidth, and spatial multiplexing assumes a best-case scenario with orthogonal channel vectors \mathbf{h}_1 and \mathbf{h}_2 . The parameter α represents the fraction of time (for timesharing) or power (for CDMA or spatial multiplexing) devoted to the first receiver.	37
3.1	Summary of array processing algorithms for a variety of scenarios with M transmit antenna elements and K receiver antenna elements. Shown are the corresponding matrix factorizations, as well as a measure of performance in terms of the order of the Erlang distribution of idealized SNR for the i th stream. “Hybrid” refers to beamforming at the transmitter and interference cancellation at the receiver side.	62
4.1	Summary of data types, organized by delay tolerances	99
4.2	Summary of when timesharing strategies are sufficient for maximizing sum capacity.	114

List of Commonly Used Mathematical Symbols

The following symbols come up often throughout the thesis, and care has been taken to define them consistently. Other symbols come up less often, some of which may occasionally be used for more than one purpose.

\mathbf{H}, \mathbf{h}	Channel matrix, channel vector
K	Number of receivers
M	Number of transmit antenna elements
\mathbf{s}	Input data stream symbols
\mathbf{x}	Outputs of transmit antenna elements
\mathbf{w}	Additive noise at receivers
\mathbf{y}	Received symbols
\mathcal{P}	Maximum transmit power
\mathcal{N}_0	Noise power (at each receiver)
\mathcal{E}	Expectation operator
\dagger	Conjugate transpose operator
\ddagger	Conjugate transpose and time reversal operator

Chapter 1

Introduction

Wireless communication has been expanding at an impressive rate for some time now. Yet in this climate, engineers still struggle with fundamental questions about network architecture and the underlying physical limitations of communicating over airwaves. This thesis hopes to contribute to this discussion with improvements in the understanding and design of antenna array systems to address these issues.

Many wireless network architectures are amenable to the limited use of arrays. In cellular systems, mobile devices are divided among geographic cells and only communicate directly with a base station associated with their current cell. Wireless ad-hoc networks do not have such central control; a local set of devices is able to self-configure. Many times, however, it is still useful to route communications through a single node that has internetwork connectivity and lack of battery-life constraints. In these and other examples, users are divided into relatively simple, inexpensive devices and a smaller number of more powerful nodes. The latter type, with their less stringent constraints on power, size, and computation, become natural candidates for the use of a multiple-element array.

In this thesis, we consider such a model and focus on the interactions between a single array device and its associated wireless users. Furthermore, we concentrate on the less-understood “downstream” direction (that is, from the base station toward the various receivers). Since the receivers are battery-limited and typically do not have a great amount of coordination, responsibility for ensuring high rates and avoiding interference falls mainly on the transmitter and is the main subject of our research. Global issues such as handoff among base stations are important, but will be considered beyond the scope of this thesis. We will see that the single-array con-

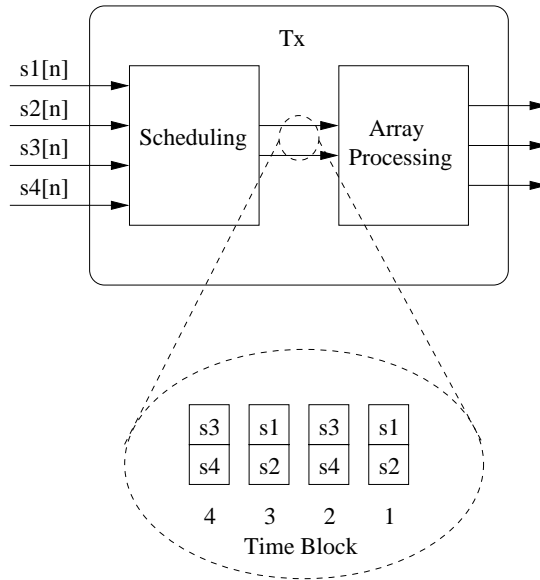


Figure 1-1: Block diagram of transmitter with an antenna array, illustrating scheduling and array processing system tasks.

figuration alone offers many opportunities for performance improvements as well as difficult design decisions.

A major reason for both the complexity and the potential of a transmitter antenna array is that it can direct data to multiple receivers simultaneously. Such a strategy affects many system components and is reflected in our transmitter system architecture, shown in Fig. 1-1. We use a packet-based, streaming data model, where some streams may be intended for individual receivers while other are common to more than one. Many types of data, such as voice, video, and file transfers, can be modeled in this way. We partition the processing of this data into two system tasks, denoted scheduling and array processing. The scheduler divides time into blocks and decides which data will be sent over each block. In the example shown, packets from the first two streams are sent in the first time block, etc. Once this has been decided, the transmitter must then map the data onto the physical antenna outputs in a way that will allow the receivers to understand the messages with sufficient fidelity. This is the function of the second task, which we call array processing, and can encompass multiple-input, multiple-output processing; modulation; coding; and other elements at the signaling level. Scheduling and array processing roughly correspond to the standard medium access control (MAC) and physical layers, although some elements

of both layers will be present in each of the two tasks.

The standard approach to these types of problems has been through layered protocols, where functions at different levels of abstraction are considered separately. For example, the networking community often concentrates on scheduling while assuming a reliable, interference-free channel. Array processing research, on the other hand, generally does not consider how the streams are selected or what their different properties may be. For array systems in particular, however, performance will depend strongly on the interaction among the data, scheduling, array processing, and physical channel. This suggests both a more comprehensive design process and greater integration, or at least awareness, among the different system components. Recently, there has been some interest in the 802.11 community in designing scheduling algorithms that are more aware of the physical channel and array processing (see [51] and references therein), though the emphasis for the most part has been on incremental upgrades of existing systems. In this thesis, we hope to develop a more complete understanding of scheduling, advanced array processing techniques, and their interactions as they relate to different system goals.

We investigate both scheduling and array processing with an eye toward helping the two tasks reinforce one another. An important part is incorporating knowledge, at both levels, of the state of the physical channel and the goals and destinations of individual data streams. Alternatively, we also consider tradeoffs in complexity between the two, where computation can be placed in one task or the other depending on implementation concerns. In many cases, a good portion of the potential gains are available when only one side incorporates a high degree of sophistication. For example, we adapt signaling-level precoding techniques to satisfy different kinds of data goals, and develop channel-aware scheduling techniques that enable high performance under lower-complexity choices for array processing.

1.1 Outline of Thesis

Chapter 2 lays the groundwork with an overview of several concepts related to transmitter antenna arrays. We discuss how elements of the fading channel model relate to the challenges and performance goals with which the rest of the thesis is concerned. Different signaling strategies lead to two basic performance criteria, outage probability and ergodic capacity, which are important to keep concrete and distinct. We

also provide motivation for scheduling several streams simultaneously and summarize some of the well-known array processing techniques on which later chapters build.

In Chapter 3, we focus on the array processing side while assuming a simple scheduler that divides streams into sequential or random groups. We primarily build upon the spatial precoding techniques described by Caire and Shamai [7] and Ginis and Cioffi [30], which in turn were adapted from precoding for intersymbol interference and information embedding. Recent results have shown that this family of techniques achieve the maximum sum capacity across all receivers (in [7] for the two-receiver channel, and [82, 71, 75] for any number of receivers).

We introduce precoding with a matrix formulation that emphasizes the connection to other strategies and makes evident various options and extensions. We then develop implementation aspects, such as robustness, constellation design, and meeting different types of performance criteria. For example, the maximum sum capacity solution can cause a large asymmetry in performance among receivers; we show how a modified order of operations results in a more equitable distribution. We conclude with a unified method of precoding for interference across both time and different streams and compare it to the multitone solution advocated in [30].

Chapter 4 shifts the focus to channel-aware scheduling and how it can improve performance. Such schedulers must be in tune with goals and constraints of the data streams; we develop algorithms for three data classes distinguished by their delay tolerance relative to certain physical parameters. Although further development is required before these algorithms can provide some standard quality of service guarantees, they do show some dramatic potential improvements. Especially promising is their ability to select subsets of streams that induce very low interference. In one example, under beamforming from an 8-element array, the medium-delay algorithm exhibits a 20 dB gain at 1% outage and more than double the ergodic capacity compared with a random grouping of streams. This places performance in the range of precoding, with much lower complexity at the array processing level. Because precoding systems start off better, scheduling can not provide as dramatic an improvement, but still pushes performance toward certain idealized limits and improves robustness.

In Chapter 5, we take a closer look at multicast scenarios where streams are intended for more than one receiver. In these cases, the scheduler and array processing can work together to transmit to all recipients simultaneously and avoid the redundancy of duplication. Unfortunately, benefits decrease as the number of re-

recipients grows and it becomes more difficult to direct the stream simultaneously to all of them. Using a single such stream for illustration, we describe a way to think about the balance of competing objectives in terms of efficient operating points. We then discuss methods that achieve these operating points, which we call space-time multicast codes, as well as more practical implementations. When the number of recipients is small or ergodic capacity is most important, we determine that beamforming strategies are a good choice. In the more general case, we show how to adapt ordinary space-time codes to this multicast scenario. In our example, these gain up to 6 dB at 1% outage over methods that do not use channel information and instead spread transmission out to all possible receivers. Furthermore, the channel information allows these multicast groups to fit more naturally into the larger picture of a system with heterogeneous sets of data and receivers.

We provide some concluding remarks and directions for future research in Chapter 6.

Chapter 2

Background on Transmitter Antenna Arrays

The recent interest in wireless communication has resulted in a large number of system models, algorithmic structures, and channel assumptions. In this chapter, we describe elements from our framework and introduce notation and concepts that will be used in later discussion.

We build up our channel model from a single link to timesharing to spatial multiplexing of multiple streams. Although we will mainly deal with transmitter arrays, the single link system is enough to illustrate different signaling approaches toward fading channels. This directly relates to the way we will classify data and judge performance throughout the rest of the thesis. We then introduce arrays, and quickly review some major issues and traditional array processing techniques. For a more comprehensive description of wireless communications systems, the reader is referred to Jakes' book [39] or the more recent review article by Biglieri, et al. [6].

2.1 Notational Conventions

Scalars are given by lowercase letters (a), vectors by boldface lowercase letters (\mathbf{a}), and matrices by boldface uppercase letters (\mathbf{A}). Certain constants or parameters are given by standard uppercase letters (A). When appropriate, explicit time dependences are shown using square brackets $a[n]$. Complex conjugation is denoted a^* , and \mathbf{A}^\dagger is the matrix Hermitian (conjugate transpose). Elements of vectors or matrices are denoted using subscripts (a_1 or $A_{1,3}$), with the first element indexed by 1. If a is a random

variable, then $\mathcal{E}[a]$ is its expectation.

2.2 Communication with Single-Element Antennas

When one talks about “wireless communication,” what is usually meant are electromagnetic information-bearing signals, transmitted and received from some kind of antennas, and propagating without waveguides. Therefore, they are subject to thermal noise, propagation loss that increases with distance, and interference from other wireless signals. Also important are the self-interference effects of reflections that depend greatly on the particular geometry of buildings, walls, and other objects in and around the path between the transmitter and receiver.

This last effect requires more discussion since it introduces a random element called fading that is the reason for much of the research in wireless communications. Reflections are received as multiple copies of the same signal, and cause different effects depending upon the difference in arrival times. If the receiver samples the signal quickly enough, the different arrivals will become resolvable as separate delays. In this thesis, however, we will usually assume a narrowband model with symbol-spaced sampling so that multipath arrivals are not resolvable. The arrivals can then combine constructively or destructively, resulting in amplitude variations. The maximum bandwidth to ensure this *flat fading* behavior is called the coherence bandwidth. Unfortunately, no exact formula exists to compute its value, although one rule of thumb is $1/\tau_{\text{rms}}$, where τ_{rms} is the RMS delay spread of the arrivals [54]. Observed values of this parameter vary, but some studies place it in the tens of nanoseconds for indoor environments, and on the order of a few millisecond for urban environments. Even when the fading is not precisely flat, many of our general findings still apply when receivers compensate with equalization techniques or the transmitters use more generalized precompensation such as discussed in Section 3.3.

The essential elements of this channel model can be expressed in the equivalent complex discrete-time baseband model (where all time dependencies have been suppressed)

$$y = h^* x + w, \tag{2.1}$$

where y is the received symbol, x is the transmitted symbol, h is the channel or fading

coefficient, and w is additive noise that encompasses thermal noise and any unmodeled background interference. All of the variables in (2.1) are complex-valued scalars. Unless specified otherwise, the additive noise w will be a zero-mean, independent, identically-distributed, circularly-symmetric Gaussian random sequence with variance \mathcal{N}_0 . Both transmitter and receiver are assumed to know the noise variance \mathcal{N}_0 , but not the particular realization w . Throughout this thesis, we enforce a constraint on the expected transmitted power, $\mathcal{E}[|x|^2] \leq \mathcal{P}$, and investigate how various scheduling and array processing approaches improve received performance. This constraint is meant to incorporate physical limitations, government regulatory issues, and the practical issue of keeping interference to a local set of receivers such as one cell in a cellular environment. (Wider network-level issues involving multiple transmitters are beyond the scope of this thesis.) Alternatively, one could use our results to achieve the received performance of current systems at reduced power.

The fading coefficient h itself is a random variable that, depending on the channel environment, can be modeled with various distributions. We will most often employ the commonly-used Rayleigh model, where the real and imaginary components of h have independent, zero-mean Gaussian distributions. Equivalently, the magnitude of h has a Rayleigh distribution (and its square magnitude has a distribution that is equivalently exponential, chi-square with two degrees of freedom, or first-order Erlang), while the phase has a uniform distribution. This is valid when there are a large number of scatterers and no direct line of sight between transmitter and receiver, and accurately models many indoor or urban environments. The *coherence time* is the duration over which h stays approximately constant. One popular model places the coherence time at about [54]

$$T_c = \frac{0.423\lambda}{\nu},$$

where λ is the wavelength of the signal and ν is the speed of the receiver. For example, the coherence time for a receiver traveling at 60 miles per hour with a 900 MHz signal will be about 6.8 ms. However, even with both transmitter and receiver are stationary, the fading will typically exhibit some time variation. Whether the fading stays constant or varies over a block of symbols depends on the physical parameters and signaling format. The current cellular and cordless phone standards DAMPS, GSM, and DECT use block durations on the order of hundreds of microseconds to several milliseconds, but sometimes also interleave over several blocks.

2.2.1 Performance Measures

There are two basic approaches toward communicating over fading channels. If the fading coefficient changes relatively slowly with time, then signaling can be performed within what is essentially a single fade of random quality. At the other extreme, one can signal across more and more fades and achieve an overall performance that typically becomes deterministic. Different performance criteria are appropriate for these two scenarios; we shall designate these criteria as outage probability and ergodic capacity.

In either case, the algorithms and performance that are available will also depend on whether one or both sides have knowledge of the fading coefficients. Receiver knowledge is a fairly common assumption and is possible through training, a separate pilot channel, and/or adaptive algorithms during the data phase itself. Most current wireless standards include mechanisms for this type of channel estimation. Consequently, we will assume perfect receiver knowledge unless specified otherwise.

By contrast, transmitter knowledge (also called side information) is typically more difficult to obtain and in some situations is considered to be less crucial. However, we will see that for multiple-receiver systems, this knowledge is very important to fulfilling the potential of the array. The transmitter can attain this side information in two ways. First, the receiver may relay its information through a separate feedback channel. Alternatively, if data is being exchanged in both directions over the same frequency band, such as in time division duplex (TDD) systems, then channel estimates made for the reverse channel will be valid in the downstream direction as well.

Characterizing performance by outage and ergodic capacity is not new, although most authors choose one form or the other. An exception is the diversity–multiplexing tradeoff expressed by Zheng and Tse [87]. Comparisons to our scheme may be useful to keep in mind, and will become clearer with the spatial multiplexing techniques of Section 2.3.2. However, care must be taken in understanding the different contexts in which the two frameworks come up. Zheng and Tse deal with a transmitter that does not have channel knowledge. As discussed above, this will lead to a different set of achievable operating points. Furthermore, we will see that this leads to very different ideas of outage and error. Secondly, we consider low-complexity, uncoordinated receivers, so that the performance at the individual receivers becomes as important as the aggregate total. Capturing this new tradeoff will be addressed throughout the

thesis.

Outage Probability

If the channel coefficient h is known at both sides and is constant for the time span of interest, then the fading channel model (2.1) takes the form of an additive white Gaussian noise channel with received signal to noise ratio (SNR)

$$\text{SNR}_{\text{rec}} = \frac{\mathcal{P}|h|^2}{\mathcal{N}_0}.$$

Since both coded and uncoded techniques for this channel are well-developed and depend only on this measure, we can capture the performance over random fading with an *outage probability* curve, which we define here as

$$\text{Pr}_{\text{outage}} \equiv \text{Pr} \{ \text{SNR}_{\text{rec}} \leq \text{SNR}_0 \}, \quad (2.2)$$

where SNR_0 is a parameter that can take on any nonnegative value, and is usually given in units of dB, equal to $10 \log_{10} \text{SNR}_0$. This curve is also equal to the cumulative distribution function (CDF) of received SNR over the fading channel ensemble.

The outage curve can be considered as a measure of the reliability of communication. For any target SNR_0 , the outage curve will show the probability that the target will be met. Perhaps more in tune with the goals of a system designer, the curve can also provide the appropriate SNR operating point if a target outage probability is to be met. Usually, a fairly small level such as 10% or 1% outage or lower is desired. For this reason, it may be equally or more important to have a probability distribution with short tails than one with a large mean. In the next section, we will see how the use of an array can concentrate the SNR distribution around its mean and thus create a more desirable channel.

For coded systems, a key quantity is the mutual information of the channel, which in this case evaluates to the rate

$$R = \log_2 \left(1 + \frac{\mathcal{P}|h|^2}{\mathcal{N}_0} \right) \quad (2.3)$$

in bits per channel use when given an optimal (Gaussian) input distribution. If the channel coefficient stays constant for long enough, this mutual information represents a maximum reliable rate of communication. In principle, this rate can then be ap-

proached using the same coding and shaping techniques that have been so successful in the additive white Gaussian noise channel, including trellis coding, turbo coding and shell mapping [22, 5, 43]. Therefore, we could have defined our outage in terms of a cumulative distribution on this rate instead of received SNR. We choose the SNR version because it is also valid for uncoded systems and because scaling by a different transmitted power \mathcal{P} will only result in a horizontal shift in the outage curve (when plotted in dB).

The instantaneous rate in (2.3) brings up an important difference between our model and one where the transmitter does not have knowledge of the channel coefficients. Without side information, the transmitter will not know at what rate it can reliably encode data. Outage probabilities are still well-defined, and it was in this context that they were first introduced by Ozarow, et al. [52]. Now, however, an outage event means a failure without the opportunity to lower the rate to a level that is known to be achievable. An alternate characterization, used by Zheng and Tse [87] as well as many other authors (e.g., [61, 34]) comes about from letting the transmitter choose a fixed modulation and coding scheme and then computing probability of bit-wise or codeword error over the ensemble of possible channel realizations. The error rate can be shown graphically for different transmitted powers \mathcal{P} . This graph will be very related to our outage curves because error events of this kind are generally dominated by low-quality channel realizations. However, we will tend to avoid this perspective because a transmitter that has channel knowledge will be able to adapt its modulation and coding scheme (or choose not to send at all) depending on the realized channel.

Ergodic Capacity

If the channel coefficient h varies ergodically over time, then one could signal across these variations and hope to achieve a reliable average performance. It turns out that this idea can be made precise for a variety of situations. We concentrate on coded performance here, although systems also exist that result in deterministic uncoded performance [80].

Consider a coded system where the transmitter has knowledge of h at each time instant. The system achieves the rate in (2.3) over each realization, resulting asymp-

totically in an average rate

$$C_{\text{ergodic}} = \mathcal{E} \left[\log_2 \left(1 + \frac{\mathcal{P}|h|^2}{\mathcal{N}_0} \right) \right] \quad (2.4)$$

that is deterministic. This performance, which we will call the ergodic capacity, depends only on the distribution of h and not on its particular evolution in time. We will see below that this rate is achievable even when the channel varies too quickly to send codewords within each individual fade. Some authors refer to ergodic capacity as the (average) throughput.

Unlike the outage probability curve, which is a distribution, the ergodic capacity results in a single number. For a given fading distribution, this number depends only on the input signal to noise ratio,

$$\text{SNR}_{\text{input}} = \mathcal{E} \left[\frac{\mathcal{P}|h|^2}{\mathcal{N}_0} \right].$$

From the concavity of the function $\log_2(1+x)$, ergodic capacity must be smaller than that of a static channel with the same input SNR, but the penalty turns out not to be too severe for most fading distributions. For example, with Rayleigh fading at an input SNR of 0 dB, the ergodic capacity is 0.86 bits/channel use, as opposed to 1 bit/channel use for a corresponding static channel.

Perhaps surprisingly, the same rate in (2.4) is achievable when the transmitter does not have complete channel knowledge, but knows only the statistics of h and the input SNR. This follows because the ergodic capacity can also be achieved using a constant-rate code, as long as the codeword symbols are interleaved across many channel realizations. Later we will find that with multiple-element transmit arrays, the ergodic capacity will become higher with side information than without.

Instead of the rate in (2.4), some authors define the capacity with side information to be a somewhat higher number achieved through a procedure called temporal waterfilling. To resolve this issue, recall that in our power constraint, a limit is placed on the expected power of each symbol $x[n]$. One might call this a *peak* power constraint (in the stochastic sense; a particular realized value of $x[n]$ may have power than is higher than \mathcal{P}). A somewhat looser, *average* power constraint would allow the transmitter to send some symbols with higher power than others, as long as the time average remains below \mathcal{P} . A transmitter with channel knowledge will then use more

power on stronger channel realizations, “pouring” power over the inverse of SNR [14],

$$\mathcal{P}[n] = \mathcal{P} \cdot \left[\lambda - \frac{\mathcal{N}_0}{\mathcal{P}|h|^2} \right]^+,$$

where the Lagrange multiplier parameter λ is chosen to satisfy the average power constraint and $[a]^+ = \max(a, 0)$. Waterfilling can be used to solve a variety of parallel channel problems, and will show up again later in this role.

2.3 Transmitter Antenna Arrays

Our main results consider a transmitter antenna array and multiple receivers, bringing an increased complexity to both the channel model and the different approaches a system may use.

See Fig. 2-1 for a diagram of the channel model with a three-element array and three receivers. In general, the transmitter now has M antenna elements from which it can send a vector of symbols, \mathbf{x} . These signals arrive at the K receivers through a cross-coupled channel, where the link between each antenna element and receiver is an independent Rayleigh channel of the type described in the previous section. If we collect all of the fading coefficients $H_{k,m}$ into a matrix \mathbf{H} , then this cross-coupled channel can be succinctly modeled as a matrix multiplication,

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{w}. \tag{2.5}$$

The power constraint now becomes $\mathcal{E}[\mathbf{x}^\dagger \mathbf{x}] \leq \mathcal{P}$, so that the maximum transmitted power is the same as with a single antenna element.

The inclusion of the cross-coupled channel has both positive and negative effects. First of all, the array provides multiple paths to each receiver, so that if one link undergoes a fade of poor quality, other links are likely to be better. In this way, a more reliable overall channel can be sustained. This is an example of *diversity*, which refers to taking advantage of multiple paths to a receiver. For this to work, however, it is important that the different copies be *independently* faded, or at least nearly so. Whether this is true depends on the physical separation of the antenna elements in array, the wavelength λ , and the location of scatterers. For indoor Rayleigh environments, for instance, the necessary separation between elements can be as small as $\lambda/2$. This diversity-centered model is not to be confused with phased array transmitters,




Figure 2-1: Channel model for transmitter antenna array and multiple receivers.

which operate in a regime where the coefficients have near-perfect correlation.

The potentially harmful effect is interference. A transmitted data stream will go to all receivers, whether intended or not. To deal with this, various scheduling and array processing techniques can be used. We begin with the simplest, which is to transmit to only one receiver at a time and therefore ignore any interference that is caused. Afterward, we will consider transmitting multiple streams simultaneously using array processing to mitigate interference, a process called spatial multiplexing.

2.3.1 Array Processing Techniques Under Timesharing

If the scheduler only selects one stream and one intended receiver at a time, interference becomes irrelevant. The array processor can then select a transmission scheme based upon outage or ergodic capacity performance criteria at the intended receiver, as well as other considerations such as complexity.

The array processor must specify the transformation from the data stream, $s[n]$, to the vector of antenna outputs, $\mathbf{x}[n]$, over the time block of interest. In general, this may include block processing and any kind of vector coded or uncoded modulation that satisfies the power constraint. It turns out, however, that optimal performance in this single-receiver scenario can be achieved by separating the modulation/encoding from the multiple antenna element considerations using a technique called beamforming.

Beamforming

Assume that the data stream $s[n]$ has been modulated and, if desired, encoded as if for a scalar additive Gaussian white noise channel. The array processor can then perform a linear transformation on each data symbol, $\mathbf{x} = \mathbf{g}s$ for some set of weights \mathbf{g} , such that the signals from the different antenna elements combine coherently at the intended receiver. This coherent combining results in the maximum possible received SNR over each realization, and is therefore optimal.

We can study the performance of this solution in more detail. If the receiver's vector of channel coefficients is \mathbf{h} , it effectively experiences an additive Gaussian white noise channel from $s[n]$ with a received SNR of

$$\text{SNR}_{\text{rec}} = \frac{\mathcal{P}|\mathbf{h}^\dagger \mathbf{g}|^2}{\mathcal{N}_0}.$$

This is maximized by matching the beamforming direction to the channel vector, $\mathbf{g} = \mathbf{h}/\|\mathbf{h}\|$, leading to the optimal value of

$$\text{SNR}_{\text{rec}} = \frac{\mathcal{P}\|\mathbf{h}\|^2}{\mathcal{N}_0}. \quad (2.6)$$

The probability distribution of (2.6) under Rayleigh fading is an M th-order Erlang (or, equivalently, chi-square with $2M$ degrees of freedom, denoted χ_{2M}^2). This has M times the mean of transmission from a single antenna element, with considerably smaller tails. The implications of this will become apparent shortly.

We plot ergodic capacity and outage probability for several scenarios in Fig. 2-2 and Fig. 2-3, respectively. For normalization, we define an “input SNR per link” as

$$\text{Input SNR per link} \equiv \frac{\mathcal{P}\mathcal{E}[|h_m|^2]}{\mathcal{N}_0}.$$

This value will usually be set at 5 dB in our examples, as this leads to reasonable coded rates in multiuser scenarios and is within the usual operating range given in the literature.

The ergodic capacity improves with the number of antenna elements, mainly because of the increase in mean received SNR. Once again, we see that the random channel variations often do not decrease ergodic capacity significantly. On the other hand, the shape of the fading distribution is very important when signaling over sin-




Figure 2-2: Ergodic capacities for an M -element transmit array and input SNR per link of 5 dB. We show the curve for Rayleigh fading, and for comparison the additive white Gaussian noise channel that does not encounter fading.

gle fading realizations. In Fig. 2-3, we see how this effect can dramatically affect the outage characteristic. At 1% outage, adding a second antenna element results in a gain of over 10 dB, even though the mean only doubles (3 dB). Note also the diminishing returns that are typical of diversity techniques; most gains occur as the first few antenna elements are added.

Space-Time Coding

The performance curves above require the transmitter to have knowledge of the channel parameters. Even if there is a small amount of uncertainty in the channel measurement, it turns out that beamforming is still optimal from the point of view of maximizing channel capacity [49, 72] or expected received SNR. However, when the transmitter does not have access to channel information, beamforming in any single direction results in the same distribution as with single-element transmission. We will




Figure 2-3: Single-receiver outage probabilities for a transmit antenna array with $M = 1, 2, 3,$ and 4 elements, using a Rayleigh fading model at an input SNR per link of 5 dB.

also see in Chapter 5 that channel information becomes less useful when a stream is intended for multiple receivers, because the transmitter can not direct data to all of them simultaneously. In these cases, more complex implementations may be useful and are often given the general heading of space-time codes.

The transformation between the data stream $s[n]$ and the antenna outputs $\mathbf{x}[n]$ can take a number of forms. One common element to space-time codes is that the covariance matrix $\mathcal{E}[\mathbf{x}\mathbf{x}^\dagger]$ has rank above one; the vector of antenna element outputs at a particular time contains information from more than one input symbol. In fact, the ergodic capacity is maximized by letting this covariance be a scaled identity [63]. Practical implementations include transformations resembling either convolutional [61, 34] or block encoders [1, 62]. In some special cases, as well as under idealized assumptions, these techniques are able to achieve performance equivalent to a received SNR distribution that is M th order Erlang, but they sacrifice a factor of M in mean SNR compared with beamforming under perfect channel knowledge.

	R_1	R_2
Timesharing	$\alpha \log_2 \left(1 + \frac{\mathcal{P} \ \mathbf{h}_1\ ^2}{\mathcal{N}_0} \right)$	$(1 - \alpha) \log_2 \left(1 + \frac{\mathcal{P} \ \mathbf{h}_2\ ^2}{\mathcal{N}_0} \right)$
CDMA	$\frac{1}{2} \log_2 \left(1 + \frac{2\alpha \mathcal{P} \ \mathbf{h}_1\ ^2}{\mathcal{N}_0} \right)$	$\frac{1}{2} \log_2 \left(1 + \frac{2(1 - \alpha) \mathcal{P} \ \mathbf{h}_2\ ^2}{\mathcal{N}_0} \right)$
Spatial multiplex (orth. channels)	$\log_2 \left(1 + \frac{\alpha \mathcal{P} \ \mathbf{h}_1\ ^2}{\mathcal{N}_0} \right)$	$\log_2 \left(1 + \frac{(1 - \alpha) \mathcal{P} \ \mathbf{h}_2\ ^2}{\mathcal{N}_0} \right)$

Table 2.1: Maximum achievable coded rates for sending distinct streams to two receivers using different multiplexing methods. All methods use the same symbol duration and bandwidth, and spatial multiplexing assumes a best-case scenario with orthogonal channel vectors \mathbf{h}_1 and \mathbf{h}_2 . The parameter α represents the fraction of time (for timesharing) or power (for CDMA or spatial multiplexing) devoted to the first receiver.

2.3.2 Spatial Multiplexing of Multiple Streams

The scheduler also has the option of sending multiple streams simultaneously. Interference then becomes an issue, but if it can be dealt with effectively, spatial multiplexing has several potential advantages. Among these are:

- *Increased Performance:* We illustrate the potential improvement using a coded system example where distinct streams are directed to their intended receivers using the type of single-user beamforming described above. In the best-case scenario where the rows of the channel matrix \mathbf{H} are orthogonal, the transmitter can send the streams simultaneously without incurring any interference. With this assumption, Table 2.1 compares the maximum achievable rates to two receivers for timesharing and spatial multiplexing, as well as a third technique, code division multiple access (CDMA), whereby the streams are modulated over linearly independent waveforms. (Actual CDMA systems usually operate in a wideband regime under different channel modeling assumptions, however.) We also plot these rate regions for a sample channel realization in Fig. 2-4.

It can be shown (using Jensen's inequality) that spatial multiplexing over orthogonal channels always results in the largest rate region, and that the disparity increases as the number of antenna elements and receivers grows larger. Looking at the formulas in the table, this improvement is reminiscent of that




Figure 2-4: Achievable rate regions for a sample channel realization with orthogonal channel vectors. $\|\mathbf{h}_1\|^2 = 1$, $\|\mathbf{h}_2\|^2 = 1.5$, and $\mathcal{P}/\mathcal{N}_0 = 1$.

achieved by increasing the bandwidth of a continuous-time channel, where the capacity with bandwidth \mathcal{W} is $\mathcal{W} \log_2(1 + \mathcal{P}/(\mathcal{N}_0\mathcal{W}))$. In fact, something very similar to this is occurring: the spatial multiplexing system is able to devote its full time–bandwidth resources to each receiver simultaneously, while timesharing and CDMA divide these resources up among the receivers. In the extreme case where the number of antenna elements and receivers (set $M = K$) grows large and $\alpha = 1/M$, the sum rate across receivers for spatial multiplexing becomes $M \log_2(1 + \mathcal{P}/\mathcal{N}_0)$. This dramatic, asymptotically linear increase with the number of antenna elements recalls similar results when the receivers are able to fully coordinate [63, 27].

Of course, realistic channel matrices will not often have orthogonal rows, but the above arguments provide motivation for investigating spatial multiplexing further. We will apply array processing (Chapter 3) and then scheduling (Chapter 4) to try to approach this performance.

- *Upgrade of Existing Systems:* Spatial multiplexing provides a method for increasing the number of receivers that a system can handle. In some cases, this can be implemented into current standards with relatively few alterations, and ideally requires only adding a few antenna elements and some additional processing to an existing array. Alternative ways to increase system capacity, such as purchasing additional spectrum or base stations, may be very expensive or difficult to bring about.
- *Flexibility:* A spatial multiplexing system can incorporate a great number of algorithmic and implementation options. For example, in many cases, most of the benefits of the array are available by adding sophistication to either the scheduling or array processing task. We will also see how to select and tune algorithms to meet the goals of different types of data streams. Design choices can be made based upon implementation issues and the different situations that are likely to come up, including the number and mobility of receivers.

To effectively use spatial multiplexing, the transmitter must deal with the issue of interference. One possible element of an interference-avoidance strategy, to be discussed in Chapter 4, is to design channel-aware schedulers that select groups of receivers with nearly orthogonal channel vectors. Even with this type of scheduler, the array processing block will likely need to compensate for some interference. In this thesis, we will concentrate on so-called “zero-forcing” schemes that remove all interference, leaving the receivers with only their intended signals and the additive white noise w_k . For the systems we consider, and the regimes in which they operate, this will lead to analyzable, relatively low-complexity solutions that perform nearly as well as optimal schemes. In Chapter 3, we present a detailed development of precoding techniques that are of this vein. For the moment, however, we briefly describe a well-known linear method for array processing.

Multiple-Receiver Beamforming

We look to extend beamforming, which was sufficient for optimality under timesharing, to deal with multiple receivers. Once again, assume that each stream has been modulated and, if desired, encoded as if for an additive white Gaussian noise channel. The vector of antenna element outputs can now be selected as a linear combination of the current symbols from all of the streams, $\mathbf{x} = \mathbf{G}\mathbf{s}$, where \mathbf{G} is called the

beamforming matrix. The channel model (2.5) now specializes to

$$\mathbf{y} = \mathbf{H}\mathbf{G}\mathbf{s} + \mathbf{w}. \quad (2.7)$$

If the symbols s_k are independent and zero-mean with variance \mathcal{P} , then the appropriate power constraint on \mathbf{G} is $\text{trace}\{\mathbf{G}^\dagger\mathbf{G}\} \leq 1$.

Assume for now that each of the elements in \mathbf{s} is intended for a separate receiver. If any streams were common to multiple receivers, the scheduler can simply duplicate them. We will return to more efficient methods of multiplexing common information in Chapter 5.

In selecting the beamforming matrix \mathbf{G} , there is an inherent tradeoff between increasing signal power and reducing interference. The zero-forcing approach is to eliminate interference by finding a \mathbf{G} for which $\mathbf{H}\mathbf{G}$ is diagonal. For independent Rayleigh fading, this can be done with probability one as long as the number of antenna elements in the transmitter array is at least as large as the number of receivers. The pseudoinverse produces the best such matrix in terms of maximizing the individual SNRs, and was used by Gerlach and Paulraj [29]. Unfortunately, by concentrating so much on interference, this solution can result in reduced signal power at the receivers. For randomly-chosen data streams, we essentially lose the effect of one of the transmitter antenna elements for every receiver that had to be nulled out.

Other useful beamforming strategies exist. One can optimize received signal power by setting \mathbf{G} proportional to \mathbf{H}^\dagger , often at the expense of high interference. A balance between this “matched filter” solution and zero forcing would be to maximize the signal-to-interference-plus-noise ratio (SINR). This is particularly useful when the interference is close to Gaussian distributed. Rashid-Farrokhi, et al. [55] found a solution (later refined by Visotsky and Madhow [73]) for reaching specified SINR levels at each receiver with the minimum total transmit power. Unfortunately, the form was of an iterative algorithm, and would require even more iterations to map it to a power constraint rather than SINR constraints. For the less ambitious problem of power control to equalize SINRs given a set of beamforming directions, an analytic solution was found by Yang and Xu [81].

2.3.3 Coordinated Versus Uncoordinated Receivers

Our basic model of a base station and several low-complexity, geographically separated receivers naturally makes it difficult for these receivers to achieve a large amount of coordination. Therefore, we have assumed that they have no knowledge of each other's received signals. Before going on, however, it may be useful to say a few words about what is possible when they do coordinate.

A variation on the model (2.5) would be for a single receiver to have access to all K antenna outputs. The usual application would be if all of the receive antenna elements were located within a single array. The purpose then is to simply communicate as much total information as possible, rather than dividing the information into separate streams for the different receivers. We briefly summarize some information theoretic results for coded systems.

When both transmitter and receiver know the channel matrix \mathbf{H} , the transmitter should send on the principle directions of \mathbf{H} and waterfill over the singular values [63]. Note that this requires both transmitter and receiver to use beamforming.

When only the receiver has channel information, capacity can be achieved when the elements of \mathbf{x} are i.i.d. over both space and time [63, 27]. The capacity is then asymptotically proportional to $\min(M, K)$ at high SNR. If $M = K$, then this represents an asymptotically linear growth in capacity with the number of antenna elements at each end, a result that has generated much excitement in the field. Simplified receivers that strip off and decode one layer of x_i at a time do not seem to lose much over the theoretical capacity [25, 3]. Recent results, though, have shown that the linear growth in $\min(M, K)$ at high SNR relies heavily on having perfect channel knowledge at the receiver and may not hold up to more realistic assumptions [41].

If neither the receiver nor the transmitter knows the channel, then i.i.d. symbols over time will not suffice. All information must now be contained in the correlations between symbols. This type of signaling, then, relies on the channel not changing too quickly, so researchers often choose a block constant fading model. This channel has been studied by Marzetta and Hochwald in [45] and subsequent papers that investigated specific coding schemes. A geometrical perspective is given by Zheng and Tse [86], including a study of the relationship between the length of the block fade and the number of antenna elements that can be used effectively.

We will find that, with the proper scheduling and array processing, systems without receiver coordination will often be able to achieve most of the ergodic sum capacity

that is possible with coordination. It is important to remember, however, that this is not our only goal. Systems of the type we examine must also consider, for instance, balancing the requirements of the individual data streams, directing them to single or multiple receivers, and doing this all with reasonable complexity and robustness.

Chapter 3

Precoding with Simple Scheduling

Our first in-depth investigation comes at the array processing level, as the transmitter attempts to direct multiple streams to their respective receivers simultaneously. Very recently, precoding-based approaches to this problem have appeared in the literature that show great promise [7, 30]. Yet much work remains in understanding their properties, performance, and implementations. In this chapter, we place precoding in perspective within a general matrix-based model, and investigate some of the design choices involved with different types of data, modulation, and channel models. In the process, we add several extensions and implementation algorithms to the basic precoding structure.

The main precoding algorithm, as applied to cross-coupled matrix channels, can be understood as a refinement of the linear zero-forcing approach described previously. Instead of diagonalizing the channel matrix (thus eliminating interference) in one step, precoding adds an intermediate triangularization. The residual interference is then dealt with using a more complicated operation that combines linear and nonlinear elements, and often results in much higher overall performance. For example, the ergodic sum capacity across receivers for precoding can be several times that of zero-forcing beamforming or timesharing. Even more, this general family of precoding algorithms has been shown to achieve the maximum sum rate of any method for this channel [82, 71, 75]. In Section 3.1, we describe this view of precoding and then characterize its performance and connection with other partitioned approaches such as BLAST [25].

Section 3.2 is concerned with issues that come up when applying precoding to systems. These include organizing the processing to meet different performance criteria,

finding low-complexity modulation techniques to eliminate interference, and a preliminary consideration of robustness to imperfect channel information. By exploring such issues, we hope to begin bridging the gap between describing what is possible and addressing design choices for particular systems.

In Section 3.3, we generalize precoding to compensate for interference across both different streams and time. By considering a matrix transfer function, we determine the types of processing that should be done. We find that there is more than one possibility, depending on the ordering of the interference cancellation that is to be done. We also compare our algorithms with the discrete multitone-based method of Ginis and Cioffi [30], which converts the matrix intersymbol interference channel into a number of parallel flat channels with only multiuser interference.

As a final note, the discussions of this chapter should be taken in two ways. First is the spatial precoder's value in dealing with the narrowly-focused array processing problem at hand. Secondly is its use as one of many building blocks within a larger system, where a large number of streams are communicated with different requirements over time-varying channels. We will deal more with this second, higher-level view as we consider the impact of scheduling later in the thesis.

3.1 Precoding for Multiuser Communications

In this first section, we bring together results on precoding using a framework that emphasizes partitioning and matrix-based operations. Our development proceeds through the elements of such a system, from linear processing to multidimensional coding techniques. In a natural way, it highlights the importance of ordered interference, the range of precoding options that are available for a general multiple-receiver model, and how these relate to other types of array processing. We also set up results in later sections on implementation and combined multiuser and intersymbol interference.

Throughout, we assume a simple scheduling algorithm that selects random or sequential groups of streams for spatial multiplexing, and furthermore duplicates any streams that are intended for multiple recipients. We will consider more sophisticated schedulers in Chapter 4 and more efficient multicast approaches in Chapter 5.

3.1.1 Precoding for Triangular Channels

We first describe precoding for situations where the channel matrix is triangular. This allows us to apply existing results on layered interference, and will prove to be a vital step in dealing with arbitrary channel matrices. With this later use in mind, we formulate precoding in somewhat unorthodox terms as a matrix inverse intertwined with additional, nonlinear operations.

Precoding relies on an implied ordering in the way symbols or data streams interfere. Recall that in our channel model,

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{w}, \quad (3.1)$$

the channel matrix of fading coefficients, \mathbf{H} , represents the transformation from antenna array outputs to receivers outputs, before white Gaussian noise is added. If this matrix is lower triangular and \mathbf{x} is simply the vector of data stream symbols, then receivers only get nonzero power from their own stream and those indexed earlier within the vector \mathbf{x} . If the transmitter processes the streams in this indexed order, it will know a priori what interference is to be expected, and can precompensate for this known interference. This type of approach first appeared as Tomlinson-Harashima (TH) precoding [64, 36, 47] over the intersymbol interference (ISI) channel, where a single stream exhibits self-interference across time. More recently, researchers have used ideas from Costa's "writing on dirty paper" [13] to refine precoding and apply it to many other problems, such as information embedding and digital watermarking (see [84] and references therein). In most cases, this dirty-paper encoding and its various implementations [10, 20] can achieve the same coded rates as without any interference; i.e., had the off-diagonal elements of \mathbf{H} been set to zero. Caire and Shamai [7] and Ginis and Cioffi [30] then applied these ideas to the matrix channel with arbitrary \mathbf{H} matrix by introducing the additional triangularization step.

The intersymbol interference channel can be interpreted as a triangular matrix channel with special structure, and serves as a useful starting point for our discussion. Consider a discrete-time, linear time-invariant channel,

$$y[n] = h[n] * x[n] + w[n], \quad (3.2)$$

with a causal, monic, minimum-phase impulse response $h[n]$. If we convert the input, output, and noise sequences to vectors, then (3.2) can be written as a lower-triangular




Figure 3-1: Transmitter for TH precoding system.

matrix channel (3.1). For example, the convolution matrix \mathbf{H} for three data symbols and an impulse response of length 2 will have the form

$$\mathbf{H} = \begin{bmatrix} 1 & 0 & 0 \\ h[1] & 1 & 0 \\ 0 & h[1] & 1 \end{bmatrix},$$

where $h[0] = 1$ because the channel response was assumed to be monic.

Suppose that the transmitter uses uncoded A^2 -QAM modulation, where A is an even integer, and wishes to eliminate interference. (Extensions to odd A are straightforward.) The real and imaginary parts of each input symbol, $s[n]$, will therefore take on values from among

$$\{-(A-1)\zeta, -(A-3)\zeta, \dots, (A-3)\zeta, (A-1)\zeta\},$$

where ζ is a real constant chosen so that the transmitted symbols obey the power constraint. The TH precoding system of Fig. 3-1 has a feedback loop to determine what the interference would have been for each symbol, then subtracts this amount off to produce a net effect of zero interference. This subtraction can result in symbols with large energy, so a modulo operation is performed to correct for this. The receiver will also have to compensate for this correction, as we describe below.

To understand this system further, and to connect it to our matrix model, consider the function of the modulo operation. This box shifts the real and imaginary components of its input until both are in the range $(-A\zeta, A\zeta]$. In other words, it adds $2\zeta A \cdot m[n]$ to the input, where $m[n]$ is the unique complex integer such that the




Figure 3-2: TH precoding system using matrix model. The last box is a “slicer,” which implements nearest-neighbor detection on a modulo-extended constellation.

output is in the square region $\mathcal{A} = \{(-A\zeta, A\zeta] \times (-A\zeta, A\zeta]\}$, which we denote as the *fundamental region* of the complex plane with respect to this modulo. If $m[n]$ were known in advance, this addition could have been performed before the feedback path is subtracted, resulting in a modulo-equivalent version of the input,

$$\tilde{s}[n] = s[n] + 2A\zeta \cdot m[n],$$

a process known as constellation expansion. If we consider the entire vector of modulo-equivalent input symbols, then the remainder of the feedback loop is equivalent to a matrix inverse and the precoder takes the form shown in Fig. 3-2. Note that since we have assumed that the diagonal elements of \mathbf{H} , and therefore of \mathbf{H}^{-1} , are unity, the outputs of the precoder are in the same fundamental region as its inputs. Therefore, to first order, the precoder conserves the energy of the input symbols. We will see in Section 3.1.3 that under closer inspection, there is a “precoding power loss” that becomes noticeable for low-order modulation [23], but can be compensated for by allowing a small amount of interference through.

The received vector is a noisy version of the modulo-equivalent input, $\tilde{\mathbf{s}}$, rather than of the original input itself. To recover \mathbf{s} , the receiver needs to either perform another modulo operation prior to detection, or to use a slicer based on a modulo-extended constellation, as shown in Fig. 3-3. In either case, the receiver may make errors it would not have had the original inputs been sent over a noninterfering channel and without precoding. For example, this could happen if the “▷” symbol was sent and the noise had very strong, but nonnegative, real and imaginary components. Therefore, the equivalent noninterfering channel for a precoding system is not additive white Gaussian noise, but rather a “modulo noise” channel with somewhat different properties. This issue was studied by Wesel and Cioffi in [77] for precoding




Figure 3-3: Example of modulo-extended constellation for 4-QAM

of intersymbol interference channels. Once again, we will see in Section 3.1.3 that these differences can be overcome.

Since \mathbf{m} is not known a priori, the precoding will not actually occur in the order shown in Fig. 3-2, but rather row-by-row as an intertwined linear operation (multiplication by \mathbf{H}^{-1}) and constellation expansion (the addition by $2\zeta\mathbf{A}\cdot\mathbf{m}$). The recursive form for ISI channels, as in Fig. 3-1, comes about using a matrix factorization:

$$\begin{aligned}
 \mathbf{H}^{-1} &= \begin{pmatrix} 1 & 0 & 0 \\ h[1] & 1 & 0 \\ 0 & h[1] & 1 \end{pmatrix}^{-1} = \left\{ \begin{bmatrix} 1 & 0 & 0 \\ h[1] & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & h[1] & 1 \end{bmatrix} \right\}^{-1} \\
 &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -h[1] & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ -h[1] & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (3.3)
 \end{aligned}$$

where each matrix multiplication represents one time through the loop. At each stage, the next element of \mathbf{m} is determined based on the symbols that were previously precoded. Because of the factorization given in (3.3), the memory only needs to be as long as the channel length.

A similar row-by-row procedure applies for any finite-size, lower-triangular matrix \mathbf{H} (with non-zero diagonal entries). The transmitter chooses the constellation expansion parameters, \mathbf{m} , such that $\mathbf{H}^{-1}\tilde{\mathbf{s}}$ is in the fundamental region \mathcal{A} . The

original minimum-phase and stability restriction for ISI channels, which ensured that the process remains stable, are not necessary for finite-length data vectors. If the diagonal entries of \mathbf{H} are not all one, their value may be factored out into a separate diagonal matrix, where they can contribute directly to the SNR of the channel.

The performance of the TH precoder can be contrasted with that of a purely linear array processor that also eliminates interference. In this second case, the transmitter could simply send $\mathbf{x} = \mathbf{H}^{-1}\mathbf{s}$. However, the transmitted power, assuming an i.i.d. input vector of length M , becomes

$$\mathcal{E}[\|\mathbf{x}\|^2] = \text{trace} \{ (\mathbf{H}^{-1})^\dagger \mathbf{H}^{-1} \} \mathcal{E}[\|\mathbf{s}\|^2] \quad (3.4)$$

Note that $\text{trace} \{ (\mathbf{H}^{-1})^\dagger \mathbf{H}^{-1} \}$ is the sum of powers of the elements in \mathbf{H}^{-1} . In other words, with channel inversion, *all* of the elements of \mathbf{H}^{-1} contribute to magnifying the transmitted energy, while in precoding, only the diagonal elements do (again, to first order). We could also write the above equation (3.4) as

$$\mathcal{E}[\|\mathbf{x}\|^2] = \sum_{n=1}^M \frac{1}{\sigma_n^2(\mathbf{H})} \mathcal{E}[\|\mathbf{s}\|^2],$$

where $\sigma_n(\mathbf{H})$ are the singular values of \mathbf{H} . This shows that as the matrix \mathbf{H} gets close to singular, the increase in energy over the precoding solution can become very large.

3.1.2 Precoding Over Arbitrary Channel Matrices

We now concentrate on the more important issue, that of communicating over arbitrary matrix channels. In Chapter 2, we discussed linear solutions using a beamforming matrix \mathbf{G} to diagonalize the channel. However, at least in the case of a triangular matrix, precoding can be much more efficient. Unfortunately, in the precoding system of Fig. 3-2, a triangular matrix was crucial to providing an ordered, layered structure to the interference. For arbitrary channel matrices, we therefore follow [7] and [30] in proposing a two-step solution, to first convert the $K \times M$ matrix (where $K \leq M$) into a triangular channel, and then apply the precoding algorithm of the previous section.

This two-step solution takes the form of a matrix factorization. Instead of using a single \mathbf{G} matrix to remove interference, we use separate beamforming and precoding

parts, \mathbf{G}_B and \mathbf{G}_P , such that the combined effective channel $\mathbf{H}\mathbf{G}_B\mathbf{G}_P$ is diagonal. Putting this all together, and factoring out any power control into a diagonal matrix \mathbf{D} , we organize the system as follows:

Algorithm 1 (Zero-Forcing Precoding) *Consider transmission from an M -element array to K uncoordinated receivers (with $K \leq M$) with a given matrix of fading coefficients \mathbf{H} . A precoding solution that results in no interference is*

$$\mathbf{y} = \mathbf{H}\mathbf{G}_B\mathbf{G}_P\mathbf{D}\tilde{\mathbf{s}} + \mathbf{w}, \quad (3.5)$$

where $\mathbf{H}\mathbf{G}_B\mathbf{G}_P$ is designed to be diagonal and

- $\tilde{\mathbf{s}} = \mathbf{s} + 2A\zeta \cdot \mathbf{m}$ is the modulo-equivalent vector of symbols to be transmitted. We assume the constellation is chosen so that the transmitted symbols satisfy the power constraint.
- \mathbf{D} is a diagonal matrix with diagonal elements d_k controlling the amplitudes sent to each receiver. We apply the constraint

$$\sum_{k=1}^K |d_k|^2 \leq 1$$

- \mathbf{G}_P is a lower-triangular matrix describing the linear part of the precoding operation. The diagonal elements of \mathbf{G}_P are all 1.
- \mathbf{G}_B is the beamforming matrix consisting of orthonormal columns.

The \mathbf{G}_B and \mathbf{G}_P matrices can be easily computed using the $\mathbf{H} = \mathbf{L}\mathbf{Q}$ lower-triangular decomposition. Let $\mathbf{G}_B = \mathbf{Q}^\dagger$, and \mathbf{G}_P be a scaled version of \mathbf{L}^{-1} so that the overall product is diagonal.

Actually, there will usually be \mathbf{G}_B matrices that are not orthonormal yet still satisfy the other criteria, such as one derived from the $\mathbf{L}\mathbf{U}$ decomposition for a square \mathbf{H} matrix. However, given our insistence on a lower-triangular \mathbf{G}_P and zero interference, an orthonormal \mathbf{G}_B is sufficient to maximize received SNR.

Furthermore, an orthonormal \mathbf{G}_B makes it relatively easy to find a scaling factor ζ to satisfy the power constraint. This way, the beamforming operation leaves the total power of the precoded symbols unchanged. If we additionally use the approximation that precoding adds no energy (i.e., that the “precoding power loss” is negligible),

then one only needs to select a constellation such that \mathbf{s} satisfies the power constraint, without worrying about the precoding and beamforming at all.

The different scalings of the \mathbf{G}_B and \mathbf{G}_P matrices illustrate how precoding is more efficient than beamforming. To ensure (to first order) that neither one changes the power of the symbol vector, each column of \mathbf{G}_P is scaled so that the diagonal element is unity, while \mathbf{G}_B must be scaled down further so that the entire column has unit norm.

Assuming there are a finite number of receivers, then stability of the precoding system is only in doubt if \mathbf{H} does not have full row rank, i.e., if the receivers have linearly dependent channel vectors. For most fading models and $M \geq K$, this occurs with probability zero. The m_i coefficients can also be kept below some threshold by choosing not to transmit to particularly weak receivers.

3.1.3 Performance of Precoding

Evaluating the performance of precoding systems is complicated, and depends on the modulation, coding, and other signaling-level implementations that are used. We begin with a preliminary discussion on “idealized” performance, and later describe how to deal with various issues that cause actual performance to diverge from this.

Idealized Performance

At a basic level, the spatial precoding solution outlined above changes the arbitrary \mathbf{H} matrix into a diagonal matrix, $\mathbf{H}\mathbf{G}_B\mathbf{G}_P$. Thus, if we treat the modulo noise as Gaussian additive noise and neglect the effect of the precoding power loss, what results is a series of parallel additive noise channels. It is then straightforward to determine the SNRs of these parallel channels in terms of the \mathbf{LQ} factorization of \mathbf{H} . Recalling that we set the beamforming matrix \mathbf{G}_B equal to \mathbf{Q}^\dagger , the diagonalized

channel matrix becomes

$$\begin{aligned}
\mathbf{H}\mathbf{G}_B\mathbf{G}_P &= \mathbf{H}\mathbf{Q}^\dagger\mathbf{G}_P \\
&= \mathbf{L}\mathbf{G}_P \\
&= \begin{bmatrix} l_1 & 0 & \dots & 0 \\ 0 & l_2 & \dots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \dots & l_K \end{bmatrix},
\end{aligned}$$

where l_k are the diagonal entries of \mathbf{L} . The last equality results because \mathbf{G}_P was specifically chosen to diagonalize the product, and furthermore is lower triangular with diagonal entries of unity. Including the power control \mathbf{d} , the channel to the k th receiver takes the form

$$y_k = l_k d_k \tilde{s}_k + w_k, \quad k = 1, 2, \dots, K, \quad (3.6)$$

with received SNR equal to

$$\text{SNR}_k = \frac{\mathcal{P}|l_k d_k|^2}{\mathcal{N}_0}. \quad (3.7)$$

We will say that the system does not use power control if all of the d_k parameters are chosen equal to $1/K$.

Similarly, the idealized instantaneous capacity of a coded link to the k th receiver becomes

$$C_k = \log_2 \left(1 + \frac{\mathcal{P}|l_k d_k|^2}{\mathcal{N}_0} \right), \quad (3.8)$$

with a corresponding ergodic capacity of

$$C_{k,\text{ergodic}} = \mathcal{E} \left[\log_2 \left(1 + \frac{\mathcal{P}|l_k d_k|^2}{\mathcal{N}_0} \right) \right]. \quad (3.9)$$

As we will see below, these rates are achievable with more sophisticated dirty-paper encoding techniques, supporting our use of (3.7)–(3.9) as performance measures.

The transmitter can adjust the power control and ordering among the streams to satisfy particular criteria based upon individual-receiver or system-wide goals, outage or ergodic capacity. We will say more about these choices, and propose algorithms

appropriate for various situations, in Section 3.2.2.

To compare the performance of precoding with other methods, sum capacity strategies provide a good illustration and have been the subject of most research up until now [30, 8, 82]. In these cases, the transmitter should use a waterfilling power control policy over the parallel channels [14],

$$|d_k|^2 = \left[\lambda - \frac{\mathcal{N}_0}{\mathcal{P}|l_k|^2} \right]^+,$$

where λ is chosen to satisfy the power constraint. Using this optimal power control and, for precoding, the max sum ordering method proposed in Section 3.2.2, we show in Fig. 3-4 the ergodic sum capacity of several techniques in situations where an 8-element array communicates with up to 8 receivers. In these simulations, we assume an independent Rayleigh model whereby the elements of the channel matrix \mathbf{H} are i.i.d. complex Gaussian variables. With eight receivers, precoding achieves well more than double the throughput of a round-robin timesharing strategy. Smarter scheduling, as in [74], improves the throughput of timesharing only slightly compared with precoding. Linear array processing, in the form of zero-forcing beamforming, does well up to a point, but eventually degrades as the transmitter must send nulls to too many receivers. The top curve represents a bound on performance, showing the highest achievable rate when the receivers can coordinate their responses, using Teletar's system [63]. That precoding can get so close, at least at the selected input SNR level, suggests that having coordination at either the transmitter or receiver side is more important than having it at both sides.

Although not shown in the figure, our simulations also suggest that except at very low SNR, power control plays only a secondary in maximizing the sum capacity. This is in line with results for communicating over parallel channels in frequency (see, e.g., [12]). We will also see in Section 3.2.2 that a random ordering of streams causes some loss in sum capacity, but still performs well.

More care must be taken for situations where the individual-receiver outage is most important, because precoding often results in performance asymmetries among the various streams. This results from the triangularization step of the precoding algorithm (3.5), where the beamforming matrix \mathbf{G}_B must steer more nulls for some streams than others. Quantitatively, this is evident by looking at probability distributions of $|l_k|^2$ (from the idealized SNR (3.7)) over the random ensemble of channel




Figure 3-4: Ergodic sum capacity for spatial precoding compared with other methods, when transmitting from an 8-element array to up to 8 receivers. Results are from simulations assuming independent Rayleigh fading at 5 dB SNR per link. The top curve represents a upper bound based on coordinated receivers.

matrices \mathbf{H} . For our independent Rayleigh model with an M -element array and random ordering of K streams, we apply the \mathbf{LQ} factorization result quoted in [18] and see that $|l_k|^2$ will have a $\chi_{2(M-k+1)}^2$ distribution, or equivalently an Erlang distribution with $M - k + 1$ degrees of freedom. This means that the k th receiver has the same outage performance as in a single-receiver system with a transmit array of $M - k + 1$ elements, if we correct for the fact that it only gets a fraction $|d_k|^2$ of the total transmitted power. This still compares favorably with a system using zero-forcing beamforming, where all receivers get the weakest of these distributions (Erlang of order $M - K + 1$). However, increasing the effective order of this weakest receiver's distribution by only one or two could mean a dramatic improvement in outage (recall Fig. 2-3). Strategies that maximize sum capacity tend to only increase the asymmetry among receivers, so we will also address the issue of providing more equitable performance among receivers in Section 3.2.2.

Approaching the Idealized Performance

A straightforward TH precoding implementation will differ from the idealized performance for a number of reasons, which we have touched upon but summarize here:

- *Modulo noise:* The modulo operations modify the noise distribution from additive Gaussian noise into a modulo-noise channel. The constellation symbols adjacent to the boundary of the fundamental region \mathcal{A} will then have smaller decision regions, increasing the probability of error.
- *Precoding power loss:* The original symbols, \mathbf{s} , are usually chosen from a discrete set of points within \mathcal{A} , while the precoded symbols $\mathbf{G}_p \mathbf{D} \tilde{\mathbf{s}}$ will have a more uniform distribution over their fundamental regions. In most cases, this causes an increase in transmitted power.
- *Shaping gain:* With a QAM constellation, the constellation points are distributed over a Cartesian product of square regions \mathcal{A} . However, for maximum power efficiency, the transmitted symbols should instead be distributed over a higher-dimensional sphere [44]. The difference in transmitted power is quantified by a “shaping gain” that must be bridged to achieve optimal performance. The maximum shaping gain occurs at high SNR, where it is equal to $\log_2(\pi e/6) = 0.51$ bits per two dimensions.

These factors vary in importance depending upon the regime of operation, and can be addressed in different ways. For example, at high SNR, the first two issues become negligible and only a shaping loss remains. One can then adapt shaping techniques that were previously used for the intersymbol interference channel. Although the ISI coder of [42] is not appropriate for layered interference across streams, an alternate method, trellis precoding [24], was implemented by Yu and Cioffi [83] and shown to achieve reasonable shaping gains. For low-rate precoding, we introduce in Section 3.2.3 a method for reducing the precoding power loss in certain situations with structured interference.

Another option, potentially more complex but offering a more unified approach, is to apply recent methods from the information embedding community that are essentially implementations of Costa’s dirty-paper encoding [13]. These are reviewed in [84] and include quantization index modulation and nested lattices [10, 4, 20]. The remainder of this subsection will be a brief overview on how they apply to spatial precoding, a connection first made by Caire and Shamai [7].

Consider the transmission of a symbol s_k with interference b caused by previously-precoded streams. The precoder uses a lattice code consisting of a “coarse” sublattice and translates of this sublattice called cosets. For example, in the modulo-extended constellation of Fig. 3-3, a coset consists of all symbol points of one type, such as “▷”. As shown in Fig. 3-5, the embedding “quantizes” the interference signal to the nearest point in the coset selected by s_k . The transmitter then sends e , the difference between this quantization point and the expected interference.

Although the preceding example is just another description of TH precoding, we now add several elements to approach the idealized performance. If the distribution of the precoded symbol, e , is not already approximately uniform over the Voronoi region of the coarse lattice (the dashed box of Fig. 3-5b, but shifted to be have zero mean), then the transmitter can add a pseudonoise dither signal, known at both transmitter and receiver, to the entire lattice prior to embedding. The average transmitted power can now be easily computed from the Voronoi region. Next, more efficient transmission is made possible by coalescing several time instances of the embedding problem together and doing vector quantization. The use of good, higher-dimensional nested lattices simultaneously provides coding gain (by increasing the minimum distance in the fine lattice) and shaping gain (by making the Voronoi region of the coarse lattice more like a higher-dimensional sphere). Finally, the modulo noise and precoding power loss are overcome by a technique known as noise cooling or distortion compensation, which requires a few more words.

Recall that in many estimation problems, mean-square error can be improved by intentionally leaving in some interference. Similarly, the slicer error for precoding systems can be improved by shifting the balance between noise and interference. As described in [10] and [84], the encoder multiplies the interference it expects by a real constant α (less than or equal to one) before quantizing to it, therefore sending

$$e = \tilde{s}_k - \alpha b,$$

where again \tilde{s}_k is the modulo-equivalent message symbol. The receiver then multiplies its signal by α before slicing, producing

$$\begin{aligned} \alpha y_k &= \alpha(e + b + w) \\ &= \tilde{s}_k + [\alpha w - (1 - \alpha)e]. \end{aligned} \tag{3.10}$$




Figure 3-5: Embedding of the symbol “00” for the example of Fig. 3-3.

Now, the noise power has been reduced to $\alpha^2 \mathcal{N}_0$ at the expense of a new self-noise term of power $(1 - \alpha)^2 \mathcal{P}_k$, where \mathcal{P}_k is the average transmitted power of this stream. With optimal lattices in higher dimensions, this self-noise behaves as i.i.d. Gaussian noise, independent of the other terms in (3.10). The α for receiver k that maximizes the overall received signal to noise ratio is

$$\alpha_{k,\text{opt}} = \frac{\mathcal{P}_k}{\mathcal{P}_k + \mathcal{N}_0}$$

and increases this received SNR by one. In this way, the system achieves the idealized SNR of $\mathcal{P}_k/\mathcal{N}_0$. Distortion compensation is incorporated into our matrix formulation by multiplying the off-diagonal elements of row k of the precoding matrix \mathbf{G}_P by α_k . For example,

$$\mathbf{G}_P : \begin{bmatrix} 1 & 0 & 0 \\ G_{2,1} & 1 & 0 \\ G_{3,1} & G_{3,2} & 1 \end{bmatrix} \longrightarrow \begin{bmatrix} 1 & 0 & 0 \\ \alpha_2 G_{2,1} & 1 & 0 \\ \alpha_3 G_{3,1} & \alpha_3 G_{3,2} & 1 \end{bmatrix}.$$

Receiver k then just needs to multiply its input y_k by α_k .

Whether a system chooses to implement nested lattices and distortion compensation or the more simple TH precoding with shaping depends on the potential benefits and complexity. At high SNR, the modulo noise and precoding power loss disappear,

making distortion compensation becomes unnecessary; α just degenerates to one. At lower SNR, high-dimensional lattices with well-shaped Voronoi regions are necessary to make the self-noise term look Gaussian. This can lead to higher complexity and decoding delays. In Section 3.2.3, we will explore a different method of reducing the precoding power loss that is applicable to a few particular interference distributions that may come up in spatial precoding.

3.1.4 Improving Upon Zero-Forcing Precoding

In our matrix factorization approach to precoding, we assumed that the transmitter wished to create a diagonal effective channel, causing no interference. It has been recently shown, for the two-receiver case by Caire and Shamai [7] and for the general case by several authors [82, 71, 75], that the sum rate can be improved somewhat by allowing some amount of interference, and that this form exactly achieves the sum capacity of the channel. It is not known whether modifications of this solution can achieve the entire achievable region of rate K -tuples.

We will continue on with the zero-forcing (that is, no interference) version, for a number of reasons. First of all, it leads to easier computation and analysis; as of now, there is no known closed-form expression or provably convergent optimal iterative algorithm to compute the more general precoder. This is especially important because we are interested in the distribution of performance among receivers, not only the maximum sum rate operating point. Secondly, by considering precoding and the receiver-cooperation bound in Fig. 3-4 and other examples, it appears that the great majority of the benefit of multiple antenna elements and multiple receivers is attainable by zero-forcing precoding, at least in this SNR regime. Furthermore, Caire and Shamai showed that in the limits of high SNR (where interference matters more than any additive noise) and low SNR (where only one of the streams is sent with nonzero power), zero-forcing precoding also achieves the maximum sum rate.

3.1.5 Relation to Other Matrix Channel Problems

Our matrix factorization description applies not only to precoding systems, but also to a variety of other scenarios involving multiple antenna elements at both the transmitter and receiver sides. We will see that a wide variety of algorithms can be incorporated under this common framework. This process helps to categorize results

from the literature, place precoding within this larger context, and perhaps lead to new algorithms within this family.

The different channel scenarios under consideration all share the same mathematical channel model (3.1) but vary depending on whether there is coordination at the transmitter side, receiver side, or both. So far, we have looked exclusively at the case where only the transmitter elements can coordinate. The opposite might be true for the uplink direction, and a formal duality between this so-called multiple access channel and precoding has been recently shown [75, 71]. Below, we demonstrate how these and other techniques can be subsumed under the idea of diagonalizing the channel matrix using factors corresponding to two types of operations:

- *Linear*: Simple matrix multiplication, i.e., beamforming
- *Interference cancellation*: Intertwined matrix multiplication and nonlinear interference subtraction

For example, in precoding, the transmitter performs an \mathbf{LQ} factorization of \mathbf{H} , where the \mathbf{Q} operation is of the first type and \mathbf{L} is of the more efficient, second type.

Receiver-Side Coordination

Consider a situation where a transmitter sends an independent data stream from each antenna element to a coordinated receiver array. If there are at least as many receivers as transmit antenna elements, then with probability one, the receivers could remove interference using a single-step linear operation. This receiver-based beamforming takes the form of a left multiplication by a matrix \mathbf{G}_B ,

$$\mathbf{G}_B \mathbf{y} = \mathbf{G}_B \mathbf{H} \mathbf{s} + \mathbf{G}_B \mathbf{w}, \quad (3.11)$$

such that $\mathbf{G}_B \mathbf{H}$ is diagonal. A more efficient method, however, is to only triangularize the channel in this way, then use a final interference cancellation step. With analogy to precoding, the receiver detects and decodes the streams in the order implied by the triangularization, and at each step subtracts off the interference caused by previously-detected streams.

This two-step receiver has appeared in the literature in different contexts and under many names. For intersymbol interference channels, it is known as the decision-feedback equalizer; in multiple antenna-element wireless, the V-BLAST system [26];

and in CDMA (where \mathbf{H} represents the spreading sequences), successive interference cancellation [17]. As with precoding, the higher efficiency of the interference cancellation operation can be seen by considering the different constraints on two matrix factors. To avoid noise enhancement, each row of \mathbf{G}_B should have unit norm, while the interference cancellation factor has unit diagonal elements (since determining and subtracting off interference does not enhance the noise). Under the idealized assumption that previously-ordered streams are detected perfectly, the received SNRs and maximum coded rates of the streams correspond to those of precoding, with the modification that we now perform an \mathbf{LQ} factorization of \mathbf{H}^\dagger rather than of \mathbf{H} . Similarly, one can achieve the sum capacity by not requiring the linear factors to strictly triangularize the channel matrix [3].

The duality between precoding and receiver-side interference cancellation is apparent, and a choice between the two methods depends on where the burden of computation and coordination should lie within a system. There are important practical distinctions as well. For instance, a precoding system may fall short of the achievable performance by not using perfect dirty-paper encoding, while receiver-based array processing can fail if there is too much error propagation from previously-detected streams.

Coordination at Both Sides

When a system has coordination at both transmitter and receiver arrays, new possibilities open up. In addition to all the previous strategies, one could use an \mathbf{LQ} decomposition to perform beamforming at the transmitter and interference cancellation at the receiver. This is applicable for $K \leq M$ and achieves the same SNR or sum rate performance as precoding, but with the tradeoffs associated with receiver interference cancellation (such as having to deal with error propagation, but not extra modulo operations). More interesting, though, are different types of factorizations.

Interestingly, Teletar [63] showed that the sum rate is maximized by splitting the processing with a singular value decomposition,

$$\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\dagger,$$

where \mathbf{U} and \mathbf{V} are unitary and $\mathbf{\Sigma}$ is diagonal with nonnegative entries. Transmit beamforming is done with \mathbf{V} and receiver beamforming with \mathbf{U}^\dagger . A geometric inter-

pretation would be that we now transmit along principal directions rather than using a Gram-Schmidt (i.e., \mathbf{LQ}) decomposition.

Another interesting choice would be to perform precoding at the transmitter (to remove interference from earlier streams) and interference cancellation at the receiver (working in the opposite direction, to remove interference from later streams). This frees the beamforming to do single-user matched filtering. The precoding and interference cancellation operations amount to performing a Cholesky \mathbf{LL}^\dagger decomposition on \mathbf{HH}^\dagger . To be more specific, we assume i.i.d. precoded symbols and let the beamforming matrix be

$$\mathbf{G}_B = \frac{\mathbf{H}^\dagger}{\sqrt{\text{trace}\{\mathbf{HH}^\dagger\}}},$$

which will then satisfy the power constraint. Next, the precoder must make the effective channel matrix \mathbf{HG}_B triangular so that the receiver's interference cancellor can do its job. Since the precoding matrix itself must be triangular, the Cholesky decomposition is natural. The precoding matrix will be $(\mathbf{L}^\dagger)^{-1}$, but scaled such that the diagonal entries are one. It turns out that this \mathbf{L} is the same matrix as the \mathbf{L} from the \mathbf{LQ} decomposition of \mathbf{H} . When all of this is done, the final received SNRs apparently become

$$\begin{aligned} \text{SNR}_k &= \frac{\mathcal{P}|l_k|^4}{\mathcal{N}_0} \cdot \frac{1}{\text{trace}\{\mathbf{HH}^\dagger\}} \\ &= (\text{SNR}_k \text{ from precoding}) \cdot \frac{|l_k|^2}{\frac{1}{K} \sum_{i=1}^K \|\mathbf{h}_i\|^2} \end{aligned}$$

This can lead to some interesting SNR distributions, increasing the performance of the receivers with better channels. However, there are two main deficiencies of this method. First, the perfect information embedding that was assumed in the SNR computation above is not achievable, since distortion compensation (which involves changing the off-diagonal entries of the precoding matrix) will affect the signal power of each stream. Secondly, power control is more difficult to do because the beamforming matrix is not orthogonal.

Table 3.1 presents a summary of several of the different methods discussed here. Recall that from Fig. 3-4, coordination at only one side may actually achieve close to the same performance as coordination at both sides.

	No coordination among Rx	Coordination among Rx
No coordination at Tx	N/A	Beamforming alone $K - M + 1$ <hr/> V-BLAST — LQ of H^\dagger $K - M + i$
Coordination at Tx	Beamforming alone $M - K + 1$ <hr/> Precoding — LQ of H $M - K + i$	Hybrid — LQ of H $M - K + i$ <hr/> SVD — $H = U\Sigma V^\dagger$ Maximizes sum capacity

Table 3.1: Summary of array processing algorithms for a variety of scenarios with M transmit antenna elements and K receiver antenna elements. Shown are the corresponding matrix factorizations, as well as a measure of performance in terms of the order of the Erlang distribution of idealized SNR for the i th stream. “Hybrid” refers to beamforming at the transmitter and interference cancellation at the receiver side.

3.2 Implementation Issues

3.2.1 Overview of New Implementation Issues

Although the methods already discussed provide a theoretical basis for precoding, much work remains in the design of practical precoding solutions. For example, when adapting techniques derived from the information embedding literature, one must be aware of the different contexts in which the two problems come up. We summarize some of the key issues below:

- When information is sent to more than two receivers, a stream can be part of both an embedding and several hosts. This allows the transmitter to rearrange the ordering of streams to achieve different performance tradeoffs. Additionally, it may divide up the available power in a number of ways. We discuss these issues in more detail in Section 3.2.2, finding that the ordering and power control can play an important role.
- The would-be interference is not some arbitrary signal, but rather a linear combination of symbols from previously precoded streams. Therefore, it may have certain properties, such as a particular discrete distribution, that the precoder may be able to exploit. We look at precoding for some of these situations in Section 3.2.3.
- Precoding and information embedding often operate in different regimes due to the goals and constraints of their respective problems. Many times in embedding applications, one wishes to hide a small amount of information without a noticeable degradation in the host signal. To satisfy this maximum distortion constraint, embedding rates tend to be smaller than one bit per host dimension. By contrast, zero-forcing precoding does not cause any distortion in the earlier streams. Instead, we have a power constraint, which is often much larger and allows higher-rate transmission. This can lead to different types of modulation and encoding techniques and, as we have seen, different nonidealities in the precoding process itself that must be considered.
- To achieve high data rates for wireless applications, complexity can become a major issue. Ideally, both the transmitter and receiver should perform only simple operations. If the receivers are battery-operated, complexity there be-

comes even more important. Simple embeddings may be preferable over nested lattices.

- Because the precoder’s “host” signal involves interaction with a channel matrix, the algorithm depends critically on the transmitter having knowledge of these channel characteristics. Similarly, the receiver must have some information about the channel and the encoding. We discuss some of these issues in Section 3.2.4.

In the next few sections, we look at several system components that address one or more of these issues.

3.2.2 Ordering of Streams

The order in which the streams are precoded will have an effect on their associated receivers’ performance. This suggests the need for practical algorithms that match the ordering to specific performance goals. In this section, we concentrate on optimizing according to two basic criteria, sum capacity and individual-receiver outage.

Several authors have recognized the importance of this ordering, but so far detailed analysis and algorithms have been lacking. Caire and Shamai [7, 8] discussed this issue (for both zero-forcing and more general precoding) and stated the solution for two-receiver sum capacity. For larger numbers of receivers, Yu and Cioffi [82] proposed an iterative algorithm for approaching the sum capacity, but were not able to prove convergence nor a closed-form solution. They and others [71, 75] also discuss a rate region that encompasses all precoding solutions, but do not offer any additional algorithms for reaching specific operating points of interest. None of these works directly deal with optimizing single-receiver outage.

Consider first a random ordering of K streams. When the later streams are beamformed to avoid interference to earlier ones, they incur a loss in channel quality. As we have seen, the first receiver gets the full M th-order diversity, the second receiver $M-1$, and the K th receiver $M-K+1$. If we wish to transmit to each receiver reliably at a constant rate, however, we would prefer greater symmetry among these receivers. On the other hand, to maximize throughput (the sum capacity across all receivers), it will turn out that an asymmetrical distribution is better. In either case, the SNR distribution resulting from a stream ordering can be augmented by appropriate power control.

One basic constraint for all possible orderings can be stated as follows:

Proposition 1 *Before power control, the product of received SNRs is independent of the ordering.*

For a full-rank, square \mathbf{H} , this is equivalent to saying that the square magnitude of the determinant of \mathbf{L} from the \mathbf{LQ} decomposition is independent of the ordering. This is true because, using \mathbf{E} to specify the permutation matrix,

$$\begin{aligned} \mathbf{L} &= (\mathbf{E}\mathbf{H}) \cdot \mathbf{Q}^\dagger \\ \Rightarrow \quad \frac{1}{K} |\det \mathbf{L}|^2 &= \frac{1}{K} |\det \mathbf{E}|^2 |\det \mathbf{H}|^2 |\det \mathbf{Q}^\dagger|^2 \\ &= \frac{1}{K} |\det \mathbf{H}|^2 \end{aligned}$$

regardless of the permutation \mathbf{E} . The last equality follows because the determinant of a unitary matrix or permutation matrix has magnitude 1. For \mathbf{H} not square, we can create a square matrix with the same product of singular values by adding extra rows that are orthogonal to each other and the other rows of \mathbf{H} . For \mathbf{H} not full-rank, the product is always zero. A corollary from this proof is that this product of SNRs is also equal to the product of the square magnitudes of singular values of $\frac{1}{\sqrt{K}}\mathbf{H}$, which are the SNRs of the parallel channels used in the Teletar scheme (also before power control) where the receivers can cooperate. \square

Proposition 2 *Power control can only decrease the product of received SNRs.*

Say that the SNRs before power control (that is, sending an equal fraction of power $1/K$ to each receiver) are $\beta_k/K, k = 1, 2, \dots, K$. We then use a different power distribution to achieve the SNRs $|d_k|^2\beta_k$, where $\sum_{k=1}^K |d_k|^2 = 1$. Instead of taking the product of SNRs, we can look at the monotonic function $1/K$ times the logarithm of this number. Before power control, we get

$$\frac{1}{K} \log \left(\prod_{k=1}^K \frac{1}{K} \beta_k \right) = \log \left(\frac{1}{K} \right) + \frac{1}{K} \sum_{k=1}^K \log(\beta_k),$$

while after power control,

$$\frac{1}{K} \log \left(\prod_{k=1}^K |d_k|^2 \beta_k \right) = \frac{1}{K} \log (|d_k|^2) + \frac{1}{K} \sum_{k=1}^K \log(\beta_k).$$

Since the mean of the $|d_k|^2$'s is $1/K$, we can invoke Jensen's equality to show that the first computation is at least as large as the second. \square

We now go on to study two performance criteria in more detail.

Maximizing Sum Capacity

The goal here is to maximize the sum capacity to all receivers. The algorithm for a particular channel realization will be the same whether we consider instantaneous capacity or ergodic capacity, because in the latter case, one only hits a maximum by optimizing the sum rate over each realization. Furthermore, if the channels are all i.i.d. and vary ergodically over time, then maximizing the sum capacity at each time will also result in each receiver achieving the same average rate, thus also achieving a degree of "fairness."

As discussed earlier, for a particular channel realization and ideal embedding, the set of streams are effectively sent to their associated receivers through K parallel channels with rates $\log_2(1 + \text{SNR}_k)$, where SNR_k is the received SNR for stream k . This SNR is determined from the \mathbf{LQ} decomposition associated with a particular ordering of the rows of \mathbf{H} . Since in this chapter each stream has only a single receiver, we interchangeably talk about ordering streams or receivers. To maximize the sum rate, one must also use power control to waterfill across the different streams.

Some guidelines on ordering streams follow.

Rule 1 (sum capacity) *For two streams ($K = 2$), the one whose receiver has the larger SNR should be first.*

This was stated in [7], and a proof is given here in Appendix A. The same result holds when waterfilling is not used, and can be shown with a simple convexity argument.

When there are more than two streams, the optimal ordering is still unknown. However, one rule that must be followed is:

Rule 2 (sum capacity) *For $K \geq 2$ consider any two consecutive streams, indexed k and $k + 1$. When projected away from the first $k - 1$ receivers' channel vectors, the stream with the larger SNR of the two should be first, i.e., given index k .*

The ordering of the two streams under consideration will not affect the SNRs of the other $K - 2$ receivers. Applying the two-stream result, we can say that for every

possible way of splitting the available power between these two and all the others, waterfilling within each grouping will result in a higher sum rate with the ordering implied above rather than the reverse. By similar reasoning, this rule also holds when power control is not used.

This rule alone does not imply a unique ordering, however. The following “greedy” algorithm will satisfy the rules above:

Algorithm 2 (sum capacity) *Choose the receiver with the strongest overall channel to be receiver 1. Then, project all other channel vectors away from this direction. Choose the strongest one of these as receiver 2, and project all remaining channel vectors away from both receivers 1 and 2. Choose the strongest one of these as receiver 3, etc.*

The Matlab command $[Q,R,E] = \text{qr}(A')$ will produce this ordering. This was used, along with waterfilling, to produce the precoding curve in Fig. 3-4. Although this algorithm satisfies the rules given above, it is not always optimal. For example, let

$$\mathbf{H} = \begin{bmatrix} 0.03 & 1.1 & 0.04 \\ 1 & 0.01 & 0.02 \\ 0.8 & 0.8 & 0.1 \end{bmatrix}.$$

Leaving the streams in the given ordering is optimal for sum capacity, although our proposed algorithm would do otherwise.

To test the significance of the stream ordering, we plot simulated ergodic sum capacities in Fig. 3-6 for several possible algorithms. The proposed algorithm gains around 1 bit per channel use over a random ordering throughout most of the given range of SNRs. Apparently, concentrating higher performance toward a small number of receivers can have an impact. Along these lines, a lower-complexity approximation to this algorithm would be to simply order the receivers by their channel strengths, without regard to the interdependencies. For the simulation, this led to almost the same performance as the original algorithm. Reversing the order, from weakest to strongest (but still using waterfilling), causes a loss of up to an additional bit over the random ordering. Throughout, this simulation assumed that the receivers have identically distributed Rayleigh channel coefficients. The effect of ordering on capacity will be even greater if some receivers have stronger channels than others.




Figure 3-6: Sum capacities, from simulations, for various stream orderings with $M = 8$ transmit antenna elements and $K = 8$ receivers. Waterfilling is used.

Minimizing Individual-Receiver Outage

The sum capacity strategy potentially sacrifices the performance of some receivers in favor of the “greater good.” This is fine as long as sum capacity is of primary importance, or if the channel varies ergodically and receivers can tolerate performance fluctuations. In other situations, however, it may be more important for individual receivers to maintain strong rates through (almost) all channel realizations. This might be true in a non-adaptive uncoded system with constant-rate transmission, or if a strong sense of “fairness” across receivers is most important, or if the channel varies extremely slowly with time.

Minimizing outage calls for a more conservative strategy that maximizes the performance of receivers with the weakest channel vectors. Ideally, all receivers would achieve the same SNR, which from Propositions 1 and 2 would reach its maximum

value at

$$\text{SNR}_{\text{outage,ideal}} = \frac{\left(\frac{1}{K} |\det \mathbf{L}|^2\right)^{1/K}}{\mathcal{N}_0}, \quad (3.12)$$

that is, the geometric mean of the SNRs. However, this is not always possible, so we instead employ a kind of max-min strategy. Afterward, power control can be used to equalize the SNRs at the receivers, though at a lower level than the ideal of (3.12) (see Proposition 2).

To be more precise, this max min criterion says that the weakest performance should be maximized, and then, given this, the second weakest should be maximized, etc. Note that to find the ordering, it does not matter whether we consider SNR or $\log_2(1 + \text{SNR})$.

Once again, the exact ordering is unknown for K streams, but some insights can be developed:

Rule 3 (max-min) *For two streams ($K = 2$), project each channel vector away from the other. The receiver with the weaker result should go first.*

Note that the first receiver could still be the weaker of the two, even though it no longer has to project away from the second. In either case, the result is worse if the ordering is reversed.

Rule 4 (max-min) *For $K \geq 2$ consider any two consecutive streams, indexed k and $k+1$. Project each channel vector away from the other and those of streams 1 through $k-1$. The receiver with the weaker result should be first, i.e., given index k .*

This follows from the two-stream case because all other streams are unaffected. This also suggests a “greedy” algorithm, which obeys the rule above but which may not necessarily be optimal:

Algorithm 3 (max-min) *Project all receivers’ channel vectors away from every other. Choose as the last receiver the one with the strongest result. Next, project all remaining channel vectors away from each other. Choose as the second-to-last receiver the one with the strongest result, etc.*

This algorithm produces some interesting results when followed by power-control that equalizes all the received SNRs. Shown in Fig. 3-7 and Fig. 3-8 are simulated




Figure 3-7: Outage probability, from simulations, for precoding of $K = 8$ streams from an $M = 8$ element array using various ordering methods. We use power control at an input SNR per link of 5 dB.

SNR outage distributions for an 8-element array and 8 and 7 streams, respectively. In the second figure, which achieves a higher sum rate, the proposed algorithm has close to the same distribution as the ideal (but perhaps unattainable) goal of (3.12). For comparison, the ordering proposed to maximize sum capacity loses about 3 dB at 10% outage and 6 dB at 1% outage. Even greater gains are exhibited with 8 streams. As the number of streams is decreased from 7, the effect of the ordering will become less significant.

Through most of these cases, the outage curves from our algorithm are even *more* steep than for a single stream with 8-level diversity (not shown). This can be explained by a kind of averaging effect across the different receivers' channel qualities. Of course, maximizing the diversity in this way comes at some price in overall throughput. For the system shown, the maximum sum capacity (reached at 6 streams) is 10.6 bits per channel use, which is still more than double the 4.6 bits per channel use for a single receiver. On the other hand, the sum rate of our proposed sum capacity




Figure 3-8: Outage probability, from simulations, for precoding of $K = 7$ streams from an $M = 8$ element array using various ordering methods. We use power control at an input SNR per link of 5 dB.

algorithm with waterfilling has the even higher value of 12.6 bits per channel use (with 8 streams), which is only 2 below the bound with receiver cooperation.

The idea of the max min ordering is the opposite of maximizing sum capacity; we now boost the performance of the receivers that will have a more difficult time communicating by placing them in the more privileged early positions. This suggests an approximate algorithm of ordering the receivers in the reverse of their channel strengths, regardless of the potential interference. As shown in the simulation, this strategy improves significantly upon the max sum ordering as well. A random ordering (not shown) falls somewhere between these two.

3.2.3 Constellation Design to Reduce Precoding Power Loss

We discussed earlier how transmitted symbols often have higher average power after precoding than in the original constellation. One proposed method to overcome this

precoding power loss involved higher-dimensional lattice codes, distortion compensation, and dither [84, 10]. However, these techniques can add considerable complexity at both transmitter and receiver. This motivates a study of the importance of the precoding power loss as well as other methods of compensating for it.

Using the example of uncoded QAM modulation, we find that precoding power loss can vary by as much as 3 dB, depending on the particular alignment of signal and interference. Both the greatest variation and worst-case losses occur for very low-order modulations when there is only a single dominant interfering signal. For these situations, we investigate ways of manipulating the symbol constellation that attempt to ensure one of the better-case scenarios.

Bounds on the Precoding Power Loss

Suppose that a stream uses uncoded A^2 -QAM modulation, with constellation symbols spaced 2ζ units apart. We will measure the precoding power loss by analyzing how much more transmitted power is necessary for precoding than with a QAM constellation at the same distance and no interference. The normalized minimum squared distance $4\zeta^2/\mathcal{N}_0$ will be our baseline for performance, as this has an approximate correspondence with probability of symbol error, but is much easier to deal with. We ignore the overall effective channel gain ($l_k d_k$) here, as this only scales the output. Also note that with some algebraic manipulation, our results can be converted to instead compute the loss in minimum distance if the transmitted power is held constant.

Depending on the particular input symbol and interference signal, the precoded symbol can take on any value in the fundamental region $\mathcal{A} = \{(-A\zeta, A\zeta] \times (-A\zeta, A\zeta]\}$. To compute the precoding power loss for a particular interference realization, we should average over all possible input symbol values. It is straightforward to show that the average power for the input QAM symbol is smallest when the interference point is centered between four neighboring constellation points (which includes the case of no interference), while the worst case occurs when the interference coincides with a constellation point. Because the interference is discrete-valued, it could potentially always be at a best-case location (where we get the same performance as if there had been no host), a worst-case location, or it may take on many values in between. The range of average precoded symbol powers can be summarized as follows:

- Best-case interference (or none): $\frac{2\zeta^2}{3}(A^2 - 1)$




Figure 3-9: Precoding power loss, in terms of SNR gap from a scenario without any interference. The points shown are for uncoded A^2 -QAM input symbols.

- Uniformly-distributed interference: $\frac{2\zeta^2}{3}A^2$
- Worst-case interference: $\frac{2\zeta^2}{3}(A^2 + 2)$

The precoding power losses for the different interference possibilities are shown in Fig. 3-9. The graph indicates that the potential loss is only significant for very small constellations. Note that the pseudorandom dithering technique can be used to always ensure that the interference looks uniformly distributed. If this is not done and we encounter a worst-case alignment, the maximum precoding loss is 3 dB for 4-QAM modulation.

Considering these trends, we would like to know if something can be done to mitigate this loss for low-order constellations, in effect transforming a worst-case interference value to something better. Because the worst-case possibility comes about from discrete-valued interference and constellation points, it makes sense to try adjusting the constellation design.

Constellation Design for 4-QAM Interference

We concentrate on embedding two bits per complex interference sample because we saw that only very low-order constellations such as 4-QAM present a significant potential for improvement. Before the constellation design stage, we need to look at the distribution of the interference signal.

The interference signal will be some linear combination of the modulo-equivalent symbols sent to earlier-ordered receivers, $\tilde{s}_1, \dots, \tilde{s}_{k-1}$. The linear combination is specified by the off-diagonal entries of the \mathbf{G}_P precoding matrix, corresponding to the feedback loop in TH precoding. For later-ordered streams, this linear combination of symbols can take on quite a great many different values, and will tend toward the uniform distribution mentioned above (after the modulo is taken into account). On the other hand, earlier-ordered streams will only see a linear combination of one or two symbols, so the interference distribution will continue to look discrete, and may at times hit the worst case. These are the situations where careful constellation design is most needed and, due to the structure of the interference, where the most can be done.

2-Bit Signaling with 4-QAM Interference

Let us start with the simplest case, where both of the first two streams use 2 bits per complex symbol. Since the first stream sees no interference, its constellation will look like standard 4-QAM. The second stream will then see the first as interference, after a gain and phase shift. Similar interference could occur for later streams if the linear combination of earlier symbols heavily favors one of them over the others.

Assume that the transmitter phase-aligns the current input symbol with the 4-QAM interference. Consider first the two extreme cases of no interference and very large interference. In either case, we can just send the new symbol s as is, resulting in the received distributions shown in Fig. 3-10. Both result in the best-case performance of no precoding power loss. The large-interference case is the same as superposition coding [14], where the earlier message is strong enough that the receiver can determine what point was sent, subtract (modulo) that out, and then detect the second message. In this and the next figure, the precoded signal that was sent is the difference between the interference point and the nearest “quantizer” point corresponding to the desired symbol.

In practice, it is likely that something in between these two extremes will occur.




Figure 3-10: Simple embedding for (a) very small or (b) very large interference. Possible interference values are shown as \bullet , and embedding points with \triangleright , $*$, \circ , and \times .

To adjust for this, first consider the large interference case of Fig. 3-10b. We could relabel the constellation points and achieve the same, optimal minimum distance with slightly smaller interference, as shown in Fig. 3-11c. Interestingly, this mapping can be interpreted as a form of distortion compensation, where some of the “quantizer error” from a standard TH precoder is added back in the form of self-noise.

As the interference gets smaller, some of the constellation points merge, as in Fig. 3-11b. At some point, we would expect to go back to the no-interference method, shown in Fig. 3-11a. It turns out that this switch occurs when the interference has half the magnitude of the symbol to be embedded. If the possible interference points are spaced $2\zeta_I$ units apart, then straightforward calculations yield:

- Method of Fig. 3-11a

$$\mathcal{P}_k = \begin{cases} 2\zeta^2 + 2\zeta_I^2, & 0 \leq \zeta_I < \zeta, \\ 2\zeta^2 + (2\zeta^2 - 2\zeta_I^2), & \zeta \leq \zeta_I < 2\zeta, \\ \text{periodic,} & \text{consequent ranges of } 2\zeta. \end{cases}$$

- Method of Fig. 3-11b

$$\mathcal{P}_k = \begin{cases} 2\zeta^2 + 2(\zeta - \zeta_I)^2, & 0 \leq \zeta_I < \zeta, \\ 2\zeta^2, & \zeta_I \geq \zeta. \end{cases}$$

For a particular interference distribution, the better of these two methods would be




Figure 3-11: Modified versions of the embeddings in Fig. 3-10 for more moderate-power interference.

chosen. As shown in Fig. 3-12, using this adaptive constellation instead of standard TH precoding provides a gain of 1 to 3 dB over a relatively wide range of possible interference powers. These include typical scenarios without power control when there are two receivers and two or three transmit antenna elements (average relative interference powers of 1 and 0.5, respectively.)

This adaptive constellation does not need to significantly affect the complexity of the receiver, which must distinguish among the embedding points corresponding to different input symbols. Since all “o” points in Fig. 3-11c embed the same symbol value, the receiver can treat the whole center region as a single decision region in its slicer. The slicers for the constellations in Fig. 3-11b and c then form a continuum parametrized by an overall gain. The receiver will have to use a separate slicer for Fig. 3-11a, but it may be possible to determine which of the two to use based on the distribution of received data.

2-Bit Signaling in Larger-Order QAM Interference

The interference will not always consist of only four possible points. This is especially true because the first stream, which causes interference on later ones, will typically have the best channel quality and is therefore more likely to use higher-order modulation. The methods of the previous discussion can still be used, although the description is more difficult and less dramatic gains are possible. This is discussed in Appendix B.




Figure 3-12: Precoding power loss for the methods of Fig. 3-11, compared with using a fixed constellation that does not depend on the interference distribution.

3.2.4 Channel Information

Precoding systems of the type described in this chapter involve a unique set of processing and channel knowledge assumptions. In this section, we look at what information is needed at the transmitter and receiver, and how this affects the processing and performance.

Transmitter Processing

The transmitter needs to compute the beamforming and precoding matrices as well as appropriate information rates, so it needs to know the channel vectors of all the receivers to which it is currently communicating. We have discussed how this can be accomplished either through feedback from the receivers' knowledge or by training on the reverse channel in time-division duplex systems.

Errors in the channel estimation can lead to unintended interference at the receivers. The modulo operation makes an exact error analysis difficult, but some indicators of the sensitivity can be found. Recall that the beamforming step converts the channel matrix \mathbf{H} into a lower-triangular form using the \mathbf{LQ} factorization. If the true \mathbf{H} differs from the estimated value by a perturbation, then the corresponding perturbations in \mathbf{L} and \mathbf{Q} can be magnified by about $\kappa_2(\mathbf{H})$ [59], where $\kappa_2(\mathbf{H})$ is the condition number of \mathbf{H} ,

$$\kappa_2(\mathbf{H}) = \frac{\sigma_{\max}(\mathbf{H})}{\sigma_{\min}(\mathbf{H})}.$$

This means that the precoding step will be based on a different \mathbf{L} from the true one, producing interference at the receivers. This effect can be lessened by ensuring that \mathbf{H} is well-conditioned. One way to do this is to send to fewer than the maximum number of receivers. A different approach, discussed in Chapter 4, is to specifically select receivers that lead to well-conditioned channel matrices.

Another analysis technique is to use a model for the perturbation. Suppose that the true \mathbf{H} differs from the estimated value by a $K \times M$ matrix Δ_1 that has i.i.d. complex Gaussian elements. The received vector before the additive noise will be

$$(\mathbf{H} + \Delta_1) \mathbf{Q}^\dagger \mathbf{G}_P \mathbf{D} \tilde{\mathbf{s}} = \mathbf{L} \mathbf{G}_P \mathbf{D} \tilde{\mathbf{s}} + \Delta_2 \mathbf{G}_P \mathbf{D} \tilde{\mathbf{s}},$$

where $\Delta_2 = \Delta_1 \mathbf{Q}^\dagger$ is a $K \times K$ matrix whose elements are i.i.d. Gaussian with

the same variance as those in Δ_1 . Since the first term on the right-hand side is the intended output, the uncertainty in \mathbf{H} apparently adds a “noise” term, which for any particular input vector is Gaussian distributed. With perfect coding and embedding, the vector of precoded symbols $\mathbf{G}_p \mathbf{D} \tilde{\mathbf{s}}$ will look i.i.d. Gaussian, so the characteristics of this noise term for each receiver can be computed; its variance will clearly be proportional to the variance of the perturbation. Unfortunately, this does not provide a complete characterization, since the precoded symbols are dependent on the original symbols. For example, the constellation expansion in $\tilde{\mathbf{s}}$ will likely be largest when the previous precoded symbols happen to be largest, suggesting that the new noise term will be at its worst when there are large modulo terms.

To ensure numerical stability of the precoding algorithm, care must be taken in the choice of \mathbf{LQ} factorization. The straightforward implementation, basically the Gram-Schmidt procedure, can lead to a severe loss in orthogonality among the beamforming vectors [32]. The “modified Gram-Schmidt” procedure is more careful about internal scaling, and $\mathbf{Q}\mathbf{Q}^\dagger$ differs from identity by a matrix of approximate norm $\epsilon\kappa_2(\mathbf{H})$, where ϵ is the machine precision. Once again, we see that using better-conditioned channel matrices helps. A different \mathbf{LQ} algorithm using Householder transformations takes about twice the number of computations but achieves still better orthogonality (approximately ϵ from identity).

Receiver Processing

The relevant receiver processing consists mostly of locking on to the gain and phase of its symbol stream and then detecting the symbols with a slicer. To do this effectively requires estimating the complex effective channel gain, $l_k d_k$, or developing systems that work around this step. It is at this level that spatial precoding presents some unique channel knowledge and sensitivity issues.

Although this type of processing is common in digital communication systems, many of the usual methods may not be appropriate for spatial precoding. For example, cellular systems often use phase shift keying constellations that do not require fine gain estimation. However, the modulo-equivalent constellation points of precoding make gain control a necessity; even if we started with a 4-QAM constellation, the modulo-extended version would expand to a constellation of higher order. Point-to-point digital subscriber line (DSL) systems use a rather involved training phase that allows the receiver to estimate its channel and determine the rates and gains on the

various subchannels [2]. Once again, however, our channel is different in that these decisions must be made in a centralized manner based on all of the receivers' channels and not each individually.

We are therefore left with two choices: all channel information can be distributed to all receivers, or the transmitter can inform or train the receivers for their individual complex gains and rates. The second option appears to be easier and involve the transfer of less information. Because different constellations can look the same under the modulo operation, the transmitter could inform the receivers of their streams' modulation and coding scheme using a few highly-protected symbols, as is common for this type of header information (see, e.g., [21]). On the other hand, it may make sense to train the receivers on the complex gain so that they can do their own adaptive gain control and continue to adjust it as the channel varies slightly from the transmitter's estimates.

To see the importance of an accurate gain estimate, consider a modulo-extended 4-QAM constellation, where the constellation points have odd real and imaginary integer coordinates. The upper-right constellation point (the triangle in Fig. 3-3) will have real coordinate $4n + 1$ for some integer n . If the receiver multiplies its input by too large or small a gain before slicing, then it could cause an error in the modulo-extended slicer even in the absence of noise. For instance, for positive n , multiplying by a gain that is a factor

$$\frac{4n + 2}{4n + 1}$$

too large will cause an error. Note that this gets steadily stricter with more severe constellation expansions: 2, 6/5, 10/9, etc. This type of effect happens for a slicer on any higher-order constellation; the new wrinkle here is that the modulo makes a low-order constellation act like a higher-order one. This provides another argument for choosing well-conditioned channel matrices, since this will help limit the constellation expansion.

Spatial precoding does provide some immunity to the amplifier saturation problem that can occur in typical TH precoding systems. These issues come up because gain control is often done with linear amplifiers or other devices that only provide good results over a limited range of inputs. With TH precoding, the modulo-equivalent symbols that are received can sometimes be large and may cause the input to go beyond this range [9]. Fortunately for spatial precoding, the symbols that are most

likely to undergo large constellation expansion, belonging to the later-ordered receivers, are also attenuated the most by the effective channel. To be more specific, assume a random ordering and no power control. This means that the vector of precoded symbols, $\mathbf{G}_P \mathbf{D} \tilde{\mathbf{s}}$, will be of equal maximum power. Through beamforming and the actual channel, this vector will be multiplied by the lower-triangular matrix $\mathbf{L} = \mathbf{H} \mathbf{G}_B$. The first receiver will therefore get the first precoded symbol multiplied by l_1 , whose power was previously determined to have an Erlang distribution with M degrees of freedom. The second receiver will get a mixture of the first precoded symbol and its own. Its own precoded symbol will be multiplied by a smaller factor than before, with $M - 1$ degrees of freedom, but the other symbol will arrive with Erlang-distributed power with 1 degree of freedom. If the two symbols add up coherently, then the overall maximum power has the same distribution as at the first receiver. This continues: the k th receiver will get a superposition of its own symbol with an Erlang-distributed power distribution with $M - k + 1$ degrees of freedom, and $k - 1$ other symbols, each with first-order Erlang. Therefore, the maximum power will have the same distribution for every receiver regardless of its ordering placement.

3.3 Precoding for Combined Multiuser and Inter-symbol Interference

This chapter has been primarily about taking techniques previously used to combat intersymbol interference and applying them to cross-channel interference between streams intended for different receivers. One would expect that when both types of interference are present, a generalization of these methods should follow.

One approach, used by Ginis and Cioffi in [30], is to perform a discrete multitone transform (DMT) to convert the time-dispersive channels into number of parallel, one-tap channels. Multiuser precoding can then be performed on each of these sub-channels. We seek a more unified treatment, using precoding directly for canceling all interference. This will take the form of two separate algorithms, representing causal and noncausal processing, depending on the order in which interference is canceled. After we develop our single-tone algorithms, we compare them with the DMT-based methods.

3.3.1 Multiple-Receiver Dispersive Channels

In describing these more general channel models, we take some of our notation from the multiple-input, multiple-output (MIMO) model of [70]. A discrete-time M -input, K -output linear time invariant (LTI) system can be represented by a $K \times M$ matrix $\mathbf{H}(z)$, whose entries $H_{km}(z)$ are the z -transforms of the channel from the m th input to the k th output. If the regions of convergence all contain the unit circle, then a matrix Fourier transform $\mathbf{H}(e^{j\omega})$ can be similarly defined. We define the “ \dagger ” operator to perform a conjugate transpose and additionally reverse the time sequence, so

$$\mathbf{H}^\dagger(z) \equiv \mathbf{H}^\dagger((z^*)^{-1}).$$

One special type of MIMO system is called *paraunitary*. This means that

$$\mathbf{H}^\dagger(z)\mathbf{H}(z) = c \cdot \mathbf{I},$$

a scaled identity matrix. If $\mathbf{H}(z)$ is defined on the unit circle, then we can similarly define a *lossless* system as a causal, stable system for which

$$\mathbf{H}^\dagger(e^{j\omega})\mathbf{H}(e^{j\omega}) = c \cdot \mathbf{I}.$$

This is the MIMO analogue to an allpass filter. It turns out that when $\mathbf{H}(z)$ is defined on the unit circle, then the two equations above imply each other, so a lossless system is the same as a causal, stable paraunitary system.

The basic channel model, before any processing, is

$$\mathbf{y}(z) = \mathbf{H}(z)\mathbf{x}(z) + \mathbf{w}(z),$$

where $\mathbf{y}(z)$ are the channel outputs and $\mathbf{w}(z)$ represents a realization of the white Gaussian noise sequence. We will assume that all the entries in the channel matrix $\mathbf{H}(z)$ are causal and stable, but not necessarily minimum phase. For a precoding system, the antenna inputs $\mathbf{x}(z)$ are the precoded and beamformed symbols, so we have

$$\mathbf{y}(z) = \mathbf{H}(z)\mathbf{G}_B(z)\mathbf{G}_P(z)\mathbf{D}\tilde{\mathbf{s}}(z) + \mathbf{w}(z).$$

$\mathbf{G}_P(z)$ is the precoding matrix, chosen so that $\mathbf{H}(z)\mathbf{G}_B(z)\mathbf{G}_P(z)$ is diagonal, where

each diagonal element of this product consists of only a single tap that is not a function of z . This way, a receiver sees no interference across time or from other streams. We also set the diagonal elements of the precoding matrix $\mathbf{G}_P(z)$ to be monic, which moves all power control into \mathbf{D} .

The beamforming matrix, $\mathbf{G}_B(z)$, specifies the transformation between precoded symbols and antenna inputs. In our discussion on flat fading channels, we constrained the elements of this matrix to be single-tap filters. In our extended model, we now allow each “beamforming weight” to be an LTI filter. Therefore, each antenna element output will be a linear combination of filtered precoded symbols from the different streams. We will also impose an orthogonality constraint on the beamforming (as we did in Section 3.1.2 for flat fading channels), so that the beamforming matrix $\mathbf{G}_B(z)$ must be paraunitary with $c = 1$,

$$\mathbf{G}_B^\dagger(z)\mathbf{G}_B(z) = \mathbf{I}. \quad (3.13)$$


Note that this imposes an orthogonality across both time and different streams.

For flat fading channels, \mathbf{G}_P was made to be lower triangular. This was necessary so that the intertwined constellation expansion and matrix multiplication operations of the precoder could be performed recursively over the different streams. In the more general model, we must precompensate for interference across both streams and time. We will find that this leads to more than one type of constraint on $\mathbf{G}_P(z)$, each corresponding with a different sequence of interference cancellation operations.

3.3.2 Canceling Multiuser Interference with Causal Processing

The k th row of $\mathbf{H}(z)\mathbf{G}_B(z)$ determines the linear combination of symbols that the k th receiver would see from both its own and other streams, if the precoding step had been omitted. With precoding, the coefficients of every power of z of each entry of this row will multiply some precoded symbol. If all of the precoded symbols corresponding to nonzero coefficients in this row are known when the current symbol is ready to be processed, then the interference can be computed and subtracted off so that that the receiver will get only the desired symbol (or a modulo-equivalent version).

The ordering in which symbols are processed, with respect to both time and the different streams, will determine the necessary structure of $\mathbf{H}(z)\mathbf{G}_B(z)$. Suppose



(a) Causal algorithm of Section 3.3.2 (b) Noncausal algorithm of Section 3.3.3

Figure 3-13: Order of processing symbols for precoding

for now that one symbol from the first stream is precoded, then one from the next stream, etc., until all of the symbols at time $n = 0$ have been processed. Next comes the $n = 1$ symbol of the first stream, and so forth. A graphical view of this processing of symbols is shown in Fig. 3-13a.

For the k th stream to precode its current symbol using this procedure, it needs to know the past precoded symbols of all streams and the present precoded symbols of the streams with lower indices. This means that in the k th row of $\mathbf{H}(z)\mathbf{G}_B(z)$, the first k entries should be causal, and the last $K - k$ entries should contain only negative powers of z . Since $\mathbf{H}(z)$ is already causal, we just need to triangularize the set of zero-lag taps. This can be done by collecting them into a matrix $\hat{\mathbf{H}}$, performing an \mathbf{LQ} decomposition $\hat{\mathbf{H}} = \hat{\mathbf{L}}\hat{\mathbf{Q}}$ on that, and using $\hat{\mathbf{Q}}^\dagger$ as the beamforming matrix $\mathbf{G}_B(z)$. In this way, the beamforming matrix still ends up as a set of single-tap, zero-lag filters, even though this was not specified a priori.

The SNR performance of this method is straightforward to calculate, since only the diagonal, zero-lag terms of $\mathbf{H}(z)\mathbf{G}_B(z)\mathbf{G}_P(z)$ contribute to the received signal. Using the same reasoning as in Section 3.1.3, the received SNR for stream k (to first order with TH precoding, or exactly with optimal information embedding) becomes

$$\text{SNR}_k = \frac{\mathcal{P}|\hat{l}_k d_k|^2}{\mathcal{N}_0},$$

where \hat{l}_k is the k th diagonal entry of $\hat{\mathbf{L}}$.

Note that performance-wise, the terms of $\mathbf{H}(z)$ with negative-powers of z were

essentially ignored, with their energy getting canceled in the precoding. This is fine when most of the energy in the channel responses resides in the zero-lag terms, but this is not generally the case.

A solution to this problem for single-user ISI channels is to use a whitened matched filter front end to make the channel minimum phase. Given an allpass constraint, this filter forces as much energy as possible into the first tap. For the multiple-receiver ISI channel, we might like to make all matrix entries minimum phase, but unfortunately there is no way to filter all of the entries of $\mathbf{H}(z)$ independently. Even if this were possible, it is not clear that it will necessarily lead to the largest SNRs. What we need is a more general $\mathbf{L}(z)\mathbf{Q}(z)$ decomposition of $\mathbf{H}(z)$ that concentrates as much energy as possible to the front of the final responses.

3.3.3 Precoding with Noncausal Filtering

For a scalar channel with impulse response $h^\dagger(z)$, the whitened matched filter starts with a matched filter $h(z)$, resulting in the conjugate-symmetric response $h^\dagger(z)h(z)$. This is followed by a filter that makes the overall response minimum phase (and makes the combined filter allpass). These ideas can be extended to transmit arrays, and eventually, multiple users.

Let us start with a single-user example, with M transmit antenna elements. This receiver's channel model is

$$y_1(z) = \mathbf{h}_1^\dagger(z)\mathbf{g}_1(z)s_1(z) + w_1(z),$$

where the elements of $\mathbf{h}_1^\dagger(z)$ are assumed to be causal. Ignoring the paraunitary constraint (3.13) for now, the received signal energy is maximized by making $\mathbf{g}_1(z)$ proportional to a bank of matched filters,

$$\mathbf{g}_1(z) = \gamma_1(z)\mathbf{h}_1(z),$$

for some scalar filter $\gamma_1(z)$. This turns the vector channel into a scalar channel, with only a single scalar filter left to be determined.

The power constraint comes down to

$$\mathbf{g}_1^\dagger(z)\mathbf{g}_1(z) = 1,$$

which is now equivalent to

$$\gamma_1^\dagger(z)\gamma_1(z)\mathbf{h}_1^\dagger(z)\mathbf{h}_1(z) = 1. \quad (3.14)$$

Finding $\gamma_1(z)$, then, is equivalent to finding a whitening filter for a random process with autocorrelation $\mathbf{h}_1^\dagger(z)\mathbf{h}_1(z)$. From well-known results in statistical signal processing [11], if $\mathbf{h}_1^\dagger(z)\mathbf{h}_1(z)$ is factorizable, then it has a canonical form,

$$\mathbf{h}_1^\dagger(z)\mathbf{h}_1(z) = |c_1|^2 t_1^\dagger(z)t_1(z), \quad (3.15)$$

where $t_1(z)$ is causal, monic, and minimum phase. The technical conditions for factorizability are that both $\|\mathbf{h}_1(e^{j\omega})\|^2$ and $\ln \|\mathbf{h}_1(e^{j\omega})\|^2$ are integrable over $-\pi < \omega \leq \pi$. These conditions hold for many functions $\mathbf{h}_1(z)$ of interest, such as FIR and rational z -transforms. The constant c_1 can be found with

$$\ln |c_1|^2 = \frac{1}{2\pi} \int_{-\pi}^{\pi} \ln \|\mathbf{h}_1(e^{j\omega})\|^2 d\omega.$$

In general, the solution for $\gamma_1(z)$ in (3.14) is not unique, but can contain factors from both $t_1(z)$ and $t_1^\dagger(z)$. From a total SNR standpoint, any of these solutions would give equal performance. However, because we want to use this system for precoding, the equivalent channel $\mathbf{h}_1^\dagger(z)\mathbf{g}_1(z)$ should also be causal and minimum phase. This means that any non-minimum-phase factors must be removed, so we set

$$\mathbf{g}_1(z) = \frac{\mathbf{h}_1(z)}{c_1^* t_1^\dagger(z)} \quad (3.16)$$

and get the equivalent channel

$$y_1(z) = c_1 t_1(z) s_1(z) + w_1(z). \quad (3.17)$$

We call this solution “noncausal” because the filter in (3.16) is in general not causal.

At this stage, no optimality has been lost by using this vector whitened matched filter. In fact, a frequency-domain version of this type of single-user, transmit array processing was derived by Zangi and Kransy [85], but without the minimum-phase constraint. Instead of precoding, they assumed an optimal receiver and showed that this system reaches the channel capacity if waterfilling across frequency is also performed.

We, on the other hand, choose to do transmitter precoding, so performance is determined by the zero-lag term of the equivalent channel. Since $t_1(z)$ is monic, the received SNR is $c_1^2 \mathcal{P}/\mathcal{N}_0$. Before we go on to multiple-receiver generalizations, it is useful to go over an example, and also ask whether this method is a significant improvement over the causal method of the previous section.

Comparison with Causal Precoding

As a simple example, take a two-antenna system, with monic channels

$$\mathbf{h}_1^\dagger(z) = \begin{bmatrix} 1 - \alpha z^{-1} & 1 - \beta z^{-1} \end{bmatrix}.$$

Let the input SNR be $\mathcal{P}/\mathcal{N}_0 = 1$. The causal processing method of Section 3.3.2 ignores the z^{-1} terms for the beamforming part and will simply combine the two channels, each with weight $\sqrt{2}/2$, to get the composite channel

$$\sqrt{2} \left(1 - \frac{\alpha + \beta}{2} z^{-1} \right).$$

Since the performance is determined by the zero-lag term, this system will always have a received SNR of 2, regardless of the values of α and β .

The noncausal solution of this section instead first performs a matched filter and attempts the spectral decomposition of (3.15). Expanding this formula out, we get

$$\begin{aligned} \mathbf{h}_1^\dagger(z) \mathbf{h}_1(z) &= -(\alpha + \beta)z + (2 + |\alpha|^2 + |\beta|^2) - (\alpha + \beta)z^{-1} \\ &= |c_1|^2 (1 - dz^{-1})(1 - d^* z) \end{aligned}$$

for some constant d . Since it is $|c_1|^2$ that determines the SNR, we solve for it algebraically:

$$|c_1|^2 = \frac{2 + |\alpha|^2 + |\beta|^2}{2} + \frac{1}{2} \sqrt{4 + |\alpha|^4 + |\beta|^4 + 2|\alpha|^2|\beta|^2 - 8\mathcal{R}e\{\alpha\beta^*\}}. \quad (3.18)$$

We see immediately that if $|\alpha|^2 + |\beta|^2 \geq 2$, the noncausal method will perform at least as well as the causal method described earlier, and will be much better as $|\alpha|^2 + |\beta|^2$ increases. This is not surprising, because it means that at least one of the two channels was not minimum phase, so pushing energy toward the beginning of the responses will help the zero-forcing precoder. At first this situation may seem trivial, since the

transmitter could individually filter the two channels to be minimum phase. However, when we move on to multiple receivers, recall that in general, it is not possible for the transmitter to make all of the channels for all of the receivers to be minimum phase.

When $|\alpha|^2 + |\beta|^2 < 2$, then for the noncausal method to be better, we need whatever is under the square root sign in (3.18) to be larger than $[2 - (|\alpha|^2 + |\beta|^2)]^2$. Subtracting this number from what is under the square root sign, we get

$$4|\alpha|^2 + 4|\beta|^2 - 8\mathcal{R}e\{\alpha\beta^*\} = 4|\alpha - \beta|^2,$$

which is always nonnegative, and equal to zero only when $\alpha = \beta$, that is, when the two channel vectors are the same. Therefore, the noncausal precoding never does worse than the causal method, and almost always does better.

Even when both channels are minimum phase, it turns out that the received SNR of the noncausal method can be higher by as much as a factor of two. (This happens when α and β are near the unit circle and differ in phase by π .) It can also be shown that this noncausal filtering method does no worse than the causal method for any two-antenna, two-tap system, whether or not the channels are monic or minimum phase.

Multiple Streams

Once again, (3.16) provides the filtering to be done on the first stream. What must be done with the second stream? If we use the same method, then this will cause interference to the first receiver, which should be avoided.

Conceptually, we can perform a similar procedure to what was done with flat channels and make use of a kind of LQ decomposition. Previously, this amounted to the Gram-Schmidt procedure of finding a set of orthogonal vectors \mathbf{Q} that span the same space as those in \mathbf{H} . Now, instead of letting $\mathbf{g}_2(z)$ be proportional to $\mathbf{h}_2(z)$ as for the first receiver, we need to make it proportional to the component of $\mathbf{h}_2(z)$ that is orthogonal to $\mathbf{h}_1(z)$. From linear algebra, this component can be written as

$$\mathbf{h}_2(z) - \frac{\mathbf{h}_1^\dagger(z)\mathbf{h}_2(z)}{\mathbf{h}_1^\dagger(z)\mathbf{h}_1(z)}\mathbf{h}_1(z).$$

(One can also think of the above as a separate orthogonalization for each frequency.)

We now need to normalize this function and make the overall response minimum

phase. With a little algebra, the autocorrelation is shown to be

$$\frac{\left(\mathbf{h}_1^\dagger(z)\mathbf{h}_1(z)\right)\left(\mathbf{h}_2^\dagger(z)\mathbf{h}_2(z)\right) - \left(\mathbf{h}_1^\dagger(z)\mathbf{h}_2(z)\right)\left(\mathbf{h}_2^\dagger(z)\mathbf{h}_1(z)\right)}{\mathbf{h}_1^\dagger(z)\mathbf{h}_1(z)}.$$

We know from the previous subsections that the denominator has the canonical factorization $|c_1|^2 t_1^\dagger(z)t_1(z)$. Let the numerator factorization be denoted $|c_2|^2 t_2^\dagger(z)t_2(z)$. Then, once we normalize out the maximum-phase terms, we get

$$\mathbf{g}_2(z) = \frac{\left(\mathbf{h}_1^\dagger(z)\mathbf{h}_1(z)\right)\mathbf{h}_2(z) - \left(\mathbf{h}_1^\dagger(z)\mathbf{h}_2(z)\right)\mathbf{h}_1(z)}{c_2^* t_2^\dagger(z)c_1 t_1(z)}. \quad (3.19)$$

The set of beamforming vectors

$$\mathbf{G}_B(z) = \begin{bmatrix} \mathbf{g}_1(z) & \mathbf{g}_2(z) \end{bmatrix}$$

now satisfies the paraunitary constraint. Using (3.16) and (3.19), we see that the new effective channel becomes

$$\mathbf{H}(z)\mathbf{G}_B(z) = \begin{bmatrix} c_1 t_1(z) & 0 \\ \frac{\mathbf{h}_2^\dagger(z)\mathbf{h}_1(z)}{c_1^* t_1^\dagger(z)} & \frac{c_2 t_2(z)}{c_1 t_1(z)} \end{bmatrix} \quad (3.20)$$

and the SNR at the second receiver will be

$$\frac{\mathcal{P}}{\mathcal{N}_0} \cdot \frac{|c_2 d_2|^2}{|c_1|^2}.$$

The procedure for adding yet more streams follows easily, though we omit the details here since the equations become more cluttered.

Operation of Precoding Algorithm

It is worth taking a minute to consider the precoding algorithm implied by the effective channel of (3.20). Recall that $t_k(z)$ are monic and minimum phase, while $\mathbf{h}_k^\dagger(z)$ are causal. This means that the diagonal elements will be causal, but the entries below the diagonal will not. What does this say about how the precoding operation must proceed?

Imagine that the transmitter wants to precode the current data symbol of stream




Figure 3-14: Simulated outage probability for received SNR for the second of two receivers. We compare the methods of Section 3.3.2 and Section 3.3.3, with a 4-element array, 4 i.i.d. Rayleigh-distributed taps each of variance 0 dB, $\mathcal{P}/\mathcal{N}_0 = 0$ dB, and equal power distributed between the two receivers.

k. It needs to compute the interference that will appear for this data symbol, subtract it out, then perform a modulo on the result. From the structure of (3.20), the interference will depend on this stream's own past precoded symbols, and past, present, and future precoded symbols of earlier streams. Therefore, before the transmitter can precode this stream, it needs to wait for all earlier streams to be precoded. What results is the algorithm flow of Fig. 3-13b. The transmitter precodes all the symbols of the first stream, then all the symbols of the second stream, etc. Realistic implementations would probably truncate the responses of (3.20), so that the processing of each stream only needs to stay a specific number of symbols ahead of the next one.

Performance


We expect that this noncausal precoding method will exhibit a performance improvement over the causal method of Section 3.3.2, which only takes advantage of the first

tap of each channel filter and simply cancels energy from the other taps. The simulated outage curves of Fig. 3-14, for a 4-element array and four taps per channel, bear this out. Shown is the performance for the second of two receivers; the first receiver's performance is similar but is 1 to 2 dB higher. The causal curve is the same as the usual third-order diversity for flat channels. (Recall that with M antenna elements, the k th receiver gets diversity order $M - k + 1$.) Because the filtering in the non-causal method attempts to use energy from all the taps, it achieves not only better average performance, but also has smaller tails resulting in a sharper outage curve. At low outage, the gain is almost 10 dB. Even the noncausal method can not gain back all of the energy from all of the taps, but it does come close: the first stream's mean received SNR is 90 percent of the matched filter bound. Simulations for ergodic capacity are given in the next section.


3.3.4 Comparison with DMT Method

The DMT-based method of Ginis and Cioffi [30] takes a very different approach, as summarized in Fig. 3-15. Each stream is broken into blocks of N symbols, and each block is put through an inverse discrete Fourier transform and then prepended with a cyclic prefix before being sent through the channel. A receiver waits for the entire block to be received and takes the N -point discrete Fourier transform (DFT). The overall effect is transforming the ISI channel into a series of N parallel single-tap channels, with the taps equal to the N -point DFT coefficients of the channel impulse response. For this to work, the cyclic prefix, which carries no useful information, must be as long as the ISI. In the context of the multiple-receiver problem, this whole procedure transforms the $\mathbf{H}(z)$ matrix into N parallel matrices with single-tap entries. Now, the beamforming/precoding procedure for flat fading channels can be applied to each of these separately. We will continue to assume that the system uses the precoding method that results in zero interference.

A comparison between our precoding method with noncausal filtering and the DMT-based method reveals many features of the classic single tone versus multitone discussion that has traditionally centered around scalar ISI channels. One way to think of frequency-selective channels with Gaussian noise is as an infinite number of parallel channels at different frequencies that can be optimized separately. Multitone methods try to approximate this with a finite number of parallel channels N , and break up the input into this many substreams. Different power, modulation, coding,



(a) Transmitter



(b) Typical receiver

Figure 3-15: Block diagrams for DMT method with two receivers. Shown are the processing at the transmitter and at a typical receiver.

and now, precoding, can be used on the different substreams. Single-tone methods instead encode the whole stream together and use a filter to spread each symbol over all frequencies so that the transmit spectrum is optimal. Multitone methods suffer from the overhead of the cyclic prefix and from the finite number of subchannels approximation. Single-tone methods typically lead to more complex receivers. This complexity, along with the necessity of receiver cooperation, was alleviated for the most part by using precoding, but this solution leads to its own set of issues.

For both methods, precoding forces the ordering among streams to be set at the transmitter. In the multitone solution, this ordering can be done separately for each subchannel. This is not quite the same as being able to reorder for each frequency, both because of the finite number of subchannels and also that the DMT causes some leakage of energy across frequency bands. As the block size gets large, these issues should disappear. In any case, there more control over the ordering than the single-




Figure 3-16: Ergodic sum capacity across two receivers for the different methods. Parameters are the same as for Fig. 3-14, except the input SNR, $\mathcal{P}/\mathcal{N}_0$, is made variable. The DMT method used 32 tones, and the rate penalty from the cyclic prefix is ignored.

tone solution. There, the channel triangularization was performed over the whole band, so the same ordering is used over all frequencies.

Another issue is the manner with which the interference is dealt. Both methods use zero-forcing precoding to eliminate the interference across the different streams. For interference across time, however, the DMT codes separately at the different frequencies, while the single-tone method again uses precoding. Zero-forcing precoding is known to be optimal at high SNR [7, 12], but is not in general, so the DMT seems to have an advantage here.

In our preliminary simulations, these details do not seem to result in major differences in performance. For example, Fig. 3-16 shows ergodic sum capacity for two receivers, four antenna elements, and four taps per channel. The difference between the DMT method (which includes optimal ordering at each tone, and waterfilling over both streams and tones) and our noncausal single-tone method (with a random

ordering, and waterfilling only over streams) is negligible except at very low SNR. Apparently, the filtering in the single-tone method is able to shift most of the power to the zero-lag taps and does not suffer from the zero-forcing approach to ISI. It also suggests that the ordering and waterfilling issues are of secondary importance here. As expected, the causal precoding method of Section 3.3.2 lags in performance. The multitone method did not seem to be very sensitive to the number of tones chosen, either. However, one reason to choose a larger number of tones would be to lower the overhead of the cyclic prefix: if this had been included, the single-tone method would have been better at most SNRs. As the number of receivers is increased and the system becomes more constrained, all of these second-order effects may gain in importance.

Changing from the zero-forcing to the more general multiuser precoder that maximizes sum capacity, as in [82], would be straightforward for the DMT method, although as of now there is no provably optimal algorithm for finding the optimal beamforming matrix. Similarly modifying the single-tone solution to allow just the right amount of interference, but now over both streams and time, is likely to be possible in principle but difficult in practice. Recall, though, that our earlier results suggested that at reasonable input SNRs, zero-forcing precoding (at least across receivers) does seem to achieve a large part of the potential gain. Similarly, waterfilling across both streams and frequencies may be easier for the DMT method (where they combine to form a single, larger power control problem), but our simulations and those of others [12] suggest that waterfilling does not play a major role except in cases of blocking off particularly bad channel segments.

The single-tone solution does have additional practical advantages. Each receiver's stream is sent with a single modulation and channel code, as opposed to potentially different ones for each of the N DMT subchannels. The single-tone beamforming filters and precoder are somewhat more complex than in multitone, but once again there is only one set of these. The DMT must run a separate set of beamformers and precoders, with different coefficients, for each subchannel. The transmitter must not only operate all of these different functional units, with corresponding added complexity at the receiver, it also needs to figure out the correct parameters. Furthermore, since precoding is somewhat sensitive to channel estimation errors, having so many separate precoders may require either more accurate channel estimation or more conservative rates.

Both types of systems require a certain amount of delay. For the single-tone system, the delay is in the noncausal filtering and waiting for earlier-ordered streams to be precoded first. (Recall the algorithm operation of Fig. 3-13b.) Both of these can be made finite by truncation. The DMT method processes signals block-wise, so both transmitter and receiver must wait for entire blocks to appear before processing them.

3.4 Concluding Remarks

The array processing described in this chapter provides another example of the power of precoding/dirty paper coding approaches in a variety of applications. These methods have the ability to layer information at rates that were previously only available with additional receiver processing and coordination. We view its application to array processing in terms of a factorization between linear and nonlinear operations. We saw in both the flat fading and frequency-selective fading scenarios that this partitioning can be done in many ways, leading to different types of processing, multiple orderings among streams, and various performance tradeoffs.

Our discussion has concentrated at the level of understanding these partitionings and their implementations in practical systems. Several open questions remain along these lines, many of which are active research directions. These include developing coding techniques that “close the gap” to capacity, further characterizing robustness to imperfect channel knowledge, finding operating points for various performance criteria, and exploring the role of the distortion compensation parameter for specific signaling schemes.

Chapter 4

Informed Data Scheduling

We now focus on how to improve performance by using schedulers that are aware of the physical channel state and other system components. This is in contrast to traditional layered architectures, where the two problems of selecting which streams to send at a particular time and of communicating those selected streams with highest efficiency have usually been considered separately. For example, cellular systems often have a medium access control (MAC) layer that assigns slots or waveforms relatively independently of the channel state, then a physical layer that may apply adaptive techniques based on properties of the links. The analysis of precoding in the previous chapter, although incorporating some amount of data stream awareness, was in this tradition in the sense that the channel vectors were random and presumably selected by an independent upper layer. However, the strong roles that interference and fading play in spatial multiplexing suggest that further integration between layers may be fruitful.

We have seen how both the overall system throughput and the reliability of individual links can be improved by scheduling more than one stream simultaneously. However, an array of a given size can only spatially multiplex a limited number of streams effectively. When more than this many streams have data to send, the scheduler must make decisions on how they are to be grouped. If this scheduling process is informed by the state of the channel vectors, the general scheme of the physical layer, and a small amount of information about the data itself, then the system as a whole will benefit. In addition to performance gains, the system may even be able to reduce the computational requirements of the array processing, so that less-intensive techniques such as beamforming can be used instead of precoding.

Any discussion of optimizing performance must be sensitive to the particular goals and constraints of the system. In Section 4.1 we describe a classification scheme based on the delay tolerances of data streams and explain what kinds of scheduling are appropriate for each class. Once again, we assume that each stream is intended for a unique receiver. Then, in Sections 4.2–4.5, we study scheduling algorithms for these classes in more detail. For these scenarios, we discuss the key roles of channel orthogonality and magnitude, and how our algorithms attempt to optimize these to improve performance. Finally, in Section 4.6, we bring together some ideas on combining different data classes within the same system. Although more research is required if systems must give rate and delay guarantees to individual streams, our results show the promise of channel-aware scheduling for array systems.

4.1 Data Model: Classification by Delay Tolerance

Scheduling algorithms can operate at a variety of levels, depending on the features of the channel and data streams they choose to model. On one side are algorithms from the networking community such as weighted fair queuing [53, 16] that typically assume a reliable channel and seek to ensure certain qualities of service for a heterogeneous set of data streams. A refined model called service curves [15, 60] enables a system to satisfy both rate and delay guarantees simultaneously by having each stream specify an entire set of rate goals at various delays. Unfortunately, these types of results are difficult to apply to our wireless channel of interest, where the total system rate depends highly on the particular set of streams selected at each time. Other approaches pursue less ambitious service guarantees but include a greater consideration of the physical channel. For example, in the multiuser uplink channel with single-element antennas, Tse and Hanley [67] derive scheduling and power control algorithms for maximizing the instantaneous weighted sum of rates among the different streams. Okamoto [51] and Shan, et al., [57] describe some scheduling algorithms for adding spatial multiplexing to array systems while maintaining SINR goals. For a downlink array system, Viswanath and Tse [74] suggest an adaptive timesharing strategy that transmits a stream whose associated channel realization has high quality with respect to its mean value.

In this chapter, we consider scheduling algorithms that take advantage of channel knowledge and lower-layer spatial multiplexing, yet still respect essential differences

Type	Delay tolerance
tight delay	one to several packet lengths
medium delay	several packet lengths to several coherence times
large delay	more than several coherence times

Table 4.1: Summary of data types, organized by delay tolerances

between classes of data. Since the overall performance will depend on the flexibility of the scheduler to rearrange packets and set rates, we classify data streams based on a few general levels of delay tolerance, as summarized in Table 4.1. The data with the tightest delay, such as critical sensor data, must be sent within a small number of packet lengths. The next level of data, perhaps modeling voice traffic, is more tolerant but still useless if not received within a few coherence times. In other words, the scheduler can rearrange the data streams into different groups, but can not count on waiting for channel realizations to change. We will see that a good strategy here is to select groups of receivers with nearly orthogonal channel vectors. Finally, data with the largest delay tolerance, such as background file transfers, is concerned only with long-term average rates, meaning that the scheduler has the freedom to send only those streams with good instantaneous channel realizations. Although this classification, based on when data must be sent, should not be confused with the discussion of signaling strategies in Section 2.2.1, similar performance criteria are appropriate. We will primarily look at individual-receiver outage for medium-delay data and ergodic sum capacity for large-delay data. Data with tight delay constraints is of a different nature, and is concerned with the delay of a particular packet.

Rather than a detailed source model, we consider a simple mechanism whereby each stream delivers data into a separate buffer. When the buffer reaches some minimum threshold, it places an entry into the queue of “ready” streams with data to send. When this stream is selected for transmission, it passes data from the buffer to the array processor functions and removes the entry from the queue. We first consider appropriate scheduling techniques for queues of a single data class, then in Section 4.6 discuss some ideas for systems with multiple classes.

These results are a starting point for making scheduling more aware of the channel state and array processing. They show the potential for improved performance, and of how a channel-aware scheduler can reduce the complexity requirements of

other system components. Our specific algorithms are not meant to replace quality of service-based methods without further development. Elements such as source distributions and admission control are needed before they can hope to make the same types of guarantees. However, our contention is that an efficient synthesis of the two approaches must proceed from a firm basis in the lower-level awareness rather than just a small adjustment off of existing quality of service-based methods.

4.2 Spatial Multiplexing Performance for Data with Medium Delay Tolerance

Before developing scheduling algorithms, we must first identify the key ways in which scheduling can impact performance. Let us say that for data with medium delay tolerance, the scheduler must send each stream in the “ready” queue within a given bounded waiting time, and that the realized channel vectors remain constant within this time period. The primary impact of the scheduler, then, will be in how these channel vectors are grouped together.

In this section, we study the performance of spatial multiplexing methods under different assumptions on the set of channel vectors. The concentration is on zero-forcing beamforming, with some discussion on precoding as well. We will see that the angle between channel vectors plays a major role, and that the scheduler should therefore choose groups of receivers with nearly orthogonal channel vectors. At the limit of a purely orthogonal set, precoding reduces to beamforming, suggesting that with channel-aware scheduling, the computational requirements of array processing will be reduced. These findings will inform the scheduling algorithms developed in Section 4.3.

4.2.1 Diversity Analysis with Random Channel Vectors

An interesting way of looking at the tradeoff between the number of receivers and performance is in terms of the diversity benefit of the array. For a single receiver and Rayleigh fading, as we saw in Fig. 2-3, the effect of adding transmit antenna elements is both to increase the average SNR and also to change the distribution to one with considerably less variation relative to the mean (in particular, an M th-order Erlang).

When there are K receivers, we expect the power constraint to limit the average

SNR per receiver to only $M\mathcal{P}/K\mathcal{N}_0$ rather than $M\mathcal{P}/\mathcal{N}_0$. However, the interference issue turns out to have a severe effect as well. Leveraging a result from the zero-forcing receive diversity solution [79, 65], it can be shown that in the absence of power control (i.e., if the transmitter sends signals of equal power to each receiver),

$$\text{SNR} \sim \text{Erlang}(M - K + 1), \quad \text{mean} = \frac{M - K + 1}{K} \frac{\mathcal{P}}{\mathcal{N}_0}. \quad (4.1)$$

The important observation is that there is effectively a tradeoff between the diversity benefit of the array and the number of receivers to be multiplexed. For each receiver the transmitter has to null out, a stream loses one degree of diversity. For $K = M$, for example, (4.1) suggests that, once the K -fold loss in average transmitted power is normalized out, each receiver's SNR distribution is the same as if we were transmitting to only that one receiver using a single antenna element.

One might think that using power control on the zero-forcing solution to equalize the SNRs of the different receivers, as was done with precoding in Section 3.2.2, might help increase the effective diversity (perhaps at the expense of peak performance). With power control, the SNR for each receiver will be

$$\text{SNR}_{\text{pc}} = \frac{\mathcal{P}}{\mathcal{N}_0 \sum_{k=1}^K \frac{1}{\sigma_k^2(\mathbf{H})}} \leq \frac{\mathcal{P}}{\mathcal{N}_0} \sigma_{\min}^2(\mathbf{H}), \quad (4.2)$$

where $\sigma_k(\mathbf{H})$ are the singular values of \mathbf{H} . This is shown by starting from the fact that the beamforming matrix \mathbf{G} must be a scaled pseudoinverse of \mathbf{H} and finding that scaling factor:

$$\begin{aligned} \text{Constraints:} \quad & \mathbf{G} = c\mathbf{H}^\dagger (\mathbf{H}\mathbf{H}^\dagger)^{-1} & \text{trace}\{\mathbf{G}^\dagger\mathbf{G}\} = \mathcal{P} \\ \Rightarrow & c^2 \text{trace}\left\{(\mathbf{H}\mathbf{H}^\dagger)^{-1}\right\} = \mathcal{P} \\ & c^2 \sum_{k=1}^K \sigma\left((\mathbf{H}\mathbf{H}^\dagger)^{-1}\right) = \mathcal{P} \\ & c^2 = \frac{\mathcal{P}}{\sum_{k=1}^K \frac{1}{\sigma^2(\mathbf{H})}}. \end{aligned}$$

When $K = M$, the upper bound in (4.2) has an exponential distribution [18] and once again equals the single-user, single antenna element distribution (except with a loss of K in average power).

The diversity loss will be less severe when there are more antenna elements than receivers, and performance will tail off more gracefully for good non-zero-forcing strategies, but in all of these cases, we see the fundamental conflict between sending to more receivers and the benefits of diversity.

4.2.2 Diversity With Orthogonal Channel Vectors

The discussion above assumed that we must transmit to a group of randomly-selected receivers at a single time. It is exactly this random selection of receivers that causes the loss in diversity. When selecting streams to spatially multiplex, one solution would be to choose only those streams whose receivers have nearly orthogonal channels. Before going on to propose specific systems that attempt to do this, we investigate the performance potential when sending to a random set of orthogonal channels.

Suppose that K receivers are multiplexed using an M -element array ($K \leq M$). The channel coefficients have the same distribution as before, but now assume that the channels are orthogonal, so that the transmitter can beamform perfectly to each receiver without adding interference. The distribution in (4.1) now becomes,

$$\text{SNR}_{\text{orth}} \sim \text{Erlang}(M), \quad \text{mean}_{\text{orth}} = \frac{M \mathcal{P}}{K \mathcal{N}_0}. \quad (4.3)$$

As expected, each receiver now gets the full M -level diversity, with just the $1/K$ factor in average SNR due to multiplexing among K streams. Precoding can be seen as achieving a compromise between (4.1) and (4.3), in that the k th receiver sees an $\text{Erlang}(M - k + 1)$ distribution in SNR.

With orthogonal channel vectors and power control,

$$\text{SNR}_{\text{orth,pc}} = \frac{\mathcal{P}}{\mathcal{N}_0 \sum_{k=1}^K \frac{1}{\|\mathbf{h}_k\|^2}} \leq \frac{\mathcal{P}}{\mathcal{N}_0} \|\mathbf{h}_{\min}\|^2.$$

Because each receiver, before power control, has an equal or greater SNR than if the channel vectors had not been orthogonal, this value is necessarily larger than (4.2). The difference tends to be significant, since the harmonic mean in these formulas is usually dominated by the weaker elements, and $\|\mathbf{h}_{\min}\|^2$ will typically be much larger than the minimum singular value $\sigma_{\max}^2(\mathbf{H})$. The outage distribution for an 8-element array is shown in Fig. 4-1, where the difference in both average SNR and the shape of the distribution are substantial.




Figure 4-1: Outage probability for an 8-element array transmitting to 8 receivers, with an SNR per link of 5 dB and power control.

The improvement in using orthogonal receivers can also be seen in the deterministic asymptotic performance (as in [68] for receive diversity) for large systems. When K and M grow to infinity according to a certain ratio $\beta = K/M$, the performance with random and orthogonal receivers are

$$\begin{aligned} \text{SNR} &\rightarrow \frac{\mathcal{P}}{\mathcal{N}_0} \frac{1 - \beta}{\beta} \\ \text{SNR}_{\text{orth}} &\rightarrow \frac{\mathcal{P}}{\mathcal{N}_0} \frac{1}{\beta}. \end{aligned}$$

These asymptotic results are plotted in Fig. 4-2 and show a 3 dB advantage for orthogonal channel vectors at $\beta = 0.5$, 6 dB at $\beta = 0.75$, and rapidly increasing after that. For large systems with small β , even randomly selected channels will most often be nearly orthogonal, so there is not much to be gained in using selected receivers. This does not completely carry over to small systems with small β , since the non-deterministic performance may still result in channels with bad correlations, but the




Figure 4-2: Deterministic received SNR (per receiver) for a large system with $K/M = \beta$ and an input SNR per link of 5 dB.

general result still holds that orthogonal receiver selection is much more important to systems with higher K/M .

A related advantage of trying to find orthogonal channel vectors is in added numerical stability of the beamforming algorithms. For example, when the zero-forcing beamforming matrix under power control (i.e., pseudoinverse) is computed, relative perturbations in \mathbf{H} can be magnified by a factor bounded by $2\kappa_2(\mathbf{H})$, where $\kappa_2(\mathbf{H})$ is the condition number,

$$\kappa_2(\mathbf{H}) = \frac{\sigma_{\max}(\mathbf{H})}{\sigma_{\min}(\mathbf{H})},$$

of \mathbf{H} , as long as the rank is not changed [59]. Similarly, if instead precoding using the \mathbf{LQ} factorization is performed, perturbations magnified by about $\kappa_2(\mathbf{H})$ are seen in \mathbf{Q} and \mathbf{L} , and $\mathbf{Q}^\dagger \mathbf{Q}$ differs from identity by a matrix of approximate norm $\epsilon \kappa_2(\mathbf{H})$, where ϵ is the machine precision. Using nearly orthogonal rather than random channel vectors results much better conditioned \mathbf{H} matrices (i.e., with smaller $\kappa_2(\mathbf{H})$), thus

producing more stable computations.

We have demonstrated some of the advantages of using receivers with uncorrelated channels. However, unless $K \ll M$, then these channels are unlikely to occur among randomly chosen receivers, and this is the regime that provides the smallest advantage. The solution is to look at a wider view of a system, which will likely contain more streams than can be spatially multiplexed at any one time. We propose an amount of integration between the physical and MAC layers, so that the correlations between channels can inform the transmitter on how to intelligently group streams to help achieve better overall performance.

4.3 Scheduling Algorithms for Data with Medium Delay Tolerance

We now go on to develop scheduling algorithms and evaluate their performance. We begin with an example where all streams are grouped into subsets, and then consider a more dynamic queuing model whereby the set of streams with enough data to send changes over time. As more data streams enter the queue, performance should increase because the scheduler has more flexibility to select appropriate groupings. We will see that not only does this expected behavior occur, but also that most of the improvement can happen with a fairly small number of streams in the queue.

The purpose of this study is to determine the potential for channel-aware scheduling. We do not make an attempt to optimize for delay, but rather to minimize outage while ensuring that all streams in the queue get scheduled before their channel parameters are likely to change. A more complete characterization of tradeoffs between delay and outage or throughput performance remains for future study. We do provide some analysis of delay characteristics, and will revisit this issue in Section 4.6, but further research is required if precise delay guarantees are necessary. However, because of the outage improvement seen with only a small amount of grouping flexibility, we expect that systems may be able to support high rates even with additional delay constraints. Conversely, a system without any additional constraints may be able to use simplified schedulers and still achieve most of the available gains.




Figure 4-3: Simulated outage curves when transmitting from a 4-element array to groups of 3 receivers using zero-forcing beamforming, at an input SNR per link of $\mathcal{P}/\mathcal{N}_0 = 5$ dB. Streams are partitioned into groups using a “greedy” algorithm.

4.3.1 Static Model Example

Consider a system with K' streams, all of which send continuous data. Therefore, the scheduler must divide all of them into spatial multicast groups for each channel realization. If these groups are of size K , then $\lceil K'/K \rceil$ groups are necessary. Given the discussion in the previous section, we would like the scheduler to select groups in which the angles between channel vectors is large. A good scheduling algorithm should be able to approach the bound of orthogonal channel vectors as K' increases.

Unfortunately, the optimal scheduling for this problem is unknown, and in any case appears to be combinatorial in nature. More promising are “greedy” algorithms, which select the optimal result at each step rather than doing a global search. Fig. 4-3 shows simulation results for a 4-element array and subsets of 3 receivers using an algorithm of this type:

1. The first $\lceil K'/K \rceil$ streams are put into separate groups

2. The next $\lceil K'/K \rceil$ streams are placed, one by one, into the groups to which they are “most orthogonal” (that is, the largest angle between the associated channel vectors), until all groups now contain two streams.
3. The last set of streams are placed similarly, but now to the group with the maximum of angles to those already in the group.

Within each group, we use zero-forcing beamforming without power control. As the population size increases, the performance grows steadily from the second order diversity of random groupings to the fourth order diversity of perfectly orthogonal channel vectors. At 10% outage, 4 dB out of the potential 5 dB gain is achieved with 32 groupings. At 1% outage, a similar portion of the total 7.5 dB gain is achieved. The ergodic capacity (without power control or waterfilling) increases as well, from 4.5 bits/channel use for random selection to 6.2 for 32 groups, out of a potential 6.8 for orthogonal channel vectors.

A first-fit algorithm and its variations in [51] and [57] also multiplexed a set number of users into timeslots in a greedy-type manner. However, those efforts sought to maximize the number of users in each slot given SINR constraints rather than optimize outage given a number of users per slot, making comparisons difficult. Additionally, they did not directly emphasize achieving orthogonality between channel vectors, but only implicitly through the SINR constraint.

4.3.2 Dynamic Queuing Model

A more realistic and dynamic model considers streams queuing up and the transmitter when they have data to send. The scheduler could just group the first K streams at the head of the queue together, but in the spirit of this section, higher performance can be achieved if there is more freedom in choosing how streams are grouped together. We quantify this idea by allowing a window of K' streams at the front of the queue from which K must be selected. As K' increases, we expect performance to increase as a more orthogonal set of channel vectors can be chosen.

To be more specific, imagine a replenishable queue of “ready” streams, each associated with a random channel vector. A diagram is shown in Fig. 4-4. To ensure a bounded waiting time, the first entry in the queue must be sent at the current time, but the other $K - 1$ streams to be multiplexed can be chosen from anywhere in the $K' - 1$ remaining streams within the window. For the simulation, these are




Figure 4-4: Diagram of queueing model

chosen in a similar manner to the static channel, but now we select the channel vector that is most orthogonal to the subspace of the channels already chosen for the group. Equivalently, this is the receiver that will lose the smallest fraction of its SNR upon zero-forcing beamforming. After this set of K entries is sent, they are removed from the queue, and K new packets with random channel vectors are added to the end of the queue. By adding more random channel vectors each time, we either assume a population size much larger than K' , or that by the time new packets from the same streams reach the window, their receivers' channel vectors have changed.

Fig. 4-5 shows simulation results for an 8-element array that schedules $K = 8$ streams at a time. Because the data has medium delay tolerance, we evaluate performance by individual-receiver outage and use power control. As in Fig. 4-3, the curve when the scheduler just selects the eight streams at the head of queue is far from the bound for orthogonal channel vectors. However, even a very small amount of freedom, selecting eight of the first nine streams, leads to gains of 5 to 10 dB for outages in the range of 1% to 10%. By the time the window size has reached twenty, the outage curves are starting to approach the orthogonal bound.

Other array processing methods will also benefit from channel-aware scheduling. For instance, orthogonality between channel vectors is also desirable for the precoding solutions of Chapter 3. There, later streams must direct nulls to earlier-ordered receivers, although the reverse is not true. The scheduling technique described above therefore selects the new stream that stands to lose the least by having to precode off of the streams already selected. In Fig. 4-6, we compare the earlier beamforming




Figure 4-5: Outage probabilities for a queuing system with $M = 8$ transmit antenna elements and $K = 8$ simultaneous receivers chosen from a window size varying from 8 to 20. We use zero-forcing beamforming, power control, and an input SNR per link of 5 dB. Shown for comparison is a bound on outage corresponding to orthogonal channel vectors.

curves side-by-side with those for precoding, where power control and the max min ordering method of Section 3.2.2 are used. Note how precoding with a small window size achieves similar outage performance to beamforming with a large window size. Precoding improves still further with larger window sizes, but by this time the incremental gains are smaller. We see similar trends looking at ergodic sum capacity in Fig. 4-7. This illustrates one of our main themes, that to get most of the benefits of a transmitter array, a system designer often has a choice between sophisticated scheduling or array processing and does not necessarily have to use high complexity at both sides.

For various reasons, including numerical stability of the beamforming and LQ operations, it might be desirable at times to spatially multiplex only 6 or 7 streams using the 8-element array. Outage and sum capacity curves will show the same




Figure 4-6: Same as Fig. 4-5, but we now add similar curves for precoding, with power control and proposed max min ordering.

general trends under these circumstances, though with less relative improvement as the window size increases. The particular $\mathcal{P}/\mathcal{N}_0$ value will also affect the exact rate tradeoffs associated with multiplexing more or fewer streams.

It is desirable to have a low-complexity method of grouping streams. The approach used in this section can be implemented as follows: Find an orthonormal basis (using the Gram-Schmidt procedure, for example) for the channel vectors of the streams already in a group. Then, multiply a candidate's channel vector by the matrix of this basis, and determine the fraction of energy that remains. Note that the matrix stays the same for all candidates, and once one candidate is chosen, only one new element of the updated basis needs to be computed. The group selection appears to be somewhat robust to different methods as well. A different grouping based on minimizing the condition number achieved almost the same performance as this one (at higher complexity). More ad-hoc methods may achieve similar performance at lower complexity. With most reasonable methods, complexity grows with the window




Figure 4-7: Ergodic sum capacity for zero-forcing beamforming and precoding, for the same simulation as in Fig. 4-5 except without power control.

size; fortunately, it appears that most of the performance gains are achievable with relatively small windows.

Delay Characteristics

This selective grouping procedure can result in longer delays than a simple first-in-first-out (FIFO) model, but this delay is bounded. Consider a total population of P streams intended for distinct receivers, backlogged so that the queue always contains exactly one packet from each stream. (That is, a stream will always have data, but is not allowed to put another packet on the queue until its previous stream has been sent.) The transmitter sends K streams at a time. In FIFO, a packet therefore always jumps K places toward the head of the queue at each time. In the selective grouping presented above, the packet always moves at least one spot, but perhaps not more. We can summarize the maximum and minimum delays, in terms of turns in the queue, as follows:

- *FIFO*: Maximum delay of

$$\left\lfloor \frac{P-1}{K} \right\rfloor.$$

The minimum delay (assuming all streams are backlogged) is only one less than this. At each time, K new streams are added to the queue. The delay can vary by one depending on where the packet of interest is among these K .

- *Selective grouping, window size P* : The maximum delay is

$$P - K,$$

but this will occur very rarely. The minimum delay is zero, because this packet could be chosen as soon as it enters the queue.

- *Selective grouping, window size K' , where $K \leq K' \leq P$* : This provides a compromise, where packets jump K places each turn until they reach the window, and may move slower after that. By judiciously choosing a relatively small window size K' that achieves most of the available performance gains, one can improve delay as well as complexity. The maximum delay is

$$1 + \left\lfloor \frac{P - K' - 1}{K} \right\rfloor + K' - K.$$

The first two terms are the time it takes to enter the window, while $K' - K$ is the maximum time spent within the window. Note that this reverts to the other two cases (in delay and algorithmically) when $K' = K$ or $K' = P$.

Our experience from the preceding simulations suggests that the worst-case delay occurs only rarely. Also, since selective grouping increases the amount of information that can be transmitted at each time, the disparity in delay per information bit will not be as great as that of delay in terms of turns in the queue as given above. However, we do suggest that if minimizing delay is of greatest importance, the scheduling algorithm should be modified somewhat.

4.4 Large Delay Tolerance

While the scheduling delay in the previous section was on the order of several packet lengths, other types of data may tolerate much longer waiting times. For example, when doing file transfers or system backups, achieving a high average throughput may be much more important than the delay on any particular packet. In these cases, the system can simply maximize the sum rate over each channel realization, and over time the rates for the different streams will even out.

This idea relates to a growing body of literature on “multiuser diversity,” in which each stream communicates when its associated channel vector is near its peak strength. However, most of these results are for timesharing strategies where only one stream can transmit at a time. Below, we provide a discussion of when such timesharing strategies are optimal and go on to develop scheduling algorithms for spatial multiplexing.

4.4.1 Relation to Timesharing Strategies

One can gain a perspective on timesharing versus spatial multiplexing by placing our problem within a larger context of multiterminal wireless scenarios. The channel may be in the uplink or downlink direction, and the base station may or may not have a multiple-element array. By looking at these different cases, we can gain an appreciation of the roles that waterfilling, power constraints, and spatial multiplexing play. In some cases, timesharing will be sufficient for maximizing the sum capacity, while in others, the gains associated with spatial multiplexing will far outweigh those achieved by simply using a stream during a good channel realization.

Table 4.2 summarizes some results for these different scenarios. One important factor is the form of power constraint used. In this thesis, we have concentrated on a peak power constraint, so that at each time, the expected power is below some prescribed limit, $\mathcal{E}[\mathbf{x}^\dagger \mathbf{x}] \leq \mathcal{P}$. A system could also potentially allow a transmitter to save up unused power for a later time, so \mathcal{P} becomes a constraint on average power over all time. The first case is more appropriate for satisfying regulatory limits or minimizing out-of-cell interference, while the second may be a better model for maximizing battery life. With an average power constraint, each fading realization may be considered as a kind of parallel channel over time [66] over which power can be waterfilled. In either case, the uplink power constraint is for each user individually,

Scenario	Peak Power Constraint	Average Power Constraint
uplink	No	Yes [40, 67]
downlink	Yes [66]	Yes [66]
uplink, array	No	No [76]
downlink, array	No [8]	No

Table 4.2: Summary of when timesharing strategies are sufficient for maximizing sum capacity.

while the downlink power constraint is for all streams combined. The results shown all assume ergodic variation in the channel parameters and equal distributions for all users, but do not rely on a Rayleigh fading model.

When there are multiple users in the system with separate streams, the system has a choice of sending multiple streams at once with signaling-level techniques such as beamforming (if there is an array), dirty-paper coding, and interference cancellation. Perhaps surprisingly, then, Knopp and Humblet [40] reported that in a basic uplink scenario, with an average power constraint and no array, timesharing is sufficient to achieve the ergodic sum capacity. Simply put, the user with the best instantaneous channel realization gets a chance to communicate. It then waterfills power over all such situations in which it expects to be selected (so at some times, there may be no active streams). Similar results were shown for the downlink [66] and have resulted in a timesharing mode called HDR for the CDMA 2000 cellular specification [69]. This idea does not carry over as well to a peak power constraint on the uplink, since a corner point of the rate region (see [14]) with more than one active user will often result in the best sum rate for a particular realization.

Things change significantly when the base station has a multiple-element array. With array processing techniques such as those discussed in this thesis, it can often separate the signals to or from the various users enough that the channel starts to look more like parallel streams than additive interference. At this point, the system can distribute power among the different streams and achieve spatial multiplexing gains, as discussed in Section 2.3.2 and elsewhere. As we will see, this effect can become even more important than hitting each user at its peak channel strength. For example, if there are two users and a base station with a large number of antenna elements, then the channel vectors will usually be nearly orthogonal and interference will not be a major issue. Except at very low SNR or very high channel quality variation (in which




Figure 4-8: Ergodic sum capacity for channel-aware timesharing strategies with various numbers of transmit antenna elements. The curves were computed using numerical integration over independent Rayleigh fading at an input SNR per link of 5 dB.

case the users want to concentrate the streams in a small portion of the available time), there are likely to be times when the base station will communicate with both at once.

To illustrate, we plot in Fig. 4-8 the ergodic sum capacity for downstream timesharing under a peak power constraint. Without an array, the overall system performance increases noticeably with the number of receivers, as the transmitter can select a receiver whose channel realization is near its peak strength. As the array size increases, we see a lesser relative benefit. This is because the array enables single-user beamforming, which results in a received SNR distribution with considerably less relative variation over time. On the other hand, spatial multiplexing can achieve much higher rates even under simple scheduling methods, as shown earlier in Fig. 3-4. This motivates our emphasis on scheduling algorithms for spatial multiplexing rather than timesharing in the next subsection.

There are still situations where timesharing to different receivers from an array may still make sense. In some cases, the channel strength varies considerably even after the array processing. For example, the Infostations proposal [33] models receivers moving relative to the transmitter, so that closer ones may have significantly better instantaneous channels than those farther away. Another reason is the possibility of attaining diversity gains without as detailed channel information. This is the subject of the so-called “dumb antennas” scheme of [74], in which the transmitter sends along random, time-varying beamforming directions. As the number of receivers increases, the transmitter can approach the ideal timesharing performance discussed above while only knowing the instantaneous SNRs of the receivers and not their full channel vectors.

4.4.2 Scheduling for Spatial Multiplexing

We now proceed to develop scheduling that incorporates spatial multiplexing for data streams with long delay constraints. The goal, once again, is to maximize the sum capacity over each channel realization and let the rates for individual streams average out over time. Although the incremental gains are not always significantly greater than those discussed earlier for medium-delay data, these new strategies do lead to increased performance and in some situations points to lower-complexity scheduling algorithms.

We know that precoding maximizes the sum capacity when the number of streams was less than the number of transmitter antenna array elements M [82]; perhaps some extension is possible when there are greater numbers of streams. One might conjecture that this would involve a selection of no more than M receivers getting information at each time, since the transmitter can send no more than this many precoded streams at once and still completely null out interference.

This selection process recalls the “greedy” max sum ordering discussed in Section 3.2.2, and indeed involves many of the same issues. As a practical approach, we could use the same method and simply stop after M streams have been selected. The results of this procedure, including the subsequent waterfilling across streams, are shown in Fig. 4-9. There is a clear improvement with precoding over the grouping method that only considers orthogonality between channel vectors. Zero-forcing beamforming does not improve as much, although a method more tuned toward this transmission strategy may be able to achieve somewhat better gains.




Figure 4-9: Ergodic sum capacity for zero-forcing beamforming and precoding with an 8-element array and large delay constraints. “Max sum proposed” uses the method of Section 3.2.2, while “Orthogonality info only” uses only orthogonality information, as in Section 4.3. Waterfilling across streams was used once the receivers were selected.

One way to think about this problem of user selection/ordering is to say that there are two issues that affect multiuser performance:

1. Orthogonality among receivers’ channel vectors
2. Instantaneous channel strength, ignoring potential interference

In Section 4.3, we did not effectively make use of the second of these factors because each receiver had to get information during each channel realization. With fewer constraints, we now see some improvement by taking this new information into account.

To see the relative importance of the second factor, compare Fig. 4-9 with Fig. 4-10. The second figure shows the performance of a four-element array using the “pre-code order” method described above, which takes both factors into account, as well as a simpler method that only selects receivers based on their single-user SNRs, i.e.,




Figure 4-10: Ergodic sum capacity for zero-forcing beamforming and precoding with a 4-element array and large delay constraints. “Max sum proposed” uses the method of Section 3.2.2, while “Channel strength” uses only single-user SNR information. Waterfilling across streams was used once the receivers were selected.

the second factor. Comparing the two figures, it appears that for zero-forcing beamforming, most of the gains achieved by increasing the window size are due to selecting more orthogonal channel vectors, while most of the precoding gain is from choosing the strongest channels. (The two precoding curves in Fig. 4-10 would be a little further apart for eight antenna elements, but “channel strength” still achieves better performance than the orthogonality selection method.) This is because in precoding, only the last couple receivers (out of those receiving data) have to sacrifice significant performance to avoid interference, while all receivers have this problem with zero-forcing beamforming.

This has some promising implications for precoding. If most of the gains are achieved by selecting receivers by their channel strengths without regard to interference, then the complexity of the user selection and ordering methods can be significantly reduced.

4.5 Tight Delay Constraints

So far, we have characterized data by whether the allowable delay is greater or less than the coherence time of the channel. The practical distinction was whether it is reasonable for those receivers with weak instantaneous channels to wait for their channel strengths to improve before starting communication. In either case, it was assumed that this allowable delay was greater than several packet lengths, so that some rearranging and spatial multiplexing is tolerable.

At the other extreme is data that needs to be received as soon as possible, with delay constraints on the order of packet lengths. This might be true for very time-dependent information, such as control signals for a physical system or critical sensor data. One approach would be to transmit this delay-critical stream by itself, and then resume with the usual scheduling procedures. However, if this stream does not need quite all of the available resources, it might be possible to take advantage of some of the throughput improvement inherent with spatial multiplexing.

Suppose that receiver one needs to receive a packet of a certain size by some given delay. Equivalently, it needs to achieve some average rate over that time span. If the transmitter wishes to communicate simultaneously to a second receiver, it should find the solution that maximizes the rate of the second stream given the constraint on stream one's rate. Unfortunately, as previously discussed, the multiple-receiver rate region and the strategies that achieve it are unknown. Still, practical methods such as beamforming or precoding may be able to increase the total throughput while satisfying the first stream's requirements.

One way to visualize this would be to look at capacity regions. Alternatively, we could take a more direct view and consider delay. If the streams to both receivers have the same amount of data, they could be sent at the same rate and finish simultaneously. But if the first stream has the tighter delay constraint, it may require a higher rate than this. Its packet will finish first, then the transmitter can send the remaining bits of the second stream's packet at its highest possible rate, at full power along the single-user beamforming direction. The opposite could be done if the second stream's packet finishes earliest. As shown in the example of Fig. 4-11, the delays at the two receivers can be plotted against each other for various transmission strategies and power distributions among the two streams. Any $(\text{delay}_1, \text{delay}_2)$ pair that is exterior to the curves is achievable (as opposed to capacity regions that are achievable if they are interior to some boundary). Now, given a minimum delay constraint on stream




Figure 4-11: Typical delay regions for certain two-receiver strategies, computed for a sample channel matrix realization. A $(\text{delay}_1, \text{delay}_2)$ point is achievable if it is on or *outside* (i.e., up and to the right) of the boundaries shown.

one, we can see how fast we can get stream two’s packet across.

This perspective of communicating with both receivers until one of them is finished, then sending any remaining bits to the second receiver, was inspired by the “static broadcasting” setup of Shulman and Feder [58]. Their information theoretic description was for a very general channel and dealt with sending common information to both receivers, for which we will have more to say in Chapter 5. Instead of delay, they plotted its inverse, corresponding to a kind of average rate. We find that plotting delay relates more closely to the goals of time-sensitive data, and furthermore avoids potential confusion over average versus sustainable rates.

Other interesting properties come up in the delay plot. For example, timesharing between the two single-user beamforming strategies does not result in a convex combination of delay pairs, as it would with rate, but rather in the rectangular-shaped curve shown as the dotted line in the figure. This timesharing therefore only results in less desirable delay pairs. When multiplexing more than two streams, a higher-

dimensional plot can capture all of the achievable delay K -tuples.

4.6 Multiplexing Different Classes of Data

We have explored different scheduling strategies based on a few general levels of delay tolerance. For this to be useful in many realistic systems, these ideas should be combined into a single framework capable of dealing with a mixture of data classes, or perhaps even a continuum of priorities. Additionally, this system would ideally be more amenable to including more concrete service guarantees. In this section, we give some ideas on the direction such an effort may take, inspired by the previously-mentioned weighted fair queuing algorithm [53, 16].

Although a straightforward application of weighted fair queuing to spatial multiplexing would not take proper account of the physical channel, it does provide a starting point for incorporating different data priorities and rate guarantees. Given a constant-rate data channel and set of weights ϕ_k on the streams, this algorithm attempts to guarantee stream k a fraction

$$\frac{\phi_k}{\sum_l \phi_l}$$

of the overall rate, where the summation is over all streams that have data to send. For packet-based serial transmission, Parekh and Gallager [53] describe a “virtual time” implementation that guarantees that no packet will be delayed from a continuous-flow ideal by more than the largest packet length. Suppose that a packet of length $L_{k,0}$ arrives from stream k at virtual time $t_{k,0}$. This packet is given the timestamp

$$t_{k,0} + \frac{L_{k,0}}{\phi_k}, \tag{4.4}$$

which specifies the finishing virtual time. If the queue already contains a packet from this stream, then $t_{k,0}$ in (4.4) is replaced by the timestamp of the earlier packet. Packets are serviced in increasing order of timestamp.

To apply this idea to spatial multiplexing over fading channels, we interpret (4.4) and then extend it to this new context. The timestamp can be seen as weighting the time it will take to send the packet by ϕ_k and giving a credit for time spent waiting in the queue. These ideas can also apply to our fading channel model, although we lose some of the strict quality-of-service guarantees. If each transmission segment is

the same length of time (but contains a different number of bits), let $t_{k,0}$ be the time segment number in which an entry arrives in the queue. Multiple submissions are not an issue because we allow a stream to submit only one entry at a time. Next, let L be some constant, perhaps a threshold number of bits that must be buffered before a stream can send an entry to the queue. To establish a “finishing time” for this packet, we need to divide by the rate at which data will be sent, taking into account the channel state and the other receivers to be spatially multiplexed. With precoding (and no power control), this rate can be computed from the channel realization and previously selected streams. Under beamforming, the system can estimate the value based on this information. After normalizing the first term in (4.4) with respect to the current time t , we thus select the receiver with the smallest

$$(t_{k,0} - t) + \frac{L}{r_k(t)\phi_k}. \quad (4.5)$$

This again represents a weighted sum between the time spent in the queue and the potential rate. Delay-tolerant streams will set a relatively small ϕ_k so that they will be transmitted only when performance is very high or when the system is not very busy. Streams that are more delay constrained will set a larger ϕ_k so that they will not have to wait very long, even if the channel is not very strong. Note how this scheme allows for a continuum of delay tolerances, rather than just a discrete number of classes. However, it will require calibrating the weights to achieve a proper balance between the two terms in (4.5).

The High Data Rate (HDR) mode in the CDMA 2000 wireless standard [69] and a related system for transmitter arrays [74] include many of the same issues for their timesharing-based systems. These systems attempt to schedule each stream near its peak channel quality while providing “proportional fairness” that channels with higher average quality do not receive more than their share of timeslots. Mechanically, they penalize for data recently sent (rather than crediting for time spent in the queue) and maximizing on the rate (rather than minimizing on its inverse). A version of this form of weighting could be formulated for our spatial multiplexing setup, though again the higher total rates may come at the expense of some guarantees.

Either of these two directions, inspired by weighted fair queuing and HDR, or a more direct composite of the strategies from Sections 4.3–4.5 could serve as the foundation for a scheduling algorithm that is more integrated across different types of data. The two ideas discussed in this section would result in a smoother distri-

bution of delay times than the moderate-delay queuing scheme of Section 4.3.2, and may provide an easier base on which more complex networking-oriented algorithms could be developed. For instance, to decouple the delay and rate priorities, it may be possible to add in some of the ideas from service curves [15, 60]. On the other hand, the schedulers from the bulk of this chapter dealt more directly with the appropriate optimization criteria for each data type. Due to the random nature of the channel, any algorithm will have a hard time providing strict quality-of-service guarantees. However, with enough potential receivers, the previously-discussed robustness of scheduling suggests that well-designed algorithms may be likely to achieve reasonable goals in practice.

Chapter 5

Multicasting of Common Information

We now add a deeper consideration at the scheduling and array processing blocks of whether data streams are intended for single or multiple receivers. Previously, we assumed that the scheduler would simply duplicate any streams that had multiple recipients. However, it would seem that in such multicast scenarios, it may be more efficient to transmit data only once rather than repeating it in this way. The drawback is that the transmitter must now satisfy the goals of all the recipients of this stream simultaneously. Therefore, an investigation into the potential performance and implementations of multicasting is needed.

To facilitate analysis, we first consider the array processing of a single stream in isolation, and later describe how to incorporate these ideas into the larger system context. A useful exercise here is to consider the two extremes where the stream is intended for a single receiver or for all possible receivers. In the first case, we have seen that an optimal strategy, implementable by beamforming, is to ensure that the signals from the different antenna elements coherently combine at the receiver. In the latter case, the transmitter can not effectively make use of its channel information since the data must be received at all possible locations. This presents a good application for space-time codes that do not take into account any channel side information that the transmitter may have. Using our usual independent Rayleigh fading model, the two extremes result in the same shape of the received SNR distribution, but with a factor of M difference in magnitude, where M is the number of transmitter antenna elements.

A multicast scenario is concerned with what happens in between, when a stream is directed to a finite number of receivers. This leads to two fundamental questions: Where does the performance fall within the spectrum of possibilities given above? What transmission schemes are optimal or most useful in these cases? To even begin to answer these questions requires a more precise concept of performance, since schemes that are good for some receivers may not be good for others. In Section 5.1, we provide such a discussion of performance and efficient operating points, setting up the analysis for the rest of the chapter.

We then examine techniques for different regimes and types of signaling. Beamforming strategies, analyzed in Section 5.2 and Section 5.3, are most useful when the number of receivers is small or when the transmitter can signal over many channel variations. For other scenarios, a potentially more complex class of schemes, of which beamforming is a subset, may be necessary. We investigate properties and implementations of this more general class, which we refer to as space-time multicast coding, in Section 5.4. In one example with an eight-element array and eight receivers, they achieve up to a 6 dB gain over ordinary space-time codes that do not incorporate channel knowledge.

Finally, we connect multicasting back to the larger system point of view in Section 5.5. In many cases, it is possible to transmit several multicast streams simultaneously, and these can be sent alongside receiver-specific streams.

5.1 Overview of Multicast

At the heart of multicast is an attempt to satisfy the goals of a number of receivers using a single transmission strategy. Because the receivers will experience distinct realized channel vectors, with correspondingly distinct optimal strategies, selecting the multicast parameters often requires a balance among conflicting objectives.

A related issue shows up when separate streams are directed to different receivers. For this scenario (often called a broadcast channel), researchers in information theory have long used the concept of rate regions, which describe all achievable rate K -tuples to the K receivers. Without arrays, superposition coding [14] or dirty-paper coding [13, 10, 84] is sufficient to achieve all points in the rate region, while the region is not completely known when transmitting from an array. Earlier in this thesis, we described different array processing methods for various types of tradeoffs

among receivers. As we begin to consider sending common information, we encounter important differences from these situations. For example, with multicast, interference is no longer an issue; while with multiplex, the transmitter can do more optimization on the signaling to different receivers. These details will lead to different tradeoffs, though many of the same fundamental concepts will appear.

For a given fading channel realization, consider the following hypothetical tool for capturing the benefits of different multicast strategies. Imagine a graphical plot with separate axes for each of the K intended receivers, denoting some appropriate measure of performance. This might be SNR for an uncoded system or rate for a coded one. Then, for a given fading channel realization, every transmission strategy would correspond to a K -dimensional point in “performance space.” Once all admissible strategies (or strategies of a given type) have been plotted, the various tradeoffs among the different receivers should become clear and the transmitter can select an appropriate operating point based upon system goals. By repeating this procedure across many fading channel realizations, one could also compute statistics over the random ensemble. In this way, decisions can be made based on individual receiver or system-wide goals, outage or ergodic capacity measures.

In the execution of this plan, care must be taken to ensure that the performance characterization is well-defined. The relationship between SNR and uncoded performance may only be clear in certain specialized cases, such as an additive Gaussian noise channel. For a coded system, the transmitter must choose codewords at a particular rate, even though different receivers may have the potential to reliably receive a range of rates. Therefore, a simple rate region interpretation is not sufficient. We will address these concerns with careful definitions and, at times, special cases.

With coded transmission, we resolve the issue by defining the performance axes in terms of mutual information rather than capacity. Given a particular input distribution and channel realization, this mutual information can be computed for each receiver and represents the maximum reliable rate of communication over that link. The achievable region then takes on different interpretations depending on the type of signaling used:

- If the transmitter signals over a particular fading realization at coded rate R , any receivers with mutual information of at least R for this signaling scheme will be able to reliably decode the data. Timesharing over different strategies (within the same channel realization) is also possible.

- If the transmitter signals over an ergodically varying channel, it can use a code with rate R and interleave symbols over all channel realizations. Any receiver with *expected* mutual information of at least R will be able to reliably decode the data. (To get this ergodic behavior, the schemes used at each realization must employ the same input distribution [31, 46], but we will see that this is automatically satisfied for our encoding schemes.)

Looking at the above descriptions, two particular strategies stand out. For a single channel realization, one could maximize the minimum mutual information among receivers and therefore achieve the highest rate that all can reliably decode. Alternatively, for ergodic signaling, one could maximize the *sum* of rates among receivers at each realization. This strategy then maximizes the rate of common information if all receivers undergo i.i.d. channel variations over the same fading distribution. We will discuss other operating points of interest throughout this chapter as well.

Our first investigation, however, will be over the particular subset of transmission schemes corresponding to beamforming. These perform well for small numbers of receivers (e.g., they are optimal for transmitting to two receivers, as we will show in Section 5.4.2) or with ergodic capacity goals. Furthermore, they lead to low-complexity transmission and reception techniques and are compatible with both coded and uncoded modulation. In Section 5.4, we will return to the more general scenario and discuss optimal strategies and useful implementations.

5.2 Operating Points for Beamforming

In this section, we consider multicast solutions where the transmitter using a beamforming strategy to send identical information to K receivers. The vector of antenna element outputs consists of a single input symbol multiplied by a vector of weights \mathbf{g} , resulting in coherent combining at some potential location. Each receiver gets a scaled copy of the input stream plus noise, so received SNR is a valid measure for uncoded performance (or $\log_2(1 + \text{SNR})$ for mutual information in coded systems).

As described in the previous section, when the antenna weights \mathbf{g} are selected, there is an associated point $(\text{SNR}_1, \text{SNR}_2, \dots, \text{SNR}_K)$ in K dimensional “SNR-space” that describes the associated SNRs experienced at the receivers in the system. Moreover, given the transmitter power constraint, there is a well-defined surface that defines the boundary of those points that are attainable. We refer to this frontier of

achievable points as the “transmitter operating characteristic” (TOC) for the realized channel and power constraint. As will become apparent, a transmitter operates efficiently if and only if it results in an SNR vector lying on the TOC.

Using $K = 2$ receivers for illustration, the TOC can be described as the set of received SNR pairs $(\text{SNR}_1, \text{SNR}_2)$ for which SNR_1 is maximized subject to various thresholds on SNR_2 . This frontier is equivalently traced out by maximizing

$$\alpha_1 \text{SNR}_1 + \alpha_2 \text{SNR}_2 \quad (5.1)$$

with various nonnegative weights α_1 and α_2 , again subject to the system power constraint. In enumerating points on the TOC, also note that it is only useful to send energy in a direction in the span of \mathbf{h}_1 and \mathbf{h}_2 .

It can be shown from the characterization of the TOC that the two pairs,

$$\left(\|\mathbf{h}_1\|^2 \frac{\mathcal{P}}{\mathcal{N}_0}, \frac{|\mathbf{h}_1^\dagger \mathbf{h}_2|^2 \mathcal{P}}{\|\mathbf{h}_1\|^2 \mathcal{N}_0} \right) \quad \text{and} \quad \left(\frac{|\mathbf{h}_1^\dagger \mathbf{h}_2|^2 \mathcal{P}}{\|\mathbf{h}_2\|^2 \mathcal{N}_0}, \|\mathbf{h}_2\|^2 \frac{\mathcal{P}}{\mathcal{N}_0} \right), \quad (5.2)$$

which correspond to beamforming directly to each of the first and second receivers, respectively, must lie on the TOC. This follows because one of the receivers experiences the maximum possible SNR in each case.

Returning to the weighted sum of SNRs formulation of (5.1), these two points correspond to $\alpha_2 = 0$ or $\alpha_1 = 0$. On the other hand, when $\alpha_1 = \alpha_2$, we wish to find a beamforming vector \mathbf{g} to maximize the sum of SNRs,

$$\max_{\mathbf{g}: \|\mathbf{g}\|^2=1} |\mathbf{h}_1^\dagger \mathbf{g}|^2 \frac{\mathcal{P}}{\mathcal{N}_0} + |\mathbf{h}_2^\dagger \mathbf{g}|^2 \frac{\mathcal{P}}{\mathcal{N}_0},$$

which is equivalent to

$$\max_{\mathbf{g}: \|\mathbf{g}\|^2=1} \|\mathbf{H}\mathbf{g}\|^2 \frac{\mathcal{P}}{\mathcal{N}_0}.$$

This is the well-known matrix norm problem, which is solved by performing a singular value decomposition,

$$\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\dagger,$$




Figure 5-1: The curve shown describes a typical frontier of achievable received SNR pairs when transmitting common information from an 8-element array to two receivers, at an input SNR per link of 5 dB. SNR pairs are achievable if and only if they lie on or inside this transmitter operating characteristic (TOC). Various operating points of interest are also shown.

and letting \mathbf{g} be the first column of \mathbf{V} . Then

$$\text{SNR}_k = |u_{k,1}|^2 \lambda_{\max} \frac{\mathcal{P}}{\mathcal{N}_0}, \quad (5.3)$$

where λ_{\max} is the largest eigenvalue of $\mathbf{H}^\dagger \mathbf{H}$, or equivalently the square of the largest singular value of \mathbf{H} , and $u_{k,1}$ is the k th entry of the first column of \mathbf{U} . For any other combination of (α_1, α_2) weights, the same procedure can be performed after first premultiplying \mathbf{H} by $\text{diag}(\sqrt{\alpha_1}, \sqrt{\alpha_2})$.

The TOC curve for a particular channel realization and power constraint is depicted in Fig. 5-1. Note that the SNRs are plotted in normal units rather than in dB. This will aid in geometric interpretations and properties; for example, the two-receiver achievable region for beamforming, when plotted in this way, is always convex. We

will prove this statement and discuss its implications in the later discussion on optimality (in Section 5.4.2). For now, this says that timesharing among two operating points on the TOC can not improve the time-average SNR.

By its location on the boundary of attainable SNR pairs, any point on the TOC represents a strategy where the transmitter is operating efficiently. To select among them, it is up to the system designer to supply a particular performance criterion that is appropriate to the given application.

Many operating points of interest can be developed from the TOC. The two circles ‘ \circ ’ correspond to the single-user beamforming points (5.2). To maximize the *minimum* performance among receives, discussed in the previous section as a way to ensure that both receivers achieve sufficient quality, one operates at the intersection of the TOC with the line $\text{SNR}_1 = \text{SNR}_2$; in Fig. 5-1 this point is indicated via the symbol ‘ ∇ ’. In scenarios where the line does not intersect the solid TOC curve, we operate at the nearest of the points (5.2). In other cases, maximizing the *average* (or, equivalently, total) SNR over all receivers is more appropriate. This is achieved by operating at the point where the TOC has slope -1 ; in Fig. 5-1 this point is indicated via the symbol ‘ \diamond ’, and corresponds to weights in (5.1) satisfying $\alpha_1 = \alpha_2$. A similar operating point can be found for maximizing the sum of mutual information across receivers by regraphing the TOC in terms of $\log_2(1 + \text{SNR})$; we saw how this is useful for maximizing ergodic capacity when the transmitter can code across many channel realizations.

Operating points other than the max-min point (‘ ∇ ’) are useful even when signaling over individual channel realizations. In some applications like voice transmission, one receiver may have higher fidelity requirements than the other. Other times, it may be important that information gets across to one receiver very quickly. In these cases, after the data is sent at a high rate that the first receiver can understand, additional symbols can be sent (perhaps along a different beamforming direction) to the second receiver. This is the idea behind Shulman and Feder’s static broadcasting [58], which was developed in a very general, information theoretic model. A practical implementation may include the use of rate-compatible punctured codes [35]. First, a high-rate, punctured code is transmitted that the first receiver can decode. Then the missing bits are sent, which combine with the first set to form the lower-rate code for the second receiver. These rate-compatible codes sacrifice very little optimality over the best known codes of the same rates (at least as of the publication of [35]).

5.3 Maximizing Average SNR per Receiver

We now proceed to answer some quantitative questions about the performance of multicast, using the operating point that maximizes the sum of SNRs to the K receivers. This point, achievable with beamforming, is amenable to analysis and gives an upper bound on the average per-receiver SNR. In this way, it provides information on where multicast scenarios may fall between the extremes of single-receiver transmission and communication with all possible receivers. Throughout, we assume that all receivers have the same Rayleigh fading distribution.

This discussion also serves to illustrate advantages and disadvantages of beamforming strategies. If the channel coefficients undergo independent, ergodic channel variations then the operating point under consideration maximizes the time-average SNR among all receivers. This offers a low-complexity approximation to maximizing the common ergodic capacity among receivers and, as we will see, achieves significant gains over scenarios where the transmitter does not have channel knowledge. On the other hand, when signaling over a single fading realization, these strategies often provide some receivers with very good performance at the expense of others. This makes the outage characteristic degrade rapidly as more receivers are added.

5.3.1 Average Performance Per Receiver

We begin by investigating average SNR per receiver, without regard to how performance is actually distributed among the different receivers. This will help characterize the potential of multicasting, and in particular the value of using channel information available at the transmitter as the number of receivers grows.

The properties of interest can be derived by analyzing the eigenvalues of certain random matrices. To achieve the maximum sum of SNRs, the beamforming vector \mathbf{g} is set to the eigenvector corresponding to the maximum eigenvalue of $\mathbf{H}^\dagger \mathbf{H}$. The average SNR per receiver then scales with the largest eigenvalue,

$$\text{Average SNR per receiver} = \frac{\lambda_{\max}(\mathbf{H}^\dagger \mathbf{H})}{K} \cdot \frac{\mathcal{P}}{\mathcal{N}_0}. \quad (5.4)$$

We then exploit that with Rayleigh fading, the matrix $\mathbf{H}^\dagger \mathbf{H}$ has a complex Wishart distribution [48] when $K \geq M$; when $K < M$ it is the matrix $\mathbf{H} \mathbf{H}^\dagger$ that is Wishart distributed.

When the number of antennas M and receivers K are moderate to large, we can take advantage of asymptotic properties of Wishart matrices. It can be shown that when M and K approach infinity in such a way that the ratio M/K of transmitter antenna elements per receiver approaches a positive constant, then the largest eigenvalue of the associated Wishart matrix converges almost surely [28, 18], resulting in

$$\text{Average SNR per receiver} \xrightarrow{a.s.} \left(1 + \sqrt{\frac{M}{K}}\right)^2 \frac{\mathcal{P}}{\mathcal{N}_0}. \quad (5.5)$$

This asymptotic behavior is shown by the solid curve in Fig. 5-2, from which we see that the SNR growth is effectively linear in the numbers of antenna elements/receiver ratio M/K for moderate to high ratios. Moreover, when the number of antenna elements M is significantly larger than the number of receivers K , there is a gain of approximately 3 dB in SNR for every doubling of M . We stress that the limit in (5.5) is no longer random, but rather a deterministic result for all channel realizations.

It is also worth emphasizing that a ratio of $M/K = 0$ means that M grows much more slowly than K , i.e., $M = o(K)$. A special case corresponds to using a fixed number of transmit antenna elements M while allowing the number of receivers K to increase to infinity. Because a transmitter can not effectively tailor a beamforming strategy to a very large number of receivers, it is not surprising that this ratio leads to an average value of $\mathcal{P}/\mathcal{N}_0$, the same as if channel information were not available.

Also shown in Fig. 5-2 are expected values for representative scenarios involving antennas with finitely many elements and finite receiver populations (using Monte Carlo simulations). As the plot reflects, the asymptotic behavior of (5.5) is approximated reasonably closely for even moderate values of M and K .

For finite values of M and K , the average SNR per receiver is a random variable whose value depends on the realized channel. If more accurate performance statistics for this random distribution are desired, it is possible to calculate the probability distribution of the possible values the SNR may take on. In particular, the joint distribution of all the eigenvalues λ_i of a Wishart matrix $\mathbf{H}^\dagger \mathbf{H}$, where \mathbf{H} has i.i.d. Gaussian entries of variance one, is [18]

$$f_{\lambda_1, \lambda_2, \dots, \lambda_M}(\lambda_1, \lambda_2, \dots, \lambda_M) = \frac{e^{-\sum_{i=1}^M \lambda_i} \prod_{i=1}^M \lambda_i^{K-M} \prod_{i < j} (\lambda_i - \lambda_j)^2}{\prod_{i=1}^M \Gamma(K - i + 1) \Gamma(M - i + 1)}, \quad (5.6)$$




Figure 5-2: Expected average SNR per receiver for various values of M/K and an input SNR per link of 0 dB. The solid line shows the deterministic asymptotic values when both M and K go to ∞ with the ratio M/K held fixed. The dashed curves denote representative points corresponding to finite M and K for $K = 4$ (\diamond) and $K = 8$ (∇), from simulations.

where

$$\Gamma(b) = \int_0^{\infty} t^{b-1} e^{-t} dt$$

denotes the usual Gamma function. Following Edelman [18], the density of the largest eigenvalue can be computed by integrating over all but one of the λ_i , and dividing by $(M-1)!$ to remove the arbitrary ordering of the eigenvalues. When $M = 2$, the resulting probability density for the largest eigenvalue is

$$f_{\lambda}(\lambda) = \frac{e^{-\lambda} \lambda^{K-2} [\lambda^K e^{-\lambda} - K \lambda^{K-1} e^{-\lambda} + (\lambda^2 + (K-1)(K-2\lambda)) \gamma(K-1, \lambda)]}{(K-1)!(K-2)!}, \quad (5.7)$$

where

$$\gamma(b, a) = \int_0^a t^{b-1} e^{-t} dt \quad (5.8)$$

is the incomplete Gamma function. From these probability functions and (5.4), it is possible to numerically calculate detailed average SNR statistics over the ensemble of possible channel realizations.

5.3.2 Individual Receiver Performance

While the average SNR per receiver may be a useful characterization of overall system performance, it does not reflect the behavior experienced by any individual receiver in the system. In this section, we focus on the individual-receiver outage and ergodic capacity.

To determine the distribution of an individual receiver's SNR under maximum sum of SNRs beamforming, we begin by repeating (5.3):

$$\text{SNR}_k = |u_{k,1}|^2 \lambda_{\max} \frac{\mathcal{P}}{\mathcal{N}_0},$$

where λ_{\max} is the largest eigenvalue in a Wishart-distributed matrix. Also, $u_{k,1}$ is an entry from the random circular unitary matrix \mathbf{U} from the singular value decomposition of \mathbf{H} . The probability density of $|u_{k,1}|^2$ is (see, e.g., [50])

$$f_{|u_{k,1}|^2}(\mu) = \begin{cases} (K-1)(1-\mu)^{K-2} & 0 < \mu < 1, \\ 0 & \text{otherwise.} \end{cases} \quad (5.9)$$

The marginal distribution for SNR_k can then be computed since random variables λ_{\max} and $|u_{k,1}|^2$ are independent — the principal eigenvector of $\mathbf{H}\mathbf{H}^\dagger$ has no preferred direction [19]. In the limiting case of $K \rightarrow \infty$ and M finite, it is straightforward to verify that SNR_k has the same exponential distribution as for a beamforming strategy that ignores side information. In principle, the SNR distribution can be computed analytically for any number of antennas or receivers. These computations quickly become very cumbersome, however, so in the discussion below, we plot results from simulations.

Maximum sum of SNRs beamforming performs well when signaling over many fading realizations. In Fig. 5-3, we plot the ergodic capacity (equal for all receivers) when multicasting from an 8-element array to as many as twenty receivers. For comparison, we also plot the performance of an ideal space-time code that does not take channel information into account. Note that at this input SNR level, the transmitter




Figure 5-3: Single-user ergodic capacity when multicasting a stream from an 8-element array to a number of receivers, at an input SNR per link of 5 dB. Also shown is a curve for a space-time code that does not make use of channel knowledge and achieves received SNR = $\|\mathbf{h}\|^2/M \cdot \mathcal{P}/\mathcal{N}_0$.

can communicate with twenty receivers simultaneously at a higher rate than is available by repeating the stream to two receivers separately with round-robin scheduling (at half the single-user rate for each).

The outage experienced during individual channel realizations does not fare as well. In Fig. 5-4, we see that the outage probability for maximum sum of SNRs beamforming degrades considerably as more receivers are added, although it does remain superior to transmission from a single antenna element. Other beamforming solutions may do somewhat better, but the outage characteristic will still suffer as the number of receivers gets large. This is because the coherent combining that occurs with beamforming also induces nulls at one or more geographic locations. Therefore, beamforming strategies are most useful when the number of receivers is fairly small or when performance is averaged across many channel realizations.




Figure 5-4: Single-user outage probabilities when multicasting a stream from an 8-element array to a number of receivers, at an input SNR per link of 5 dB. Also shown are curves for a space-time code that does not make use of channel knowledge.

5.4 General Space-Time Multicast Coding

We now turn to more general transmission schemes, which we refer to as space-time multicast coding. They can have higher complexity than beamforming, but are able to achieve a more equitable distribution of performance among the different receivers. Furthermore, we show that they achieve all possible operating points from the mutual information point of view.

5.4.1 Optimal Structures

In space-time multicast coding, the outputs at the different antenna elements can be described using an arbitrary covariance matrix. There are many possible implementations, but the structure of Fig. 5-5 is particularly useful for analysis. The data stream is encoded to produce a complex Gaussian sequence of coded symbols that is i.i.d., zero-mean, circularly symmetric, and has variance \mathcal{P} . Such encoders appear




Figure 5-5: Possible structure for space-time multicast coding.

often in the information theory literature. This sequence is then split into a number of parallel sequences and then undergoes a linear transformation described by an arbitrary matrix \mathbf{G} to produce the antenna element outputs. Note that this procedure reduces to beamforming in the special case where \mathbf{G} is a single column vector. To satisfy the power constraint, we impose $\text{trace}\{\mathbf{G}^\dagger \mathbf{G}\} \leq 1$.

We first show that this structure is sufficient to achieve all possible operating points for a coded system, and then go on to describe properties and interpretations.

Proposition 3 *Suppose a transmitter sends information from an M -element array to K receivers. The entire frontier of efficient operating points, in terms of mutual information K -tuples, is achievable by space-time multicast coding as described in Fig. 5-5.*

To prove this, note that space-time multicast coding sends a zero-mean, jointly Gaussian vector of antenna element outputs \mathbf{x} with covariance $\Gamma_{\mathbf{x}} = \mathbf{G}\mathbf{G}^\dagger\mathcal{P}$. The mutual information at receiver k is then equal to [50]

$$\log_2 \left(1 + \frac{\mathbf{h}_k^\dagger \Gamma_{\mathbf{x}} \mathbf{h}_k}{\mathcal{N}_0} \right). \quad (5.10)$$

Now consider any other scheme. The mutual information at receiver k will be

$$\begin{aligned} I(y_k; \mathbf{x}) &= H(y_k) - H(y_k | \mathbf{x}) \\ &= H(y_k) - \log_2(2\pi e \mathcal{N}_0), \end{aligned} \quad (5.11)$$

where $H(\cdot)$ denotes the entropy of a random variable. The vector of antenna element outputs \mathbf{x} should be zero-mean, because any power that is used in the mean will not contribute to the mutual information. Using the Cholesky factorization of the covariance matrix $\Gamma_{\mathbf{x}}$, the vector \mathbf{x} can always be written as

$$\mathbf{x} = \mathbf{L}\mathbf{s},$$

where \mathbf{s} is a length- M vector of uncorrelated random variables, each with variance \mathcal{P} , and \mathbf{L} is a lower-triangular matrix such that $\Gamma_{\mathbf{x}} = \mathbf{L}\mathbf{L}^\dagger\mathcal{P}$ and $\text{trace}\{\mathbf{L}^\dagger\mathbf{L}\} = 1$. Receiver k 's output will have variance

$$\sigma_{y_k}^2 = \sum_{m=1}^M |\mathbf{h}_k^\dagger \mathbf{l}_m|^2 \mathcal{P} + \mathcal{N}_0,$$

where \mathbf{l}_m is the m th column of \mathbf{L} . Among all random variables with this variance, the entropy, and therefore the mutual information in (5.11), is maximized with a Gaussian distribution [14]. This can be achieved with an i.i.d. Gaussian vector \mathbf{s} . Since this same distribution maximizes the mutual information for all receivers (given a particular $\Gamma_{\mathbf{x}}$), such a Gaussian vector is optimal. The overall system then becomes equivalent to space-time multicast coding.

This structure is also optimal when coding over ergodic variations of the channel. When the optimal input distribution is equivalent for all channel realizations, the maximum achievable rate for a receiver is the expected value of mutual information (shown in [46] for general channels and applied to fading channels in [31]). For the case of space-time multicast codes, the distribution on \mathbf{s} is the same for all receivers and all channel realizations. Since an arbitrary \mathbf{G} achieves all instantaneous operating points, we also achieve all operating points on the ergodically-varying channel.

Although we have not technically defined the set of mutual information K -tuples as a rate region, we should still double check whether timesharing can expand the region. Consider two matrices \mathbf{G}_1 and \mathbf{G}_2 used for space-time multicast, resulting in mutual information vectors $\log_2(1 + \boldsymbol{\gamma}_1)$ and $\log_2(1 + \boldsymbol{\gamma}_2)$, respectively. If the first

scheme is used a fraction β of the time, we get through timesharing,

$$\begin{aligned} \beta \begin{bmatrix} \log_2(1 + \gamma_{1,1}) \\ \log_2(1 + \gamma_{1,2}) \\ \vdots \\ \log_2(1 + \gamma_{1,k}) \end{bmatrix} &+ (1 - \beta) \begin{bmatrix} \log_2(1 + \gamma_{2,1}) \\ \log_2(1 + \gamma_{2,2}) \\ \vdots \\ \log_2(1 + \gamma_{2,k}) \end{bmatrix} \\ &= \begin{bmatrix} \beta \log_2(1 + \gamma_{1,1}) + (1 - \beta) \log_2(1 + \gamma_{2,1}) \\ \beta \log_2(1 + \gamma_{1,2}) + (1 - \beta) \log_2(1 + \gamma_{2,2}) \\ \vdots \\ \beta \log_2(1 + \gamma_{1,k}) + (1 - \beta) \log_2(1 + \gamma_{2,k}) \end{bmatrix}. \end{aligned}$$

However, using Jensen's inequality, this vector is the same or inferior for every receiver to the vector

$$\begin{bmatrix} \log_2(1 + \beta\gamma_{1,1} + (1 - \beta)\gamma_{2,1}) \\ \log_2(1 + \beta\gamma_{1,2} + (1 - \beta)\gamma_{2,2}) \\ \vdots \\ \log_2(1 + \beta\gamma_{1,k} + (1 - \beta)\gamma_{2,k}) \end{bmatrix},$$

which is achievable by space-time multicast coding with the matrix

$$\mathbf{G} = \begin{bmatrix} \sqrt{\beta}\mathbf{G}_1 & \sqrt{1 - \beta}\mathbf{G}_2 \end{bmatrix}.$$

Similar reasoning shows that timesharing between more than two points does not add to the region, either. \square

The following interpretation of space-time multicast coding provides a connection to the SNR operating characteristic described earlier in Section 5.2. Consider the linear transformation in Fig. 5-5 as beamforming each of the parallel sequences along a direction corresponding to a column of \mathbf{G} . Let \mathbf{G} in turn be written as

$$\mathbf{G} = \begin{bmatrix} \alpha_1\mathbf{g}_1 & \alpha_2\mathbf{g}_2 & \dots & \alpha_N\mathbf{g}_N \end{bmatrix},$$

where \mathbf{g}_n are unit-length vectors, i.e., $\|\mathbf{g}_n\|^2 = 1$ for all n . To satisfy the power

constraint, the α_n 's must be chosen so that

$$\sum_{n=1}^N |\alpha_n|^2 \leq 1.$$

The mutual information to the k th receiver is given by (5.10). Note that the second term inside the logarithm is

$$\begin{aligned} \frac{\mathbf{h}_k^\dagger \Gamma \mathbf{x} \mathbf{h}_k}{\mathcal{N}_0} &= \mathbf{h}_k^\dagger \mathbf{G} \mathbf{G}^\dagger \mathbf{h}_k \frac{\mathcal{P}}{\mathcal{N}_0} \\ &= \|\mathbf{h}_k^\dagger \mathbf{G}\|^2 \frac{\mathcal{P}}{\mathcal{N}_0}, \\ &= \sum_{n=1}^N |\alpha_n|^2 |\mathbf{h}_k^\dagger \mathbf{g}_n|^2 \frac{\mathcal{P}}{\mathcal{N}_0}. \end{aligned} \tag{5.12}$$

This is equal to the time-average SNR had the stream been transmitted along the columns of \mathbf{G} at different times (with each being used a fraction of time $|\alpha_n|^2$). Unlike with timesharing, however, the coded rate corresponding to this “equivalent SNR” is actually achievable. Beamforming falls out as a special case when the covariance matrix $\Gamma_{\mathbf{x}}$ has rank one.

This motivates the use of “equivalent SNRs,” such as (5.12), as a convenient parametrization for mutual information. In this way, the performance of space-time multicast codes is seen as a kind of averaging between beamforming strategies. The achievable region, in terms of equivalent SNR K -tuples, becomes the convex hull of the beamforming region. It is important to keep in mind that for higher-rank covariance matrices, the equivalent SNR simply represents the SNR of an additive white Gaussian noise channel with the same mutual information, and in general is not a true SNR achievable by uncoded systems. Later, we will develop implementations that are more amenable to uncoded transmission.

This discussion also relates to the theory of space-time codes that do not incorporate channel knowledge at the transmitter. Without channel information, the transmitter can still set \mathbf{G} to be a scaled identity matrix, or any other unitary matrix, and achieve the set of equivalent SNRs

$$\text{SNR}_k = \frac{\|\mathbf{h}_k\|^2 \mathcal{P}}{M \mathcal{N}_0}. \tag{5.13}$$

It also maximizes the ergodic capacity in these situations [63]. Exact implementations

would be complex, so various space-time codes have been developed to approach this performance at lower complexity. The goal of (5.13) is achievable for space-time block codes designed for two transmit antennas [1] but has not been reached for larger sizes [62]. It also shows up as an ideal “matched filter bound” in an early version of space-time trellis coding sometimes called delay diversity [78]. We will demonstrate space-time multicast code implementations that approach the equivalent SNR to the extent that these other types of designs do.

5.4.2 Beamforming Versus Higher-Rank Covariances

We have identified beamforming as a subset of space-time multicast codes where the covariance of the antenna outputs has rank one. Such strategies have low complexity and are compatible with most types of coded or uncoded modulation, but have poor outage characteristics when the number of receivers grows large. In this section, we investigate when beamforming is sufficient from an optimality standpoint, and when higher-rank covariances are necessary to achieve certain operating points.

Two Antenna Elements or Two Receivers

In earlier sections, we found that beamforming strategies work well when the number of receivers is small. Using the concept of equivalent SNRs, we can now make this statement more precise. We show that beamforming is entirely sufficient for multicasting to two receivers, and then look at where it breaks down as the number of receivers is increased.

Proposition 4 *Suppose a transmitter sends information from an M -element array to two receivers. Then all efficient operating points can be achieved using a rank-one covariance; in other words, by a beamforming strategy.*

Space-time multicast coding is already known to be optimal. Since it averages the effective SNR for each receiver over several beamforming directions, we can prove the statement above by showing that the set of $(\text{SNR}_1, \text{SNR}_2)$ pairs achievable by beamforming is convex. One way to do this is to simply enumerate all of these points.

First, recall that the transmitter should only send in directions that are in the span of the two receivers’ channel vectors; components outside of this subspace will simply produce nulls at the receivers and waste power. Using the Gram-Schmidt

procedure, this space can be parametrized by a component in the direction of the first receiver's channel vector, and a second component orthogonal to this. Therefore, the M -antenna problem can be reduced to an equivalent two-antenna problem with lower-triangular channel matrix and channel vectors

$$\mathbf{h}_1 = \begin{bmatrix} L_{1,1} \\ 0 \end{bmatrix} \quad \text{and} \quad \mathbf{h}_2 = \begin{bmatrix} L_{2,1} \\ L_{2,2} \end{bmatrix}.$$

If the transmitter is operating at the power constraint $\|\mathbf{g}\|^2 = 1$, the beamforming vector \mathbf{g} can be parametrized as

$$\mathbf{g} = \begin{bmatrix} \cos \theta \\ e^{j\phi} \sin \theta \end{bmatrix},$$

where the two angles θ , which is in the range $[0, \pi/2)$, and ϕ , in $[0, 2\pi)$, produce the relative gain and phase. We can then enumerate all of the SNR pairs that are attainable by beamforming:

$$\text{SNR}_1 \frac{N_0}{\mathcal{P}} = |L_{1,1}|^2 \cos^2 \theta, \quad (5.14)$$

$$\text{SNR}_2 \frac{N_0}{\mathcal{P}} = |L_{2,1}|^2 \cos^2 \theta + |L_{2,2}|^2 \sin^2 \theta + |L_{2,1}L_{2,2}| \sin(2\theta) \cos(\phi - \phi_2), \quad (5.15)$$

where ϕ_2 is defined using

$$L_{2,1}L_{2,2}^\dagger = |L_{2,1}| |L_{2,2}| e^{-j\phi_2}.$$

Performance is clearly maximized by choosing $\phi = \phi_2$ so that the final cosine term in (5.15) is equal to one. The transmitter operating characteristic curve in Fig. 5-1 is produced by following this trajectory as well as a similar one when the ordering of the two receivers is reversed. The solid portion, which is equivalent to maximizing a (non-negative) weighted sum of SNRs, is the intersection of these two curve. Any point in SNR-space that is inside these boundaries can be achieved by transmitting below the power constraint. By taking second derivatives of (5.14)–(5.15) along the boundary, it can be shown that the overall region is convex. \square

This shows that beamforming is sufficient in the two-receiver case with any number of transmit antenna elements. On the other hand, we know that it is not optimal for a

two-element array and a large number of receivers. To understand where this behavior changes, the method above can be extended to two transmit antenna elements and any number of receivers. Using the same parametrization as before, we have

$$\mathbf{h}_k = \begin{bmatrix} L_{k,1} \\ L_{k,2} \end{bmatrix}, \quad k = 2, 3, \dots, K$$

and

$$\text{SNR}_k \frac{N_0}{\mathcal{P}} = |L_{k,1}|^2 \cos^2 \theta + |L_{k,2}|^2 \sin^2 \theta + |L_{k,1} L_{k,2}| \sin(2\theta) \cos(\phi - \phi_k),$$

$$k = 2, 3, \dots, K.$$

Consider this for $K = 4$ receivers. The three ϕ_k are parameters of the realized channel vectors. Take the case of $\phi_2 = 0$, $\phi_3 = 2\pi/3$, and $\phi_4 = 4\pi/3$. Holding θ constant and alternating between $\phi = 0$ and $\phi = \pi$ makes all three cosine terms average to zero, while it is impossible to make all of them simultaneously nonnegative for any single ϕ . Therefore, the equivalent point in SNR-space corresponding to this alternating strategy is achievable only through space-time multicast coding with a rank two covariance. Beamforming from a two-element array is apparently no longer sufficient when there are four or more receivers.

Arbitrary Numbers of Antennas and Receivers

For larger systems, parametrizations of all beamforming operating points such as (5.14)–(5.15) become very cumbersome. However, with some additional geometric insight, we can generalize the two antenna element results and conjecture that for M transmit antenna elements, beamforming becomes suboptimal when transmitting to $2M$ or more receivers. This is done with essentially a dimension-counting argument.

In general, beamforming with an M -element array requires specifying $2M$ real parameters: the individual gains and phases applied to the different antenna inputs. For SNR purposes, however, there are really only $2M - 1$ degrees of freedom, because an overall phase can be factored out without affecting performance. This implies that the region achievable by beamforming has dimension no greater than $2M - 1$.

Space-time multicast coding achieves any convex combination of points in SNR-space that are achievable by beamforming. Mathematically, this is a convex hull operation [56]. If we can show that the resulting region has a higher dimension than

$2M - 1$, then clearly beamforming is not sufficient. One way to do this is to show that there are at least $2M$ linearly independent points in the beamforming region.

We conjecture that this is true with probability one when the number of receivers is at least $2M$. Consider the points in SNR-space corresponding to the K single-user beamforming directions. When beamforming in the direction of the k th receiver, the vector of received SNRs is

$$\begin{bmatrix} |\mathbf{h}_k^\dagger \mathbf{h}_1|^2 \\ |\mathbf{h}_k^\dagger \mathbf{h}_2|^2 \\ \vdots \\ |\mathbf{h}_k^\dagger \mathbf{h}_K|^2 \end{bmatrix} \frac{1}{\|\mathbf{h}_k\|^2} \cdot \frac{\mathcal{P}}{\mathcal{N}_0}.$$

If we collect these vectors into a $K \times K$ matrix, its rank will be equal to the number of linearly dependent points in SNR-space achieved by these K particular beamforming directions.

Multiplying each column by a constant will not change the rank, so we therefore wish to find the rank of the matrix

$$\mathbf{B} \circ \mathbf{B}^*,$$

where

$$\mathbf{B} = \mathbf{H}\mathbf{H}^\dagger, \tag{5.16}$$

\mathbf{B}^* represents the conjugation (but not transpose) of \mathbf{B} , and “ \circ ” represents the element-by-element Hadamard product. Applying a singular value decomposition, $\mathbf{B} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\dagger$, the Hadamard product above can be taken as a particular submatrix of

$$(\mathbf{U} \otimes \mathbf{U}^*) \circ (\mathbf{\Sigma} \otimes \mathbf{\Sigma}^*) \circ (\mathbf{V}^\dagger \otimes \mathbf{V}^T),$$

where \otimes represents the Kronecker product [38]. By noting that \mathbf{B} is Hermetian (so that $\mathbf{U} = \mathbf{V}$) and carefully inspecting the individual elements, it can be shown that

$$\mathbf{B} \circ \mathbf{B}^* = \mathbf{B}_u \mathbf{B}_u^\dagger,$$

where \mathbf{B}_u consists of all K^2 possible columns of the form

$$\sqrt{\sigma_i \sigma_k} \mathbf{u}_i^* \circ \mathbf{u}_k.$$

The rank of the overall product is the same as the rank of \mathbf{B}_u . From (5.16), there can be at most M nonzero singular values σ_k , at most M^2 possible nonzero columns of \mathbf{B}_u , and therefore the maximum overall rank is $\min(K, M^2)$. Further general analysis appears difficult, but our simulations and analysis of special cases suggest that this maximum rank does hold true. We therefore conjecture that when $K \geq 2M$ then with probability one, there are at least $2M$ linearly dependent points in SNR-space, and consequently higher-rank covariance matrices are necessary to achieve all possible points in SNR-space.

5.4.3 Implementation Issues

Multicasting With Arbitrary Coding and Modulation

The points achievable by higher-rank covariances are in general not available for arbitrary signaling, but rather are equivalent SNRs for rates achieved by particular vector-coded systems. More practical implementations may take their inspiration from existing space-time codes that are adapted to take advantage of channel knowledge.

For example, orthogonal space-time block codes can easily be converted to use any covariance matrix. These codes are compatible with arbitrary modulation and scalar coding and have simple detection algorithms. In the Alamouti scheme for two transmit antenna elements [1], the transmitter sends two symbols, $s[1]$ and $s[2]$, over two time periods:

$$\begin{aligned} \text{Time 1} & : \mathbf{x}[1] = \begin{bmatrix} 1 \\ 0 \end{bmatrix} s[1] + \begin{bmatrix} 0 \\ 1 \end{bmatrix} s[2] \\ \text{Time 2} & : \mathbf{x}[2] = \begin{bmatrix} 0 \\ 1 \end{bmatrix} s^*[1] - \begin{bmatrix} 1 \\ 0 \end{bmatrix} s^*[2]. \end{aligned} \quad (5.17)$$

A receiver with channel vector $\mathbf{h} = [h_1 \ h_2]^T$ gets

$$\begin{aligned} y[1] &= h_1^*s[1] + h_2^*s[2] + w[1] \\ y[2] &= h_2^*s^*[1] - h_1^*s^*[2] + w[2]. \end{aligned}$$

The receiver, knowing the channel, can recover the input symbols by taking linear combinations and conjugations,

$$\begin{aligned} \hat{s}[1] &= h_1y[1] + h_2^*y^*[2] \\ \hat{s}[2] &= h_2y[1] - h_1^*y^*[2], \end{aligned}$$

to achieve the ideal space-time coding SNR of $\|\mathbf{h}\|^2\mathcal{P}/2\mathcal{N}_0$.

The orthogonal signaling vectors in (5.17) are used because the channel is assumed not to be known at the transmitter. However, the procedure will work just as well with arbitrary vectors, \mathbf{g}_1 and \mathbf{g}_2 :

$$\begin{aligned} \text{Time 1} &: \mathbf{x}[1] = \mathbf{g}_1s[1] + \mathbf{g}_2s[2] \\ \text{Time 2} &: \mathbf{x}[2] = \mathbf{g}_2s^*[1] - \mathbf{g}_1s^*[2]. \end{aligned}$$

Now, instead of being sent on channels h_1^* and h_2^* , the symbols are sent on $\mathbf{h}^\dagger\mathbf{g}_1$ and $\mathbf{h}^\dagger\mathbf{g}_2$. The rest of the procedure works exactly the same as before but with these substitutions, achieving the received SNR

$$\text{SNR} = (\|\mathbf{h}^\dagger\mathbf{g}_1\|^2 + \|\mathbf{h}^\dagger\mathbf{g}_2\|^2) \frac{\mathcal{P}}{\mathcal{N}_0}.$$

If $\mathbf{g}_1 = \mathbf{g}_2 = \mathbf{h}/\|\mathbf{h}\|$, then the full single-user SNR of $\|\mathbf{h}\|^2\mathcal{P}/\mathcal{N}_0$ is achieved. For multicast streams, such coherent combining will usually not be possible for all receivers, so distinct vectors \mathbf{g}_1 and \mathbf{g}_2 will be used. Also note that it is possible to use different power distributions among the two \mathbf{g}_i vectors and effectively produce any weighted average of SNRs between the two, achieving all the points we expect from space-time multicast coding that uses a rank $N = 2$ covariance (5.12).

It is important to note that although the Alamouti scheme is for a two-element antenna, our adapted version for arbitrary transmission vectors will work for any number of antenna elements, as long as we only wish to average two beamforming

directions. In this way, it can achieve the equivalent SNR of space-time multicast coding with any rank-two covariance matrix, but now with arbitrary modulation and coding. Other orthogonal space-time block codes can be adapted for covariance matrices above rank two with a procedure analogous to that outlined above. Instead of averaging the SNRs over the channel components h_i , we average over the SNRs of the inner products $\mathbf{h}^\dagger \mathbf{g}_i$. Unfortunately, all orthogonal space-time codes of this type with rank higher than two incur a rate penalty [62]. Still, optimizing the beamforming directions rather than simply using orthogonal vectors can lead to significant SNR improvement. Other techniques such as space-time trellis codes can be similarly converted to achieve “diversity” over the $\mathbf{h}^\dagger \mathbf{g}_i$.

Finding Operating Points

What remains is a method for finding good operating points for space-time multicast coding. We concentrate here on maximizing the minimum performance among receivers for each channel realization, which leads to the highest coded rate that all receivers can understand. Recall that outage-based operating points such as this are where beamforming strategies are weakest.

In general, this problem represents a maximization of a concave function over a convex set,

$$\max_{\mathbf{G}: \text{trace}\{\mathbf{G}^\dagger \mathbf{G}\} \leq 1} \min_k \|\mathbf{h}_k^\dagger \mathbf{G}\|^2,$$

implying that every local maximum is also a global maximum [56]. This suggests that iterative optimization algorithms might be useful. Still, the convex set of all achievable points is rather complicated, which might make an exact approach difficult.

This becomes more tractable if broken down into the separate problems of finding unit-length column vectors for the matrix \mathbf{G} and corresponding weights on those vectors. Given a set of unit vectors, the convex domain is polyhedral, and the optimization can be converted to a linear programming problem that can be solved with the simplex method [56, 37], a standard linear optimization tool. For the unit vectors themselves, we will find that the single-user beamforming directions lead to good results. For example, for K receivers and the three received SNR vectors, $\boldsymbol{\gamma}_1$,

γ_2 , and γ_3 , the goal is

$$\max_{\alpha_1, \alpha_2, \alpha_3} \min_{k=1, \dots, K} (\alpha_1 \gamma_{1,k} + \alpha_2 \gamma_{2,k} + \alpha_3 \gamma_{3,k}),$$

where α_1 , α_2 , and α_3 are all nonnegative and sum to one. This can be reformulated by introducing as a new variable the max min SNR goal, α_4 :

Maximize α_4

given the constraints

$$\begin{aligned} \sum_{i=1}^3 \alpha_i \gamma_{i,k} - \alpha_4 &\geq 0, & k = 1, \dots, K \\ \alpha_1 + \alpha_2 + \alpha_3 &\leq 1 \\ \alpha_i &\geq 0, & i = 1, \dots, 4. \end{aligned} \tag{5.18}$$

With at most a couple sign changes of coefficients to get all the inequalities in the same direction, this fits the form of the Matlab command `linprog` and other simplex method implementations.

Two-Receiver Illustration

For the specific case of two receivers, an interesting result illustrates the relationship between sending common information and distinct information to two receivers. With distinct information, the (often suboptimal, but tractable) zero-forcing beamforming leads to SNRs of

$$\begin{aligned} \text{SNR}_1 &= \left(\alpha \|\mathbf{h}_1\|^2 - \alpha \frac{\|\mathbf{h}_1^\dagger \mathbf{h}_2\|^2}{\|\mathbf{h}_2\|^2} \right) \frac{\mathcal{P}}{\mathcal{N}_0} \\ \text{SNR}_2 &= \left((1 - \alpha) \|\mathbf{h}_2\|^2 - (1 - \alpha) \frac{\|\mathbf{h}_1^\dagger \mathbf{h}_2\|^2}{\|\mathbf{h}_1\|^2} \right) \frac{\mathcal{P}}{\mathcal{N}_0} \end{aligned}$$

where α is the fraction of power sent to the first receiver. Using power control to equalize the SNRs leads to

$$\text{SNR}_1 = \text{SNR}_2 = \frac{\|\mathbf{h}_1\|^2 \|\mathbf{h}_2\|^2 - \|\mathbf{h}_1^\dagger \mathbf{h}_2\|^2}{\|\mathbf{h}_1\|^2 + \|\mathbf{h}_2\|^2} \cdot \frac{\mathcal{P}}{\mathcal{N}_0}. \tag{5.19}$$

On the other hand, with space-time multicast coding using the two single-user beamforming directions,

$$\begin{aligned}\text{SNR}_1 &= \left(\alpha \|\mathbf{h}_1\|^2 + (1 - \alpha) \frac{\|\mathbf{h}_1^\dagger \mathbf{h}_2\|^2}{\|\mathbf{h}_2\|^2} \right) \frac{\mathcal{P}}{\mathcal{N}_0} \\ \text{SNR}_2 &= \left((1 - \alpha) \|\mathbf{h}_2\|^2 + \alpha \frac{\|\mathbf{h}_1^\dagger \mathbf{h}_2\|^2}{\|\mathbf{h}_1\|^2} \right) \frac{\mathcal{P}}{\mathcal{N}_0}\end{aligned}$$

where α is the fraction of power sent along the first receiver's direction. (This is also suboptimal, because we showed that for two receivers, single-rank beamforming is best.) We wish to optimize α to maximize the minimum SNR. Since increasing α always improves SNR_1 at the expense of SNR_2 , the best scenario is when we can make $\text{SNR}_1 = \text{SNR}_2$. If this is possible (that is, if the solution to $\text{SNR}_1 = \text{SNR}_2$ leads to an $0 \leq \alpha \leq 1$), then

$$\text{SNR}_1 = \text{SNR}_2 = \frac{\|\mathbf{h}_1\|^2 \|\mathbf{h}_2\|^2 + \|\mathbf{h}_1^\dagger \mathbf{h}_2\|^2}{\|\mathbf{h}_1\|^2 + \|\mathbf{h}_2\|^2} \cdot \frac{\mathcal{P}}{\mathcal{N}_0}. \quad (5.20)$$

Comparing (5.19) and (5.20), the only difference is in the sign of the cross term, which is essentially the deterministic correlation between the two realized channel vectors. When multiplexing separate data, correlation between channels is bad, because it causes interference that either degrades performance or is to be avoided. For multicast, correlation improves performance, avoiding the need to send redundant information.

5.4.4 Performance of Higher-Rank Covariance Matrices

Although the value of channel information decreases as the number of receivers gets large, our space-time multicast codes still exhibit a significant performance advantage for moderate-sized systems. We illustrate this for an example where an 8-element array multicasts a single stream to 8 receivers.

To communicate with all receivers reliably, we concentrate on maximizing the minimum of equivalent SNRs among them. Once again, this optimization tends to be very difficult in general, so we will constrain space-time multicast coding to using a weighted set of single-user beamforming directions, as discussed in Section 5.4.3. Optimal weights between the vectors were computed as in (5.18). Without channel




Figure 5-6: Weakest-receiver effective SNR, from simulations, when multicasting a stream from an 8-element array to 8 receivers, at an input SNR per link of 5 dB. The schemes shown are: “No Array”: single transmit antenna element; “No Tx Knowledge”: space-time coding with orthogonal matrix \mathbf{G} , “ST Multicast”: using channel knowledge with the weights chosen by the method of (5.18) .

knowledge, ordinary space-time codes would ideally choose a \mathbf{G} matrix with orthogonal columns.

We compare the outage performance with and without channel knowledge in Fig. 5-6. As expected, there is a large gain for both methods over not using an array. Even on this scale, however, channel-aware transmission noticeably outpaces ordinary space-time codes. The outage curves have similar shapes, but are separated by about 6 dB at 1% outage. To place this into context, the improvement if the transmitter could perfectly direct the stream to all receivers simultaneously would be $10 \log_2 8 \approx 9$ dB. For similar simulations with a 4-element array and 4 receivers, about 4 dB of the possible 6 dB advantage is preserved. Since we know these bounds are unattainable, the fact that we get a good deal of the way there speaks to the effectiveness of our methods and the usefulness of channel knowledge for multicast.

5.5 Multicast Within Larger Systems

We now return to the larger picture of a system with a number of data streams, both multicast and receiver-specific. Such systems must find ways for these streams to coexist through scheduling, spatial multiplexing, or both. In our vision, the scheduler doles out streams to different array processing subblocks for precoding, beamforming, or multicasting; these subblocks in turn work together to get the data across without undue interference. At a basic level, this consists of incorporating a multicast stream into the model of earlier chapters, as a “metastream” that has a number of intended receivers in its multicast group. This discussion brings together many of the techniques developed in this thesis and provides an overall vision for how such a system may operate.

5.5.1 Integration Among Array Processing Subblocks

If the array processing task is to transmit more than one stream at once, it must find a way not only to direct the data to its intended recipients but also not to cause interference at other receivers. For individual-receiver streams, the transmitter’s channel information enabled us to use spatial precoding and beamforming to accomplish this. We will find that similar techniques can reduce interference among multicast streams or between a multicast stream and several individual-receiver streams.

Beamforming-based separation works in much the same way as before. Any individual-receiver streams must set their beamforming directions to be orthogonal to all other active receivers’ channel vectors, including those in multicast groups. A multicast stream similarly needs to transmit orthogonally to all receivers not in the group. To find the proper space-time multicast coding parameters, the multicast subblock should first project each of the channel vectors in the group away from all receivers not in the group, and then optimize the group’s transmission scheme based on these new channel vectors. Power can be redistributed among the different streams and metastreams as needed. Note that because space-time codes that do not make use of channel knowledge are designed to spread their signal throughout the entire space of possible directions, they are not appropriate for spatial multiplexing with other streams in this way. This serves as an additional advantage of our space-time multicast codes.

Precoding, which we saw achieve significant improvements over zero-forcing beam-

forming in many cases, can be adapted in a more limited way. Recall that precoding works on an ordered set of streams, where later streams precompensate for interference from the earlier ones. Because a multicast stream must send the same message to a number of receivers, yet each will each receive a different linear combination of interference from earlier streams, it is difficult to set up a precoding procedure for a multicast group without suffering a rate loss. On the other hand, it is possible for later-ordered streams to use precoding to precompensate for interference from one or more multicast streams. This ordering also has the advantage that a multicast group can compute its transmission scheme and performance before dealing with the other receivers.

Putting this all together, we can group the active receivers into a number of groups based on the array processing of their associated data streams. Individual-receiver streams may go to the precoding subblock or may instead perform zero-forcing beamforming, for instance if their receivers do not support modulo-extended slicers. Each multicast stream has its own group of receivers. We can then partition the array processing into a global preprocessing step and more local signaling done within each group. First, order the total set of receivers such that the multicast groups are first, then beamforming, and finally precoding. Then, the preprocessing step could ensure that groups constrain their transmission to be orthogonal to channel vectors in other multicast or beamforming groups. The multicast and beamforming groups need not worry about causing interference for receivers in the precoding group, since any crossover interference will be removed by precoding. At this point, processing within each group can proceed as normal. What results is an effective channel matrix (before precoding itself) that is a mix between block diagonal (for the multicast and beamforming groups) and lower triangular (for the precoding group). For example, if there is a precoding group with four receivers, a beamforming group with two single-user streams, and a precoding group of two streams, then the possible non-zero entries

of this matrix are highlighted as

$$HG_B \sim \begin{bmatrix} \times & \times & \times & \times & & & & & \\ \times & \times & \times & \times & & & & & \\ \times & \times & \times & \times & & & & & \\ \times & \times & \times & \times & & & & & \\ & & & & \times & & & & \\ & & & & & \times & & & \\ & & & & & & \times & & \\ \times & \times & \times & \times & \times & \times & \times & & \\ \times & \times & \times & \times & \times & \times & \times & \times & \end{bmatrix}.$$

This procedure allows spatial multiplexing between different types of groups while not significantly changing the operation within each array processing subblock.

5.5.2 Integration at the Scheduling Layer

If there are more than a few potential streams, the system will likely also require the integrated scheduling of multicast and individual-receiver data. The simplest transmitters would timeshare between different types of data streams in a round-robin manner. However, this ignores the possibility of capturing some of the multiplexing gains we saw in earlier chapters. Recall that the potential throughput of systems increased severalfold as the number of receivers approached the number of transmit antenna elements. A more sophisticated system would attempt to use channel information to select appropriate groups of streams for spatial multiplexing. The scheduler then feeds the selected active streams into their respective array processing subblocks, where the types of multiplexing described above can occur.

The scheduling algorithm will depend on the delay tolerance of the data, how many intended receivers there are for each multicast stream, and whether precoding is used for individual-receiver streams. Because both precoding and zero-forcing beamforming require restricting the transmission of at least some streams, a good rule of thumb for an M -element array may be to have no more than M active receivers at any one time, unless all receivers belong to the same multicast stream. Beyond this, the most important issues will once again be reducing the potential interference between streams and, if the delay tolerance allows, selecting receivers whose channels are of high instantaneous quality. Multicast streams make dealing with both of these issues more difficult because the scheduler must satisfy all receivers of a particular

stream simultaneously.

One complexity-reducing solution would employ separate queues for multicast and individual-receiver streams and communicate no more than one multicast group at any single time. Once the active multicast group is selected, the scheduler can treat these receivers as if they were getting separate streams as it selects additional streams according to the algorithms of Chapter 4. For example, with moderate delay constraints, this selection would be to find a set of channel vectors that are nearly orthogonal. For a transmitter with an 8-element array, a typical timeslot may include a multicast stream with four receivers and three or four additional individual-receiver streams. Unless the multicast groups consist of a very small number of receivers, such a system is not likely to lose much in performance compared with a fully integrated scheduler. Among the remaining challenges would include building in fairness constraints to strike the right balance between the different types of streams.

Chapter 6

Conclusions and Future Work

In this thesis, we have discussed the design of various system components for a transmitter antenna array as well as a higher-level view of how these components interact. We found that a consideration of the channel parameters and input data stream properties can be very useful at both the scheduling and array processing levels. Although this may violate some of the principles of the traditional layered approach, the gains achieved by channel-aware scheduling or sophisticated spatial multiplexing imply that a rethinking may be in order.

In an effort to make our results applicable, we have centered the development around implementations, design choices, and analyzing the key issues involved with particular system tasks. Some of the major contributions include:

- An overall framework for the integrated design of transmitter antenna array systems. Of particular importance is the partitioning into scheduling and array processing tasks, as outlined in the introduction. We found this led to convenient problem formulations yet allowed for sufficient interaction among components to approach the potential of the array. It also helped make clear the different options for placing complexity throughout the system and the associated performance of these choices.
- At the array processing level, we added new insights and extensions to spatial precoding. Building upon a series of recent results, our work concentrated on variations of the basic precoding procedure to satisfy system goals for different data classes, channel modeling assumptions, and modulation techniques. This included changing the ordering of streams for different types of data, adapting symbol constellations based on the interference distribution, and extending

precoding to multiuser intersymbol interference channels.

- We demonstrated how channel-aware scheduling techniques have the potential to increase performance for a number of data types. In particular, we used the delay tolerance of the various streams to provide the scheduler with flexibility and constraints in rearranging the ordering and grouping of streams. Even for a small amount of flexibility, an appropriate grouping can help the array processing achieve much better reliability and higher rates. More sophisticated scheduling can also enable lower complexity at the array processing level.
- For the multicasting of common data streams, we developed optimal signaling techniques as well as more practical implementations. Among these were two important methods, useful in different regimes, representing beamforming and an adaptation of space-time codes to accommodate channel knowledge. In this process, we helped define what it means for the transmitter to operate efficiently in terms of balancing performance to the multiple recipients.

In this way, we have considered many problems in detail, yet within an overall structure in which individual algorithms may be included or replaced depending upon the needs of an individual system.

Future work can continue development within this structure. In addition to numerous possible algorithmic improvements, this may take the form of expanding into additional components at either end of the signal chain.

On the physical channel side, systems may be developed to more tightly incorporate the mechanisms for attaining channel information. This channel estimation takes up system resources not accounted for in our discussion. Information about the current data streams and previous channel states could potentially be used to request when and how much channel information is needed. Another important goal would be to further characterize the effect that partial, rather than perfect, channel information has on the main components. Yet another direction is to include more detail in the channel model, such as the movement of mobile receivers relative to the transmitter, and then adapt scheduling algorithms to these models.

At the other side would be a further awareness of the data streams and their performance goals. We have attempted to maximize rate or reliability-related goals while respecting certain coarse delay constraints. The next step may be a more specific investigation into the fundamental delay/throughput tradeoffs of spatial multiplexing

systems. A different, but related, direction involves the development of scheduling algorithms to achieve more formal quality-of-service measures such as packet drop rates and delay guarantees. As explained in Chapter 4, we anticipate that although such goals are not necessarily a good match to array fading channels, in practice it may be possible to meet them due to the robustness of scheduling over a constrained set of available channel vectors.

Issues of a more global nature appear when a wireless network contains multiple array transmitters. For example, signals from one transmitter will cause interference on the communication from others. Cellular systems often mitigate this interference by partitioning receivers and bandwidth resources among separate cells. A more effective approach would use greater coordination across transmitters. At a conceptual level, the different antenna elements from all of the transmitters may be considered as one larger virtual array, upon which many of the techniques discussed in this thesis may be applied. However, as networks extend from several cells to entire metropolitan areas and more, a comprehensive implementation quickly becomes unmanageable. It is also unnecessary, because interference from a single transmitter will be negligible except in a small geographic area; mathematically, the channel matrix from the whole virtual array to all of the receivers will be very sparse. The main network-level problem is therefore to find some reasonable compromise between partitioning and full coordination.

Appendix A

Ordering of Two Streams to Maximize Sum Capacity

We wish to show that in a two-receiver scenario with precoding, the sum capacity is maximized by choosing the first receiver to be the one with the best channel, i.e., the largest $\|\mathbf{h}_k\|^2$.

Recall that before power control, the receiver that is ordered first gets its full single-user SNR and that the product of SNRs to the two receivers is independent of the ordering. Therefore, we can show the above by proving that, given a constant product of SNRs, the sum capacity is monotonically increasing with the maximum value of the two SNRs.

Suppose without loss of generality that with a particular ordering, the SNRs before power control are β_1 and β_2 , where $\beta_1 \geq \beta_2$. Power control gives a fraction α of power to receiver 1. The sum capacity, given perfect information embedding, is then

$$C = \log_2(1 + \alpha\beta_1) + \log_2(1 + (1 - \alpha)\beta_2).$$

By taking the derivative, we see that this is maximized with the waterfilling solution

$$\alpha = \max\left(1, \frac{\beta_1 - \beta_2 + \beta_1\beta_2}{2\beta_1\beta_2}\right).$$

If $\alpha = 1$, that is, if waterfilling gives all the power to one receiver, then capacity is maximized by choosing β_1 as large as possible. Therefore, this case is proved.

From now on, then, assume that nonzero power is sent to both receivers. The sum

capacity is

$$\begin{aligned}
C &= \log_2 \left(1 + \frac{\beta_1 - \beta_2 + \beta_1\beta_2}{2\beta_1\beta_2} \beta_1 \right) + \log_2 \left(1 + \frac{\beta_2 - \beta_1 + \beta_1\beta_2}{2\beta_1\beta_2} \beta_2 \right) \\
&= \log_2 \left(1 + \frac{\beta_1^2 + \beta_2^2 - 2\beta_1\beta_2 + 2\beta_1^2\beta_2 + 2\beta_1\beta_2^2 + \beta_1^2\beta_2^2}{4\beta_1\beta_2} \right). \tag{A.1}
\end{aligned}$$

Next, let c be the product of SNRs, $c = \beta_1\beta_2$. The sum capacity in (A.1) becomes

$$C = \log_2 \left(1 + \frac{\beta_1^2 + \frac{c^2}{\beta_1^2} - 2c + 2\beta_1c + \frac{2c^2}{\beta_1} + c^2}{4c} \right).$$

We wish to show that for a constant c , this is monotonically increasing in β_1 . Since terms that are only functions of c will not affect this property, this is equivalent to showing that

$$c_1 = \beta_1^2 + \frac{c^2}{\beta_1^2} + 2\beta_1c + \frac{2c^2}{\beta_1}$$

is monotonic in β_1 . Taking derivatives,

$$\begin{aligned}
\frac{dc_1}{d\beta_1} &= 2\beta_1 - \frac{2c^2}{\beta_1^3} + 2c - \frac{2c^2}{\beta_1^2} \\
\frac{d^2c_1}{d\beta_1^2} &= 2 + \frac{6c^2}{\beta_1^4} + \frac{4c^2}{\beta_1^3}.
\end{aligned}$$

The first derivative is zero at $\beta_1 = \beta_2 = \sqrt{c}$ and the second derivative is always nonnegative. Since we assumed that $\beta_1 \geq \beta_2$, this implies that in the region of interest, c_1 is monotonically increasing in β_1 , and consequently, the sum capacity is monotonically increasing in β_1 . This proves the case when both streams are sent with nonzero power.

Appendix B

2-Bit Signaling in Larger-Order QAM Interference

Consider a 4-QAM embedding (that is, two bits of information) by receiver k with power constraint \mathcal{P}_k . We look at the case when the interfering signal behaves like a higher-order QAM constellation. For simplicity, assume that this interference constellation has infinite extent, and has equal-probability points spaced $2\zeta_I$ apart in both the real and imaginary directions.

If the interference is large enough, we can surround each constellation point with an embedding constellation, as in Fig. 3-11c, and suffer no precoding power loss. As we have seen, though, this only works if the spacing between interference points is large enough. When this is not true, we can match a larger number of interference points with each quartet of embedding points. Fig. B-1 demonstrates an embedding where each 4-QAM set surrounds four interference points. An interference point will get quantized to one member of the surrounding embedding quartet, selected by the input bit pair. If there are A^2 interference points for each quartet, then the embedding tiling will will not overlap as long as

$$A \geq \left\lceil \frac{\zeta}{\zeta_I} \right\rceil, \quad (\text{B.1})$$

where 2ζ is the spacing between embedding constellation points (of different types) and $\lceil \cdot \rceil$ is the ceiling operator.

For this type of embedding, the average transmitted power is the sum of the powers of a quartet of embedding points and of the set of interference points with




Figure B-1: Sample embedding of 4-QAM inside large-order QAM interference. In this example, each “embedding constellation” surrounds four possible interference points.

with it is matched, if we assume that both groups of points are centered at the origin. With a set of A^2 interference points,

$$\mathcal{P}_k = 2\zeta^2 + \frac{2}{3}\zeta_I^2(A^2 - 1).$$

The second term in the formula represents the precoding power loss, so it is clear that we want to surround as few interference points as possible, working at the lower bound of (B.1). The precoding power loss is shown in Fig. B-2. As the set of interference points becomes more dense, its discrete distribution gets closer to a uniform distribution, so it is not surprising that precoding power loss approaches that of uniformly-distributed interference.

A slightly different perspective, perhaps more in line with system goals, would be to maximize the distance 2ζ given a power constraint \mathcal{P}_k . To satisfy both the power constraint and (B.1), we may at times have to transmit with a power lower than \mathcal{P}_k .

Note that the tiling in Fig. B-1b once again looks like a uniform quantizer but where the embedding points were pulled back to the center of each set, just as in distortion compensation. However, in true distortion compensation, the reconstruction points would be pulled back in the direction of the interference points themselves, not to the center of each set of points.




Figure B-2: Precoding power loss for various relative interference ratios.

Bibliography

- [1] Siavash M. Alamouti. A simple transmit diversity technique for wireless communications. *IEEE JSAC*, 16(8):1451–1458, Oct. 1998.
- [2] ANSI T1E1.4/98-007R3. *Network and customer installation interfaces – Asymmetric digital subscriber line (ADSL) metallic interface*, 1998.
- [3] Sirikiat Lek Ariyavisitakul. Turbo space-time processing to improve wireless channel capacity. *IEEE Trans. Comm.*, 48(8):1347–1358, Aug. 2000.
- [4] Richard J. Barron, Brian Chen, and Gregory W. Wornell. The duality between information embedding and source coding with side information and some applications. *submitted to IEEE Trans. Info. Th.*, 2000.
- [5] Claude Berrou, Alain Glavieux, and Punya Thitimajshima. Near shannon limit error-correcting coding and decoding: turbo codes(1). In *Proc. ICC '93*.
- [6] E. Biglieri, J. Proakis, and S. Shamai. Fading channels: information-theoretic and communications aspects. *IEEE Trans. Info. Th.*, 44(6):2619–2692, Oct. 1998.
- [7] Giuseppe Caire and Shlomo Shamai. On achievable rates in a multi-antenna broadcast downlink. In *38th Annual Allerton Conference on Communications, Control and Computing*, 2000.
- [8] Giuseppe Caire and Shlomo Shamai. On the achievable throughput of a multi-antenna Gaussian broadcast channel. *submitted to IEEE Trans. Info. Th.*, 2001.
- [9] Yuk-Lun Chan and Weihua Zhuang. Modified TH-precoding with constant amplitude for indoor wireless system using QPSK. In *Proc. IEEE VTC '96*, 1996.

- [10] Brian Chen and Gregory W. Wornell. Quantization index modulation: a class of provably good methods for digital watermarking and information embedding. *IEEE Trans. Info. Th.*, 47(4):1423–1443, May 2001.
- [11] John M. Cioffi, Glen P. Dudevoir, M. Vedat Eyuboglu, and G. David Forney. MMSE decision-feedback equalizers and coding — part i: equalization results. *IEEE Trans. Comm.*, 43(10):2582–2594, Oct. 1995.
- [12] John M. Cioffi, Glen P. Dudevoir, M. Vedat Eyuboglu, and G. David Forney. MMSE decision-feedback equalizers and coding — part ii: coding results. *IEEE Trans. Comm.*, 43(10):2595–2604, Oct. 1995.
- [13] Max. H. M. Costa. Writing on dirty paper. *IEEE Trans. Info. Th.*, 29(3):439–441, May 1983.
- [14] Thomas A. Cover and Joy A. Thomas. *Elements of Information Theory*. John Wiley and Sons, 1991.
- [15] R. L. Cruz. Quality of service guarantees in virtual circuit switched networks. *IEEE JSAC*, 13(6):1048–1056, Aug. 1995.
- [16] A. Demers, S. Keshav, and S. Shenkar. Analysis and simulation of a fair queuing algorithm. *J. Internetworking Res. and Experience*, 1(1):3–26, Sept. 1990.
- [17] Alexandra Duel-Hallen. Decorrelating decision-feedback multiuser detector for synchronous code-division multi-access channel. *IEEE Trans. Comm.*, 41(2):285–290, Feb. 1993.
- [18] Alan Edelman. *Eigenvalues and Condition Numbers of Random Matrices*. PhD thesis, MIT, 1989.
- [19] Yonina C. Eldar and Albert M. Chan. The independence of Wishart matrix eigenvalues and eigenvectors and its application to the asymptotic performance of the decorrelator. *submitted to IEEE Trans. Info. Th.*, 2001.
- [20] Uri Erez, Shlomo Shamai, and Ram Zamir. Capacity and lattice-strategies for cancelling known interference. In *Proc. ISITA*, Nov. 2000.
- [21] ETSI TS 100 573 V8.4.0. *GSM digital cellular telecommunications systems; physical layer on the radio path; general description*, 1999.

- [22] G. David Forney, Les Brown, M. Vedat Eyuboglu, and John L. Moran II. The V.34 high-speed modem standard. *IEEE Communications Magazine*, pages 28–33, Dec. 1996.
- [23] G. David Forney and M. Vedat Eyuboglu. Combined equalization and coding using precoding. *IEEE Communications Magazine*, pages 25–34, Dec. 1991.
- [24] G. David Forney and M. Vedat Eyuboglu. Trellis precoding: combined coding, precoding and shaping for intersymbol interference channels. *IEEE Trans. Info. Th.*, 38(2):301–314, Mar. 1992.
- [25] Gerard J. Foschini. Layered space-time architecture for wireless communication in a fading environment when using multi-element antenna. *Bell Labs Technical Journal*, pages 41–59, Autumn 1996.
- [26] Gerard J. Foschini, Glen D. Golden, Reinaldo A. Valenzuela, and Peter W. Wolniansky. Simplified processing for high spectral efficiency wireless communication employing multi-element arrays. *IEEE JSAC*, 17(11):1841–1852, Nov. 1999.
- [27] G.J. Foschini and M.J. Gans. On limits of wireless communication in a fading environment when using multiple antennas. *Wireless Personal Communications*, 6(3):311–335, Mar. 1998.
- [28] Stuart Geman. A limit theorem for the norm of random matrices. *The Annals of Probability*, 8(2):252–261, April 1980.
- [29] Derek Gerlach and Arogyaswami Paulraj. Adaptive transmitting antenna arrays with feedback. *IEEE Signal Processing Letters*, 1(10):150–152, Oct. 1994.
- [30] George Ginis and John M. Cioffi. A multi-user precoding scheme achieving crosstalk cancellation with application to DSL systems. In *Conf. Record of the Thirty-Fourth Asilomar Conference on Signals, Systems, and Computers*, pages 1627–1631, 2000.
- [31] Andrea J. Goldsmith and Pravin P. Varaiya. Capacity of fading channels with channel side information. *IEEE Trans. Info. Th.*, 43(6):1986–1992, Nov. 1997.
- [32] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, 3rd edition, 1996.

- [33] D. J. Goodman, J. Borras, R. D. Mandayam, and R. D. Yates. Infostations: a new system model for data and messaging services. In *Proc. IEEE VTC '97*, 1997.
- [34] Jiann-Ching Guey, Michael P. Fitz, Mark R. Bell, and Wen-Yi Juo. Signal design for transmitter diversity wireless communication systems over Rayleigh fading channels. *IEEE Trans. Comm.*, 47(4):527–537, April 1999.
- [35] Joachim Hagenauer. Rate-compatible punctured convolutional codes (RCPC codes) and their applications. *IEEE Trans. Comm.*, 36(4):389–400, Apr. 1988.
- [36] H. Harashima and H. Miyakawa. Matched-transmission technique for channels with intersymbols interference. *IEEE Trans. Comm.*, COM-20(4):774–780, Aug. 1972.
- [37] Frederick S. Hillier and Gerald J. Lieberman. *Introduction to Operations Research*, pages 482–485. McGraw-Hill, New York, 6th edition, 1995.
- [38] Roger A. Horn and Charles R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, 1991.
- [39] William C. Jakes. *Microwave Mobile Communications*. IEEE Press, 1993.
- [40] R. Knopp and P. A. Humblet. Information capacity and power control in single-cell multiuser communications. In *Proc. ICC '95*, pages 331–335, 1995.
- [41] A. Lapidoth and S. M. Moser. On the fading number of multi-antenna systems. In *Proc. IEEE Info. Th. Workshop*, 2001.
- [42] Rajiv Laroia. Coding for intersymbol interference channels — combined coding and precoding. *IEEE Trans. Info. Th.*, 42(4):1053–1061, July 1996.
- [43] Rajiv Laroia, Nariman Farvardin, and Steven A. Tretter. On optimal shaping of multidimensional constellations. *IEEE Tran. Info. Th.*, 40(4):1044–1056, July 1994.
- [44] Edward A. Lee and David G. Messerschmitt. *Digital Communications*. Kluwer Academic Publishers, second edition, 1994.

- [45] Thomas L. Marzetta and Bertrand M. Hochwald. Capacity of a mobile multiple-antenna communication link in Rayleigh flat fading. *IEEE Tran. Info. Th.*, 45(1):139–157, Jan. 1999.
- [46] R. J. McEliece and W. E. Stark. Channels with block interference. *IEEE Trans. Info. Th.*, 30(1):44–53, Jan. 1984.
- [47] H. Miyakawa and H. Harashima. A method of code conversion for a digital communication channel with intersymbol interference. *Trans. Inst. Electron. Commun. Eng. Japan*, 52-A:272–273, June 1969.
- [48] Robb J. Muirhead. *Aspects of Multivariate Statistical Theory*. John Wiley and Sons, 1982.
- [49] Aradhana Narula, Michael J. Lopez, Mitchell D. Trott, and Gregory W. Wornell. Efficient use of side information in multiple-antenna data transmission over fading channels. *IEEE JSAC*, 16(8):1423–1436, Oct. 1998.
- [50] Aradhana Narula, Mitchell D. Trott, and Gregory W. Wornell. Performance limits of coded diversity methods for transmitter antenna arrays. *IEEE Trans. Info. Th.*, 45(7):2418–2433, Nov. 1999.
- [51] Garret T. Okamoto. *Smart Antenna Systems and Wireless LANS*. Kluwer Academic Publishers, Boston, 1999.
- [52] Lawrence H. Ozarow, Shlomo Shamai, and Aaron Wyner. Information theoretic considerations for cellular mobile radio. *IEEE Trans. Vehicular Technology*, 43(2):359–378, May 1994.
- [53] 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. Networking*, 1(3):344–357, June 1993.
- [54] Theodore S. Rappaport. *Wireless Communications: Principles and Practices*. Prentice-Hall, 1996.
- [55] F. Rashid-Farrokhi, K.J.R. Liu, and L. Tassiulas. Transmit beamforming for cellular communications systems. In *Proc. 31st Annual Conference on Information Sciences and Systems*, pages 92–97, Baltimore, MD, March 1997.

- [56] A. Wayne Robert and Dale E. Varberg. *Convex functions*. Academic Press, New York, 1973.
- [57] Faisal Shad, Terence D. Todd, Vytas Kezys, and John Litva. Dynamic slot allocation (DSA) in indoor sdma/tdma using a smart antenna basestation. *IEEE/ACM Trans. Networking*, 9(1):69–81, Feb. 2001.
- [58] Nadav Shulman and Meir Feder. Static broadcasting. In *Proc. IEEE ISIT*, page 23, 2000.
- [59] G. W. Stewart and Ji-guang Sun. *Matrix Perturbation Theory*. Academic Press, San Diego, CA, 1990.
- [60] Ion Stroica, Hui Zhang, and T. S. Ng. A hierarchical fair service curve algorithm for link-sharing, real-time, and priority services. *IEEE/ACM Trans. Networking*, 8(2):185–199, April 2000.
- [61] V. Tarokh, N. Seshadri, and A.R. Calderbank. Space-time codes for high data rate wireless communication I: Performance criterion and code construction. *IEEE Trans. Info. Th.*, 44(2):744–765, March 1998.
- [62] Vahid Tarokh, Hamid Jafarkhani, and A.R. Calderbank. Space-time block codes from orthogonal designs. *IEEE Trans. Info. Th.*, 1999.
- [63] I. Emre Telatar. Capacity of multi-antenna Gaussian channels. *AT&T-Bell Labs Technical Memo.*, June 1995.
- [64] M. Tomlinson. New automatic equalizer employing modulo arithmetic. *Electronics Letters*, pages 138–139, March 1971.
- [65] David Tse and Ofer Zeitouni. Linear multiuser receivers in random environments. *IEEE Trans. Info. Th.*, 2000.
- [66] David N. Tse. Optimal power allocation over parallel Gaussian broadcast channels. *submitted to IEEE Trans. Information Theory*, 1998. Can be found at <http://degas.eecs.berkeley.edu/~dtse/pub.html>.
- [67] David N. Tse and Stephen V. Hanly. Multi-access fading channels: Part i: Polymatroid structure, optimal resource allocation and throughput capacities. *IEEE Trans. Information Theory*, 44(7):2796–2815, Nov. 1998.

- [68] David N.C Tse and Stephen V. Hanly. Linear multiuser receivers: effective interference, effective bandwidth and user capacity. *IEEE Trans. Info. Th.*, 45(3):641–657, March 1998.
- [69] TUA/EUA IS-856. *CDMA 2000: High rate packet data air interface specification*, Nov. 2000.
- [70] P.P. Vaidyanathan. *Multirate systems and filter banks*, chapter 13-14. Prentice Hall, Englewood Cliffs, 1993.
- [71] Sriram Vishwanath, Nihar Jindal, and Andrea Goldsmith. On the capacity of multiple input multiple output broadcast channels. In *Proc. ICC 2002*, 2002.
- [72] E. Visotsky and U. Madhow. Space-time transmit strategies and channel feedback generation for wireless fading channels. In *Conf. Record of the Thirty-Fourth Asilomar Conference on Signals, Systems, and Computers*, pages 1593–1597, 2000.
- [73] Eugene Visotsky and Upamanyu Madhow. Optimum beamforming using transmit antenna arrays. In *Proc. IEEE VTC '99*, pages 851–856, 1999.
- [74] P. Viswanath, D. Tse, and R. Laroia. Opportunistic beamforming using dumb antennas. *IEEE Trans. Info. Th.*, 48(6):1277–1294, June 2002.
- [75] Pramod Viswanath and David N. Tse. Sum capacity of the multiple antenna broadcast channel. In *Proc. ISIT*, 2002.
- [76] Pramod Viswanath, David N.C. Tse, and Venkat Anatharam. Asymptotically optimal waterfilling in vector multiple access channels. *IEEE Trans. Info. Th.*, 47(1):241–267, Jan. 2001.
- [77] Richard D. Wesel and John M. Cioffi. Achievable rates for Tomlinson-Harashima precoding. *IEEE Trans. Info. Th.*, 44(2):824–831, Mar. 1998.
- [78] Jack H. Winters. The diversity gain of transmit diversity in wireless systems with rayleigh fading. In *Proc. ICC '94*, 1994.
- [79] Jack H. Winters, Jack Salz, and Richard D. Gitlin. The impact of antenna diversity on the capacity of wireless communication systems. *IEEE Trans. Communications*, 42(4):1740–1751, April 1994.

- [80] Gregory W. Wornell. Spread-response precoding for communication over fading channels. *IEEE Trans. Info. Th.*, 42(2):488–501, March 1996.
- [81] Weidong Yang and Guanghan Xu. Optimal downlink power assignment for smart antenna systems. In *Proc. IEEE ICASSP*, pages 3337–3340, May 1998.
- [82] Wei Yu and John Cioffi. Sum capacity of gaussian vector broadcast channels. *submitted to IEEE Trans. Info. Th.*, 2001.
- [83] Wei Yu and John M. Cioffi. Trellis precoding for the broadcast channel. In *GLOBECOM '01*, 2001.
- [84] Ram Zamir, Shlomo Shamai, and Uri Erez. Nested linear/lattice codes for structures multiterminal binning. *IEEE Trans. Info. Th.*, 48(6):1250–1276, June 2002.
- [85] Kabiz C. Zangi and Leonid G. Kransy. Maximizing data rate over m -input/single-output channels. *submitted to IEEE Trans. Wireless Comm.*, 2001.
- [86] I. Zheng and D. Tse. Communicating on the Grassmann manifold: a geometric approach to the non-coherent multiple antenna channel. *IEEE Trans. Info. Th.*, 48(2):359–383, Feb. 2002.
- [87] Lizhong Zheng and David N. Tse. Diversity and multiplexing: a fundamental tradeoff in multiple antenna channels. *to appear in IEEE Trans. Info. Th.*, 2002.