Fundamental Capacity of MIMO Channels

Andrea Goldsmith, Syed Ali Jafar, Nihar Jindal, and Sriram Vishwanath
Department of Electrical Engineering, Stanford University, Stanford, CA 94305
E-mail : andrea, syed, njindal, sriram@wsl.stanford.edu

November 8, 2002

Keywords

MIMO channels, Shannon Capacity, multiuser systems, multicell systems, antenna correlation, channel side information, beamforming, multiple access channels, broadcast channels.

Abstract

We provide an overview of the extensive recent results on the Shannon capacity of single-user and multiuser multiple-input multiple-output (MIMO) channels. Although enormous capacity gains have been predicted for such channels, these predictions are based on somewhat unrealistic assumptions about the underlying time-varying channel model and how well it can be tracked at the receiver as well as at the transmitter. More realistic assumptions can dramatically impact the potential capacity gains of MIMO techniques. For time-varying MIMO channels there are multiple Shannon-theoretic capacity definitions and, for each definition, different correlation models and channel side information assumptions that we consider. We first provide a comprehensive summary of ergodic and outage capacity results for single-user MIMO channels. These results indicate that the capacity gain obtained from multiple antennas heavily depends on the amount of channel knowledge at either the receiver or transmitter, the channel SNR, and the correlation between the channel gains on each antenna element. We then focus attention on the capacity regions for MIMO broadcast and multiple access channels. In contrast to single-user MIMO channels, capacity results for these multiuser MIMO channels are quite difficult to obtain, even for constant channels. We summarize capacity results for the MIMO broadcast and multiple access channel for channels that are either constant or fading with perfect instantaneous knowledge of the antenna gains at both transmitter(s) and receiver(s). We also show that the MIMO multiple access and broadcast capacity regions are intimately related via a duality transformation. This transformation is not only useful for proving capacity theorems; it also facilitates finding the optimal transmission strategy of the nonconvex MIMO broadcast channel using convex optimization techniques applied to the dual MIMO multiple access channel. Last we discuss capacity results for multicell MIMO channels with base station cooperation. The base stations then act as a spatially diverse antenna array, and transmission strategies that exploit this structure exhibit significant capacity gains. This section also provides a brief discussion of system level issues associated with MIMO cellular. Open problems in this field abound and are discussed throughout the paper.

I. INTRODUCTION

Wireless systems continue to strive for ever higher data rates. This goal is particularly challenging for systems that are power, bandwidth, and complexity limited. However, another domain can be exploited to significantly increase channel capacity: the use of multiple transmit and receive antennas. Pioneering work by Foschini [22] and Telatar [70] ignited much interest in this area by predicting remarkable spectral efficiencies for wireless systems with multiple antennas when the channel exhibits rich scattering and its variations can be accurately tracked. This initial promise of exceptional spectral efficiency almost "for free" resulted in an explosion of research activity to characterize the theoretical and practical issues associated with MIMO wireless channels and to extend these concepts to multiuser systems. This tutorial summarizes the segment of this recent work focused on the capacity of MIMO systems for both single users and multiple users under different assumptions about channel correlation and the amount of channel side information available at the transmitter and receiver.

The large spectral efficiencies associated with MIMO channels are based on the premise that a rich scattering environment provides independent transmission paths from each transmit antenna to each receive antenna. Therefore, for single-user systems, a transmission and reception strategy that exploits this structure achieves capacity on approximately $\min(N, M)$ separate channels, where N is the number of transmit antennas and M is the number of receive antennas. Thus, capacity scales linearly with $\min(N, M)$ relative to a system with just one transmit and one receive antenna. This capacity increase requires a scattering environment such that the matrix of channel gains between transmit and receive antenna pairs has full rank and independent entries, and that perfect estimates of these gains are available at the receiver. Perfect estimates of these gains at both the transmitter and receiver provides an increase in the constant multiplier associated with the linear scaling. Much subsequent work has been aimed at characterizing MIMO channel capacity under more realistic assumptions about the underlying channel model and the channel estimates available at the transmitter and receiver. The main question from both a theoretical and practical standpoint is whether the enormous capacity gains initially predicted by Foschini and Telatar can be obtained in more realistic operating scenarios, and what specific gains result from adding more antennas and/or a feedback link between the receiver and transmitter.

MIMO channel capacity depends heavily on the statistical properties and antenna element correlations of the channel. Recent work has developed both analytical and measurement-based MIMO channel models along with the corresponding capacity calculations for typical indoor and outdoor environments [1]. Antenna correlation varies drastically relative to the scattering environment, the distance between transmitter and receiver, the antenna configurations, and the Doppler spread [2, 67]. As we shall see, the effect of channel correlation on capacity depends on what is known about the channel at the transmitter and receiver: correlation sometimes increases capacity and sometimes reduces it [17]. Moreover, channels with very low correlation between antennas can still exhibit a “keyhole” effect where the rank of the channel gain matrix is very small, leading to limited capacity gains [13]. Fortunately, this effect is not prevalent in most environments.

We focus on MIMO channel capacity in the Shannon theoretic sense. The Shannon capacity (maximum mutual information) of a single-user time-invariant channel corresponds to the maximum data rate that can be transmitted over the channel with arbitrarily small error probability. When the channel is time-varying channel capacity has multiple definitions, depending on what is known about the instantaneous channel state information (CSI) at the transmitter and/or receiver and whether or not capacity is measured based on averaging the rate over all channel states or maintaining a fixed rate for most channel states. Specifically, when CSI is known perfectly at both transmitter and receiver, the transmitter can adapt its transmission strategy relative to the channel and therefore channel capacity is characterized by the ergodic, outage, or minimum rate capacity. Ergodic capacity defines the maximum average rate under an adaptive transmission strategy averaged over all channel states (long-term average). Outage capacity defines the maximum rate that can be maintained in *all* channel states with some probability of outage (no data transmission). Minimum rate capacity defines the maximum average rate under an adaptive transmission strategy that maintains a given minimum rate in every channel state, and then averages the total rate in excess of this minimum over all channel states. When only the receiver has perfect CSI then the transmitter must maintain a fixed-rate transmission strategy based on knowledge of the channel statistics only, which can include the full channel distribution, just its mean and variance (equivalent to the full distribution for complex Gaussian channel gains), or just its mean or variance. In this case ergodic capacity defines the rate that can be achieved via this fixed-rate strategy based on receiver averaging over all channel states [70]. Alternatively, the transmitter can send at a rate that cannot be supported by all channel states: in these poor channel states the receiver declares an outage and the transmitted data is lost. In this scenario each transmission rate has an outage probability associated with it and capacity is measured relative to outage probability (capacity CDF) [22]. An excellent tutorial on fading channel capacity for single antenna channels can be found in [5]. For single-user MIMO channels with perfect transmitter and receiver CSI the ergodic and outage capacity calculations are straightforward since the capacity is known for every channel state. Thus, for single-user MIMO systems the tutorial will focus on capacity results assuming imperfect CSI at the transmitter and perfect or imperfect CSI at the receiver. Although there has been much recent progress in this area, many open problems remain.

In multiuser channels capacity becomes a K-dimensional region defining the set of all rate vectors $(R_1, \ldots, R_K)$ simultaneously achievable by all K users. The multiple capacity definitions for time-varying channels under different transmitter and receiver CSI assumptions extend to the capacity region of the multiple access channel (MAC) and broadcast channel (BC) in the obvious way [30, 49, 50, 71] However, these MIMO multiuser capacity regions, even for time-invariant channels, are difficult to find. Few capacity results exist for time-varying multiuser MIMO channels, especially under realistic assumptions of partial CSI at the transmitter(s) and or receiver(s). Therefore, the tutorial focus for MIMO multiuser systems will be on ergodic capacity under perfect CSI at the transmitter and receiver, with a brief discussion of the known results and open problems for other capacity definitions and CSI assumptions.

Note that the MIMO techniques described herein are applicable to any channel described by a matrix. Matrix channels describe not only multiantenna systems but also channels with crosstalk [80] and wideband channels [72]. While the focus of this tutorial is on memoryless channels (flat-fading), the results can also be extended to channels with memory (ISI) using well-known methods for incorporating the channel delay spread into the channel matrix [61], as will be discussed in the next section.

Many practical MIMO techniques have been developed to capitalize on the theoretical capacity gains predicted by Shannon theory. A major focus of such work is space-time coding: recent work in this area is summarized in [23]. Other techniques for MIMO systems include space-time modulation [32, 35], adaptive modulation and coding [11], space-time equalization [3, 52], space-time signal processing [4], space-time CDMA [15, 36], and space-time OFDM [51, 53, 78]. An overview of the recent advances in these areas and other practical techniques along with their performance can be found in [27].

The remainder of this paper is organized as follows. In Section II we discuss the capacity of single-user MIMO systems under different assumptions about channel correlation and channel knowledge at the transmitter and receiver. This section also describes the optimality of beamforming under partial transmitter CSI and training issues. Section III describes the capacity regions for both MIMO BCs and MIMO MACs, along with a duality connection between these regions. The capacity-achieving strategy of dirty paper coding for the MIMO BC is also discussed. The capacity of multicell systems under dirty-paper coding and opportunistic beamforming is discussed in Section IV, as well as tradeoffs between capacity, diversity, and sectorization. Section V summarizes these capacity results and describes some remaining open problems and design questions associated with MIMO systems.

A table of abbreviations used throughout the paper is given below.

II. Single User MIMO

In this section we focus on the capacity of single-user MIMO channels. While most wireless systems today support multiple users, single-user results are still of much interest for the insight they provide and their application to channelized systems where users are allocated orthogonal resources (time, frequency bands, etc.) MIMO channel capacity is also much easier to derive for single users than for multiple users. Indeed, single-user MIMO capacity results are known for many cases where the corresponding multiuser problems remain unsolved. In particular, very little is known about multiuser capacity without

CSIR	Receiver Channel State Information
CSIT	Transmitter Channel State Information
CMI	Channel Mean Information
CCI	Channel Covariance Information
DPC	Dirty Paper Coding
MAC	Multiple-Access Channel
BC	Broadcast Channel

TABLE I

TABLE OF ABBREVIATIONS

the assumption of perfect channel state information at the transmitter (CSIT) and at the receiver (CSIR). While there remain many open problems in obtaining the single-user capacity under general assumptions of CSIT and CSIR, for several interesting cases the solution is known. This section will give an overview of known results for single-user MIMO channels with particular focus on partial CSI at the transmitter as well as the receiver. We begin with a brief description of the commonly used partial CSI models.

A. Channel Models

Consider a transmitter with M transmit antennas, and a receiver with N receive antennas. The channel can be represented by the $N \times M$ matrix $\mathbf{H}$. The $N \times 1$ received signal $\mathbf{y}$ is equal to

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n}, \tag{1}$$

where $\mathbf{x}$ is the $M \times 1$ transmitted vector and $\mathbf{n}$ is the $N \times 1$ additive white circularly symmetric complex Gaussian noise vector, normalized so that its covariance matrix is the identity matrix. The normalization of any non-singular noise covariance matrix $\mathbf{K_n}$ to fit the above model is as straightforward as multiplying the received vector $\mathbf{y}$ with $\mathbf{K_n}^{-1/2}$ to yield the effective channel $\mathbf{K_n}^{-1/2}\mathbf{H}$ and a white noise vector.

With perfect CSIT or CSIR, the channel matrix $\mathbf{H}$ is assumed to be known perfectly and instantaneously at the transmitter or receiver, respectively. With partial CSI at either the transmitter or receiver, the entries of the channel matrix $\mathbf{H}$ are often modeled as complex jointly Gaussian random variables. While this allows the channel to have both a non-zero mean as well as non-zero correlations between its elements, almost all research in this area has focused on the two extreme cases: channel mean information (CMI) and channel covariance information (CCI). Under CMI, the information available is the mean of the channel distribution while the covariance is modeled as white, i.e. the channel elements are assumed to be independent random variables. Under CCI, the channel is assumed to be varying too rapidly to track its mean, so that the mean is set to zero and the information regarding the relative geometry of the propagation paths is captured by a non-white covariance matrix. Mathematically, the two cases can be described as follows:

Channel Mean Information (CMI) : $\mathrm{E}[\mathbf{H}] = \overline{\mathbf{H}}, \;\; \mathrm{E}[(\mathbf{H_{ij}} - \overline{\mathbf{H}}_{\mathbf{ij}})(\mathbf{H}^{\star}_{\mathbf{kl}} - \overline{\mathbf{H}}^{\star}_{\mathbf{kl}})] = \alpha\delta_{ik}\delta_{jl}$

Channel Covariance Information (CCI) : $\mathbf{H} = (\mathbf{R^r})^{1/2}\mathbf{H^n}(\mathbf{R^t})^{1/2}$.

Here $\mathbf{H^n}$ is an $N \times M$ matrix of i.i.d. zero mean, unit variance complex circularly symmetric Gaussian random variables, and $\mathbf{R^r}$ and $\mathbf{R^t}$ are called the receive and transmit fade covariance matrices respectively. Although not completely general, this simple correlation model has been validated through recent field measurements as a sufficiently accurate representation of the fade correlations seen in actual cellular systems [14].

A common assumption for all the results on partial CSI is that the CSIT is a deterministic function of the CSIR. In other words, the CSI feedback link from the receiver to the transmitter, if it exists, is assumed to be noiseless. In general for partial information $\mathcal{U}$ at the transmitter, codes need to be defined over an extended alphabet of functions $\mathcal{U} \to \mathcal{X}$ where $\mathcal{X}$ is the input alphabet. However when the CSIT is a deterministic function of the CSIR optimal codes can be constructed directly over the input alphabet $\mathcal{X}$ [9].

Next, we summarize the single user MIMO capacity results under various assumptions on CSI.

B. Constant MIMO Channel Capacity

When the channel is constant and known perfectly at the transmitter and the receiver, Telatar [70] showed that the MIMO channel can be converted to parallel, non-interfering single input single output (SISO) channels through a singular value decomposition of the channel matrix. Waterfilling the transmit power over these parallel channels whose gains are given by the singular values σ_i^2 of the channel matrix leads to the power allocation

$$P_i = \left(\mu - \frac{1}{\sigma_i^2}\right)^+, \tag{2}$$

where μ is the waterfill level, P_i is the power in the i^{th} eigenmode of the channel, and x^+ is defined as $\max(x, 0)$. The channel capacity is shown to be

$$C = \sum_i (\log(\mu\sigma_i^2))^+. \tag{3}$$

Interestingly, the worst-case additive noise for this scenario is shown by Diggavi to be Gaussian in [20], where the worst case noise is also given explicitly for low SNR. An algorithm to compute the worst case noise for any SNR is obtained in [63].

Although the constant channel model is relatively easy to analyze, wireless channels in practice typically change over time due to multipath fading. The capacity of fading channels is investigated next.

C. Fading MIMO Channel Capacity

With slow fading, the channel may remain approximately constant long enough to allow reliable estimation at the receiver (perfect CSIR) and timely feedback of the channel state to the transmitter (perfect CSIT). However, in systems with moderate to high user mobility, the system designer is inevitably faced with channels that change rapidly. Fading models with partial CSIT or CSIR are more applicable to such channels. Capacity results under various assumptions regarding CSI are summarized in this section.

C.1 Capacity with Perfect CSIT and Perfect CSIR

Perfect CSIT and perfect CSIR model a fading channel that changes slow enough to be reliably measured by the receiver and fed back to the transmitter without significant delay. The ergodic capacity of a flat-fading channel with perfect CSIT and CSIR is simply the average of the capacities achieved with each channel realization. The capacity for each channel realization is given by the constant channel capacity expression in the previous section. Since each realization exhibits a capacity gain of $\min(M, N)$ due to the multiple antennas, the ergodic capacity will also exhibit this gain.

C.2 Capacity with No CSIT and Perfect CSIR

Seminal work by Foschini and Gans [24] and Telatar [70] addressed the case of no CSIT and perfect CSIR. The two relevant definitions in this case are capacity versus outage (capacity CDF) and ergodic capacity. For any given input covariance matrix the input distribution that achieves the ergodic capacity is shown in [24] and [70] to be complex vector Gaussian, mainly because the vector Gaussian distribution maximizes the entropy for a given covariance matrix. This leads to the transmitter optimization problem - i.e., finding the optimum input covariance matrix to maximize ergodic capacity subject to a transmit power (trace of the input covariance matrix) constraint. Mathematically, the problem is to characterize the optimum $\mathbf{Q}$ to maximize capacity

$$C = \max_{\mathbf{Q}\,:\,\text{trace}(\mathbf{Q})=P} C(\mathbf{Q}), \tag{4}$$

where

$$C(\mathbf{Q}) \triangleq \mathrm{E}\left[\log\left|\mathbf{I_{n_R}} + \mathbf{HQH}^\dagger\right|\right] \tag{5}$$

is the capacity with the input covariance matrix $\mathrm{E}[\overline{\mathbf{uu}}^\dagger] = \mathbf{Q}$ and the expectation is with respect to the channel matrix $\mathbf{H}$. The capacity $C(\mathbf{Q})$ is achieved by transmitting independent complex circular Gaussian symbols along the eigenvectors of $\mathbf{Q}$. The powers allocated to each eigenvector are given by the eigenvalues of $\mathbf{Q}$.

It is shown in [70] and [24] that the optimum input covariance matrix that maximizes ergodic capacity is the scaled identity matrix, i.e. the transmit power is divided equally among all the transmit antennas. So the ergodic capacity is the expected value of

$$C = \log|I_M + \frac{P}{M}\mathbf{HH}^\dagger| \tag{6}$$

Telatar [70] conjectures that the optimal input covariance matrix that maximizes capacity versus outage is a diagonal matrix with the power equally distributed among a *subset* of the transmit antennas. In fact, the higher the outage probability, the smaller the number of transmit antennas used. As the transmit power is shared equally between more antennas the expectation of C increases (so the ergodic capacity increases) but the tails of its distribution decay faster. While this improves capacity versus outage for low outage probabilities, the capacity versus outage for high outages is decreased. Usually we are interested in low outage probabilities[1] and therefore the usual intuition for outage capacity is that it increases as the diversity order of the channel increases, i.e. as the capacity CDF becomes steeper. Foschini and Gans [24] also propose a layered architecture to achieve these capacities with scalar codes. This architecture, called BLAST (Bell Labs Layered Space Time), shows enormous capacity gains over single antenna systems. e.g. for 1% outage, at 12 dB SNR and with 12 antennas, the spectral efficiency is shown to be 32 b/s/Hz as opposed to the spectral efficiencies of around 1 b/s/Hz achieved in present day single antenna systems. While the channel models in [24] and [70] assume uncorrelated and frequency flat fading, practical channels exhibit both correlated fading as well as frequency selectivity. The need to estimate the capacity gains of BLAST for practical systems in the presence of channel fade correlations and frequency selective fading sparked off measurement campaigns reported in [56] [26]. The measured capacities are found to be about 30% smaller than would be anticipated from an idealized model. However, the capacity gains over single antenna systems are still overwhelming.

[1] The capacity for high outage probabilities become relevant for schemes that transmit only to the best user. For such schemes, indeed it is shown in [7] that increasing the number of transmit antennas reduces the average sum capacity.

C.3 Capacity with Partial CSIT and Partial CSIR

There has been much interest in the capacity of multiple antenna systems with perfect CSIR but only partial CSIT. It has been found that unlike single antenna systems where exploiting CSIT does not significantly enhance the Shannon capacity [28], for multiple antenna systems the capacity improvement through even partial CSIT can be substantial. Key work on capacity of such systems by several authors including Madhow and Visotsky [74], Trott and Narula [59] [58], Jafar and Goldsmith [43] [41] [39], Jorsweick and Boche [46] [47] and Simon and Moustakas [57] [68] has provided many interesting results. In this subsection our aim is to summarize these results.

The optimum input covariance matrix in general can be a full rank matrix which implies vector coding across the antenna array. Limiting the rank of the input covariance matrix to unity, called *beamforming*, essentially leads to a scalar coded system which has a significantly lower complexity for typical array sizes.

The complexity versus capacity tradeoff in relation to the amount of CSIT is an interesting aspect of capacity results under partial CSIT. The ability to use scalar codes to achieve capacity with enough CSIT, also called optimality of beamforming, captures this tradeoff and has been the topic of much research in itself. Note that vector coding refers to fully unconstrained signaling schemes for the memoryless MIMO Gaussian channel. Every symbol period, a channel use corresponds to the transmission of a vector symbol comprised of the inputs to each transmit antenna. Ideally, while decoding vector codewords the receiver needs to take into account the dependencies in both space and time dimensions and therefore the complexity of vector decoding grows exponentially in the number of transmit antennas. A lower complexity implementation of the vector coding strategy is also possible in the form of several scalar codewords being transmitted in parallel. It is shown in [39] that without loss of capacity, any input covariance matrix, regardless of its rank, can be treated as several scalar codewords encoded independently at the transmitter and decoded successively at the receiver by subtracting out the contribution from previously decoded codewords at each stage. However, well known problems associated with successive decoding and interference subtraction, e.g. error propagation, render this approach unsuitable for use in practical systems. It is in this context that the question of optimality of beamforming becomes important. Beamforming transforms the MIMO channel into a single input single output (SISO) channel. Thus, well established scalar codec technology can be used to approach capacity and since there is only one beam, interference cancellation is not needed. In the summary given below we include the results on both the transmitter optimization problem as well as the optimality of beamforming. For simplicity we start with the multiple-input single-output (MISO) channel.

MISO Channels

We first consider systems that use a single receive antenna and multiple transmit antennas. The channel matrix is rank one. With perfect CSIT and CSIR, for every channel matrix realization it is possible to identify the only non-zero eigenmode of the channel accurately and beamform along that mode. On the other hand with no CSIT and perfect CSIR, it was shown by Foschini and Gans [24] and Telatar [70] that the optimal input covariance matrix is a multiple of the identity matrix. Thus, the inability of the transmitter to identify the non-zero channel eigenmode forces a strategy where the power is equally distributed in all directions.

For a system using a single receive antenna and multiple transmit antennas, the transmitter optimization problem is solved by Visotsky and Madhow in [74] for the cases of CMI (also called CMIT or mean feedback) at the transmitter and only CCI (also called CCIT or covariance feedback) at the transmitter. For the case of mean feedback ($\mathbf{H} \sim \tilde{\mathcal{N}}(\overline{\mathbf{H}}, \alpha\mathbf{I})$) the principal eigenvector of the optimal input covariance matrix $\mathbf{Q^o}$ is found to be along the channel mean vector and the eigenvalues corresponding to the remaining eigenvectors are shown to be equal. With covariance feedback ($\mathbf{H} \sim \tilde{\mathcal{N}}(\mathbf{0}, \mathbf{R^t})$) the eigenvectors of the optimal input covariance matrix $\mathbf{Q^o}$ are shown to be along the eigenvectors of the transmit fade covariance matrix and the eigenvalues are in the same order as the corresponding eigenvalues of the transmit fade covariance matrix. Moreover, Visotsky and Madhow's numerical results indicate that beamforming is close to the optimal strategy when the quality of feedback improves, i.e. when the channel uncertainty decreases under mean feedback or when a stronger channel mode can be identified under covariance feedback. For *mean feedback*, Narula and Trott [59] point out that there are cases where the capacity is actually achieved via beamforming. While they do not obtain fully general necessary and sufficient conditions for when beamforming is a capacity achieving strategy, they develop partial answers to the problem for two transmit antennas.

A general condition that is both necessary and sufficient for optimality of beamforming is obtained by Jafar and Goldsmith in [41] for both cases of mean and covariance feedback. The result can be stated as the following theorem:

Theorem The Shannon capacity can be achieved with a unit rank matrix if and only if the following

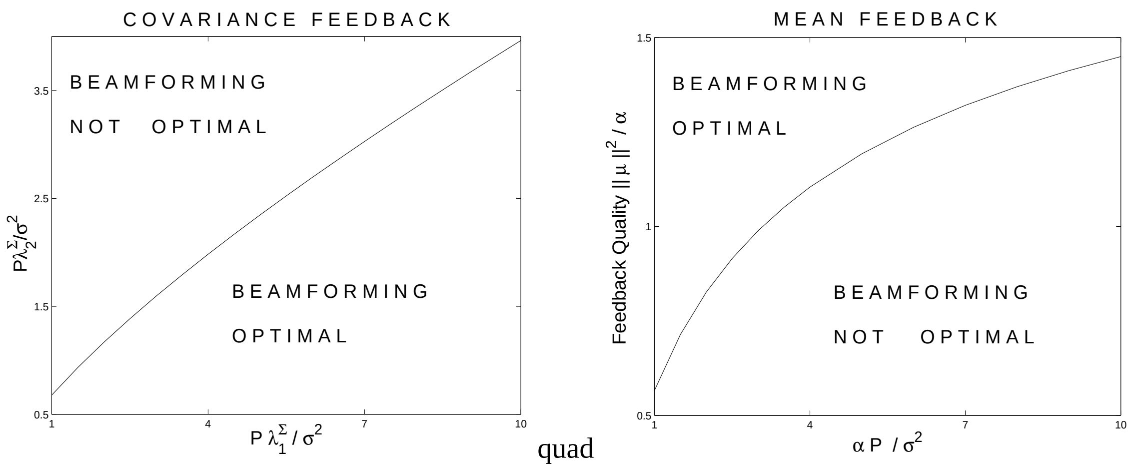

quad

Fig. 1. Plot of Necessary and Sufficient Condition (7)

condition is true:

$$\mathrm{E}\left[\frac{1}{1+P\lambda_1 w_1}\right] \leq \frac{1}{1+P\lambda_2} \tag{7}$$

where, for covariance feedback

1. $\lambda_1 > \lambda_2$ are the two largest eigenvalues of the channel fade covariance matrix $\mathbf{R^t}$,
2. w_1 is exponential distributed with unit mean, i.e. $w_1 \sim e^{-w1}$,

and for mean feedback

1. $\lambda_1 = \lambda_2 = \alpha$, and
2. w_1 has a noncentral chi-squared distribution. More precisely $w_1 \sim e^{-\frac{||\mu||^2}{\alpha} - w_1} I_0\left(2||\mu||\sqrt{\frac{w_1}{\alpha}}\right)$ where $I_0(.)$ is the 0th-order modified Bessel function of the first kind.

Further, for covariance feedback the expectation can be evaluated to express (7) explicitly in closed form as

$$\frac{1}{P\lambda_1} e^{\frac{1}{P\lambda_1}} \Gamma(0, \frac{1}{P\lambda_1}) \leq \frac{1}{1+P\lambda_2}. \tag{8}$$

The optimality conditions are plotted in Figure 1. For covariance feedback the optimality of beamforming depends on the two largest eigenvalues λ_1, λ_2 of the transmit fade covariance matrix and the transmit power P. Beamforming is found to be optimal when the two largest eigenvalues of the transmit covariance matrix are sufficiently disparate or the transmit power P is sufficiently low. Since beamforming corresponds to using the principal eigenmode alone, this is reminiscent of waterpouring solutions where only the deepest level gets all the water when it is sufficiently deeper than the next deepest level and when the quantity of water is small enough. For mean feedback the optimality of beamforming is found to depend on transmit power P and the *quality of feedback* defined as the ratio $\frac{||\overline{\mathbf{H}}||^2}{\alpha}$ of the norm squared of the channel mean vector $\overline{\mathbf{H}}$ and the channel uncertainty α. As the transmit power P is decreased or the quality of feedback improves beamforming becomes optimal. As mentioned earlier, for perfect CSIT (Quality of feedback $\rightarrow \infty$) the optimal input strategy is beamforming, while for zero CSIT (Quality of feedback $\rightarrow 0$) as shown by Telatar [70] the optimal input covariance has full rank, i.e. beamforming is necessarily sub-optimal. Note that [74], [59], [58] and [41] assume a single receive antenna. Next we summarize the corresponding results for MIMO capacity with partial CSIT and perfect CSIR.

MIMO Channels

With multiple transmit and receive antennas the covariance feedback case with correlations only at the transmitter ($\mathbf{R^r} = \mathbf{I}$) is solved by Jafar and Goldsmith in [43]. Like the single receive antenna case the capacity achieving input covariance matrix is found to have the eigenvectors of the transmit fade covariance matrix and the eigenvalues are in the same order as the corresponding eigenvalues of the transmit fade covariance matrix. Jafar and Goldsmith also presented in closed form a mathematical condition that is both necessary and sufficient for optimality of beamforming in this case. The same necessary and sufficient condition is also derived independently by Jorsweick and Boche in [46] and Simon and Moustakas in [68]. In [47] Jorsweick and Boche extend these results to incorporate fade correlations at the receiver as well. Their results show that while the receive fade correlation matrix does not affect the eigenvectors of the optimal input covariance matrix, it does affect the eigenvalues. The general condition for optimality of beamforming found by Jorsweick and Boche depends upon the two largest eigenvalues of the transmit covariance matrix and all the eigenvalues of the receive covariance matrix.

The mean feedback case with multiple transmit and receive antennas is solved by Jafar and Goldsmith in [39] and Moustakas and Simon in [68]. While the general case remains unsolved, solutions are obtained for the case where the channel mean has rank one [39] and two [68].

These results summarize our discussion of the case with partial CSIT and perfect CSIR. From these results we notice that the benefits of feeding back CSI from the receiver to the transmitter are two-fold. Not only does the capacity increase with feedback, but this feedback also allows the transmitter to identify the stronger channel modes and achieve this higher capacity with simple scalar codewords.

We conclude this subsection with a discussion on the growth of capacity with number of antennas. With perfect CSIR and i.i.d. fades, it was shown by Foschini and Gans [24] and by Telatar [70] that the channel capacity grows linearly with $\min(M, N)$. This linear increase occurs whether the transmitter knows the channel perfectly (perfect CSIT) or does not know the channel at all (no CSIT). The proportionality constant of this linear increase, called the rate of growth, has also been characterized in [33, 69, 70]. Chuah et. al. [16] proved that capacity increases linearly with min(M, N) even with correlated fading although the slope is reduced. As we will see in the next section, the assumption of perfect CSIR is crucial for the linear growth behavior of capacity with the number of antennas.

In the next section we explore the capacity when the channel is not known perfectly at either the transmitter or the receiver.

C.4 Capacity with No CSIT and No CSIR

We saw in the last section that with perfect CSIR, channel capacity grows linearly with the minimum of the number of transmit and receive antennas. However, reliable channel estimation may not be possible for a mobile receiver that experiences rapid fluctuations of the channel coefficients. Since user mobility is the principal driving force for wireless communication systems, the capacity behavior with partial CSIT and CSIR is of particular interest. In this section we summarize some MIMO capacity results in this area.

One of the first papers to address the MIMO capacity in the absence of CSIT and CSIR is [54] by Marzetta and Hochwald. They model the channel matrix components as i.i.d. complex Gaussian random variables that remain constant for a coherence interval of T symbol periods after which they change to another independent realization. Capacity is achieved when the $T \times M$ transmitted signal matrix is equal to the product of two statistically independent matrices: a $T \times T$ isotropically distributed unitary matrix times a certain $T \times M$ random matrix that is diagonal, real, and nonnegative. This result enables them to determine capacity for many interesting cases. Marzetta and Hochwald show that, for a fixed number of antennas, as the length of the coherence interval increases, the capacity approaches the capacity obtained as if the receiver knew the propagation coefficients. However, perhaps the most surprising result in [54] is the following: In contrast to the reported linear growth of capacity with $\min(M, N)$ under the perfect CSIR assumption, [54] showed that in the absence of CSI, capacity does not increase at all as the number of transmitter antennas is increased beyond the length of the coherence interval T. The MIMO capacity in the absence of CSIT and CSIR was further explored by Zheng and Tse [83]. They show that at high SNRs the optimal strategy is to use no more than $M^\star = \min(M, N, \lfloor T/2 \rfloor)$ transmit antennas. In particular, having more transmit antennas than receive antennas does not provide any capacity increase at high SNR.

Zheng and Tse also find that for each 3-db SNR increase, the capacity gain is $M^\star\left(1-\frac{M^\star}{T}\right)$.

Notice that [54], [83] assume block fading models, i.e. the channel fade coefficients are assumed to be constant for a block of T symbol durations. Hochwald and Marzetta extend their results to continuous fading in [55] where, within each independent T-symbol block, the fading coefficients have an arbitrary time correlation. If the correlation vanishes beyond some lag τ, called the *correlation time* of the fading, then it is shown in [55] that increasing the number of transmit antennas beyond $\min(\tau, T)$ antennas does not increase capacity. Lapidoth and Moser [48] explored the channel capacity in the absence of CSI further at high SNR without the block fading assumption. In contrast to the results of Zheng and Tse for block fading, Lapidoth and Moser show that without the block fading assumption, the channel capacity in the absence of CSI grows only double logarithmically in SNR. This result is shown to hold under very general conditions, even allowing for memory and partial receiver side information.

C.5 Capacity with Partial CSIT and Partial CSIR

The results in [54] and [83] seem to leave little hope of achieving the high capacity gains predicted for MIMO systems when users are highly mobile. However, before resigning ourselves to these less-than-optimistic results we note that these results assume no CSIT or CSIR. Even for a rapidly fluctuating channel where reliable channel estimation is not possible, it might be much easier to estimate the distribution of the channel fades instead of the channel realizations themselves. This is because the channel distribution changes much more slowly than the channel itself. The estimated distribution can be made available to the transmitter through a feedback channel. This brings us to the realm of MIMO capacity with partial CSI.

Recent work by Jafar and Goldsmith [38] addresses the MIMO channel capacity with CCIT and CCIR (channel covariance information at the receiver). The channel matrix components are modeled as spatially correlated complex Gaussian random variables that remain constant for a coherence interval of T symbol periods after which they change to another independent realization based on the spatial correlation model. The channel correlations are assumed to be known at the transmitter and receiver. As in the case of spatially white fading (no CSI), Jafar and Goldsmith show that with CCIT and CCIR the capacity is achieved when the $T \times M$ transmitted signal matrix is equal to the product of a $T \times T$ isotropically distributed unitary matrix, a certain statistically independent $T \times M$ random matrix that is diagonal, real, and nonnegative, and the matrix of the eigenvectors of the transmit fade covariance matrix $\mathbf{R^t}$. It is shown in [38] that the channel capacity is independent of the smallest $M - T$ eigenvalues of the transmit fade covariance matrix as well as the eigenvectors of the transmit and receive fade covariance matrices

$\mathbf{R^t}$ and $\mathbf{R^r}$. Also, in contrast to the results for the spatially white fading model where adding more transmit antennas beyond the coherence interval length ($M > T$) does not increase capacity, [38] shows that additional transmit antennas always increase capacity as long as their channel fading coefficients are spatially correlated. Thus, in contrast to the results in favor of independent fades with perfect CSIR, these results indicate that with CCI at the transmitter and the receiver, transmit fade correlations can be beneficial, making the case for minimizing the spacing between transmit antennas when dealing with highly mobile, fast fading channels that cannot be accurately measured. It is shown through a simple example that with coherence time $T = 1$ the capacity with M transmit antennas is $10 \log_{10} M$ db higher for perfectly correlated fades than for independent fades.

C.6 Frequency Selective Fading Channels

While flat fading is a realistic assumption for narrowband systems where the signal bandwidth is smaller than the channel coherence bandwidth, broadband communications involve channels that experience frequency selective fading. Research on the capacity of MIMO systems with frequency selective fading typically takes the approach of dividing the channel bandwidth into parallel flat fading channels, and constructing an overall block diagonal channel matrix with the diagonal blocks given by the channel matrices corresponding to each of these subchannels. Under perfect CSIR and CSIT, the total power constraint then leads to the usual closed-form waterfilling solution. Note that the waterfill is done simultaneously over both space and frequency. Even SISO frequency selective fading channels can be represented by the MIMO system model (1) in this manner [61]. For MIMO systems, the matrix channel model is derived by Bolcskei, Gesbert, and Paulraj in [6] based on an analysis of the capacity behavior of OFDM-based MIMO channels in broadband fading environments. Under the assumption of perfect CSIR and no CSIT, their results show that in the MIMO case, unlike the SISO case, frequency selective fading channels may provide advantages over flat fading channels not only in terms of outage capacity but also in terms of capacity vs outage. In other words, MIMO frequency selective fading channels are shown to provide both higher diversity gain and higher multiplexing gain than MIMO flat-fading channels. The measurements in [56] show that frequency selectivity makes the CDF of the capacity steeper and, thus, increases the capacity for a given outage as compared with the flat-frequency case, but the influence on the ergodic capacity is small.

C.7 Training for Multiple Antenna Systems

The results summarized in the previous sections indicate that CSI plays a crucial role in the capacity of MIMO systems. In particular, the capacity results in the absence of CSIR or with partial CSIR are strikingly different and often quite pessimistic compared to those that assume perfect CSIR. To recapitulate, with perfect CSIR MIMO channel capacity is known to increase linearly with min(M, N) even without CSIT and even with correlated fading. However in fast fading when the channel changes so rapidly that it cannot be estimated reliably (no CSIR) the capacity does not increase with the number of transmit antennas at all for $M > T$ where T is the channel decorrelation time. Also at high SNR, capacity with perfect CSIR increases logarithmically with SNR, while the capacity without CSIR increases only double logarithmically. Thus, CSIR is critical to exploiting the capacity benefits of multiple antenna wireless links. CSIR is often obtained by sending known training symbols to the receiver. However, with too little training the channel estimates are poor, whereas with too much training there is no time for data transmission before the channel changes. So the key question to ask is how much training is needed in multiple antenna wireless links. This question itself is the title of the paper [31] by Hassibi and Hochwald where they compute a lower-bound on the capacity of a channel that is learned by training and maximize the bound as a function of the receive SNR, fading coherence time, and number of transmitter antennas. When the training and data powers are allowed to vary, the optimal number of training symbols is shown to be equal to the number of transmit antennas - which is also the smallest training interval length that guarantees meaningful estimates of the channel matrix. When the training and data powers are instead required to be equal, the optimal training duration may be longer than the number of antennas. Hassibi and Hochwald also show that training-based schemes can be optimal at high SNR, but suboptimal at low SNR.

The results summarized in this section form the basis of our understanding of channel capacity under partial CSI. These results serve as useful indicators for the benefits of incorporating training and feedback schemes in a MIMO wireless link to obtain CSIR and CSIT respectively. However, our knowledge of MIMO capacity with partial CSI is still far from complete, even for single user systems. We conclude this section by pointing out some of the many open problems.

(1) Mean and covariance feedback: The transmitter optimization problem with combined mean and covariance feedback is unsolved even with a single receive antenna. With multiple transmit and receive antennas, the solution is not completely known except in special cases.

(2) CMI or CCI: Capacity is not known with CCI for completely general correlations, or with CMI for an

arbitrary channel mean matrix.

(3) Partial CSIR: Almost all cases with partial CSIR are open problems.

(4) Outage capacity: Most results under partial CSI are for ergodic capacity. Outage capacity has proven to be less analytically tractable than ergodic capacity and contains an abundance of open problems.

III. MULTI-USER MIMO

In this section we consider the two basic multi-user MIMO channel models: the MIMO multiple-access channel (MAC), and the MIMO broadcast channel (BC). Since the capacity region of a general MAC has been known for quite a while, there are quite a few results on the MIMO MAC for both constant channels and fading channels with different degrees of channel knowledge at the transmitters and receivers. The MIMO BC, however, is a relatively new problem for which capacity results have only recently been found. As a result, the field is much less developed, but we summarize the recent results in the area. Interestingly, the MIMO MAC and MIMO BC have been shown to be duals, as we will discuss in Section III-C.2.

A. System Model

To describe the MAC and BC models, we consider a cellular-type system in which the base station has M antennas, and each of the K mobiles has N antennas. The downlink of this system is a MIMO BC, and the uplink is a MIMO MAC. We will use $\mathbf{H}_i$ to denote the *downlink* channel matrix from the base station to User i. Thus, the *uplink* matrix of User i is $\mathbf{H}_i^\dagger$. A picture of the system model is shown in Fig. 2.

In the MAC, let $\mathbf{u_k} \in \mathbb{C}^{N\times 1}$ be the transmitted signal of transmitter k. Let $\mathbf{v} \in \mathbb{C}^{M\times 1}$ denote the received signal and $\mathbf{w} \in \mathbb{C}^{M\times 1}$ the noise vector where $\mathbf{w} \sim \mathcal{N}(0, \mathbf{I})$ is circularly symmetric Gaussian with identity covariance. The received signal is then equal to:

$$\begin{aligned} \mathbf{v} &= \mathbf{H_1}^\dagger \mathbf{u_1} + \ldots + \mathbf{H}_K^\dagger \mathbf{u}_K + \mathbf{w} \\ &= \mathbf{H}^\dagger \begin{bmatrix} \mathbf{u_1} \\ \vdots \\ \mathbf{u}_K \end{bmatrix} + \mathbf{w} \quad \text{where} \quad \mathbf{H}^\dagger = \begin{bmatrix} \mathbf{H_1}^\dagger & \ldots & \mathbf{H}_K^\dagger \end{bmatrix}. \end{aligned}$$

In the MAC, each transmitter is subject to an individual power constraint of $P_1, \ldots, P_K$. The transmit covariance matrix of User k is defined to be $\mathbf{Q}_k \triangleq E[\mathbf{u_k u_k}^\dagger]$. The power constraint implies $\mathrm{Tr}(\mathbf{Q}_k) \leq P_k$ for $k = 1, \ldots, K$.

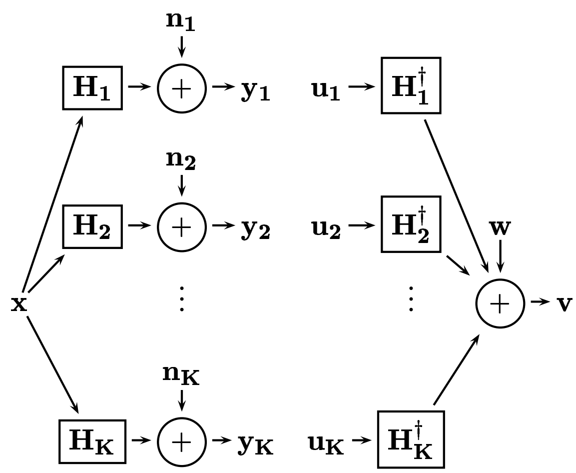

Fig. 2. System models of the MIMO BC(left) and the MIMO MAC (right) channels

In the BC, let $\mathbf{x} \in \mathbb{C}^{M\times 1}$ denote the transmitted vector signal and let $\mathbf{y_k} \in \mathbb{C}^{N\times 1}$ be the received signal at receiver k. The noise at receiver k is represented by $\mathbf{n_k} \in \mathbb{C}^{N\times 1}$ and is assumed to be circularly symmetric complex Gaussian noise ($\mathbf{n_k} \sim N(0, \mathbf{I})$). The received signal of User k is equal to:

$$\mathbf{y_k} = \mathbf{H_k x} + \mathbf{n_k}. \tag{9}$$

The transmit covariance matrix of the input signal is $\mathbf{\Sigma}_x \triangleq \mathbb{E}[\mathbf{x}\mathbf{x}^\dagger]$. The transmitter is subject to an average power constraint P, which implies $\mathrm{Tr}(\mathbf{\Sigma}_x) \leq P$.

A note on notation: we use boldface to denote matrices and vectors. $|\mathbf{S}|$ denotes the determinant and $\mathbf{S}^{-1}$ the inverse of a square matrix $\mathbf{S}$. For any general matrix $\mathbf{M}$, $\mathbf{M}^\dagger$ denotes the conjugate transpose and $\mathrm{Tr}(\mathbf{M})$ denotes the trace. $\mathbf{I}$ denotes the identity matrix and $\mathrm{diag}(\lambda_i)$ denotes a diagonal matrix with the (i, i) entry equal to λ_i.

B. MIMO MAC

In this section we summarize capacity results on the multiple-antenna MAC. We first consider the constant channel scenario, and then look at the fading channel. Since the capacity region of a general MAC is known, the expressions for the capacity of a constant MAC are quite straightforward. For the fading case, one must consider the CSI (channel knowledge) available at the transmitter and receiver. We consider four cases: perfect CSIR and perfect CSIT, perfect CSIR and partial CSIT, perfect CSIR and no CSIT, and finally the case where there is neither CSIT nor CSIR.

B.1 Constant Channel

The capacity of any MAC can be written as the convex closure of the union of pentagon regions for every product input distribution satisfying the user-by-user power constraints [19]. For the Gaussian MIMO MAC, however, it has been shown that it is sufficient to consider only Gaussian inputs and that the convex hull operation is not needed [81] [12]. For any set of powers $(P_1, \dots, P_K)$, the capacity of the MIMO MAC is:

$$\mathcal{C}_{\mathrm{MAC}}(P_1, \dots, P_K; \mathbf{H}^\dagger) \triangleq \bigcup_{\{\mathbf{Q}_i \geq 0,\ \mathrm{Tr}(\mathbf{Q}_i) \leq P_i \ \ \forall i\}} \left\{ \begin{array}{l} (R_1, \dots, R_K) : \\ \sum_{i \in S} R_i \leq \frac{1}{2} \log |\mathbf{I} + \sum_{i \in S} \mathbf{H}_i^\dagger \mathbf{Q}_i \mathbf{H}_i| \ \ \forall S \subseteq \{1, \dots, K\} \end{array} \right\}.$$

The i-th user transmits a zero-mean Gaussian with spatial covariance matrix $\mathbf{Q}_i$. Each set of covariance matrices $(\mathbf{Q}_1, \dots, \mathbf{Q}_K)$ corresponds to a K-dimension polyhedron (i.e. $\Big\{(R_1, \dots, R_K) : \sum_{i \in S} R_i \leq \frac{1}{2} \log |\mathbf{I} + \sum_{i \in S} \mathbf{H}_i^\dagger \mathbf{Q}_i \mathbf{H}_i| \ \ \forall S \subseteq \{1, \dots, K\}\Big\}$), and the capacity region is equal to the union (over all covariance matrices satisfying the trace constraints) of all such polyhedrons. The corner points of this pentagon can be achieved by *successive decoding*, in which users' signals are successively decoded and subtracted out of the received signal. For the two-user case, each set of covariance matrices corresponds to a pentagon, similar in form to the capacity region of the scalar Gaussian MAC. For example, the corner point where $R_1 = \log |\mathbf{I} + \mathbf{H}_1^\dagger \mathbf{Q}_1 \mathbf{H}_1|$ and $R_2 = \log |\mathbf{I} + \mathbf{H}_1^\dagger \mathbf{Q}_1 \mathbf{H}_1 + \mathbf{H}_2^\dagger \mathbf{Q}_2 \mathbf{H}_2| - R_1 = \log |\mathbf{I} + (\mathbf{I} + \mathbf{H}_1^\dagger \mathbf{Q}_1 \mathbf{H}_1)^{-1} \mathbf{H}_2^\dagger \mathbf{Q}_2 \mathbf{H}_2|$ corresponds to decoding User 2 first (i.e. in the presence of interference from User 1) and decoding User 1 last (without interference from User 2). Successive decoding can reduce a complex multi-user detection problem into a series of single-user detection steps [29].

The capacity region of a MIMO MAC for the single transmit antenna case ($N = 1$) is shown in Fig. 3. When $N = 1$, the covariance matrix of each transmitter is a scalar equal to the transmitted power. Clearly, each user should transmit at full power. Thus, the capacity region for a K-user MAC with $N = 1$ is given by the following region:

$$\left\{ (R_1, \dots, R_K) : \sum_{i \in S} R_i \leq \frac{1}{2} \log |\mathbf{I} + \sum_{i \in S} \mathbf{H}_i^\dagger P_i \mathbf{H}_i| \ \ \forall S \subseteq \{1, \dots, K\}. \right\} \quad (10)$$

For the two-user case, this reduces to the simple pentagon seen in Fig. 3.

When $N > 1$, however, a union must be taken over all covariance matrices. Intuitively, the set of covariance matrices that maximize R_1 are different from the set of covariance matrices that maximize the sum rate. In Figure 4 a MAC capacity region for $N > 1$ is shown. Notice that the region is equal to the union of pentagons, a few of which are shown with dashed lines in the figure. The boundary of the capacity region is in general curved, except at the sum rate point, where the boundary is a straight

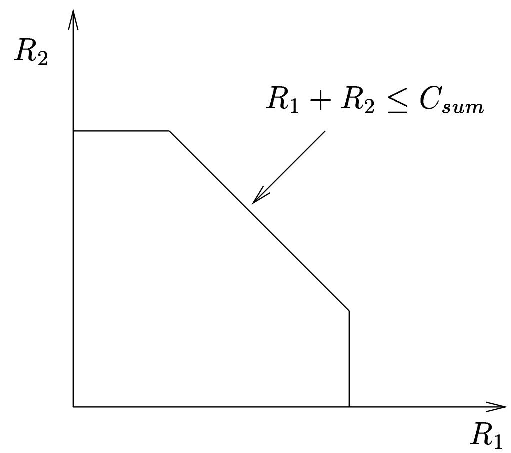

Fig. 3. Capacity region of MIMO MAC for $N = 1$

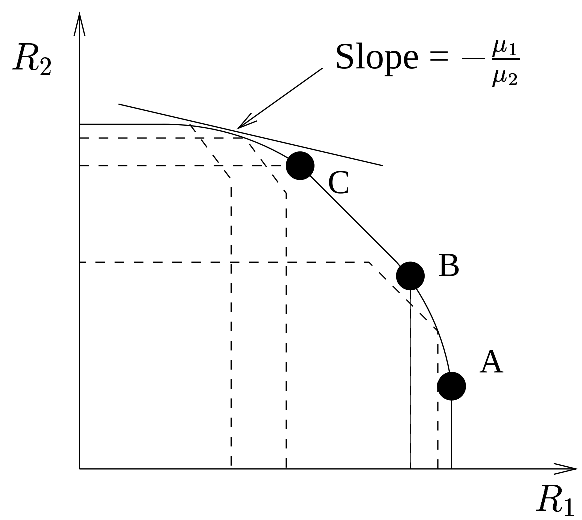

Fig. 4. Capacity region of MIMO MAC for $N > 1$

line [81]. Each point on the curved portion of the boundary is achieved by a *different* set of covariance matrices. At point A, User 1 is decoded last and achieves his single-user capacity by choosing $\mathbf{Q}_1$ as a water-fill of the channel $\mathbf{H}_1$. User 2 is decoded first, in the presence of interference from User 1, so $\mathbf{Q}_2$ is chosen as a waterfill of the channel $\mathbf{H}_2$ and the interference from User 1. At points B and C, sum-rate is maximized (by the same set of covariance matrices), but at B User 1 is decoded last whereas at C User 2 is decoded last.

Next, we focus on characterizing the optimal covariance matrices $(\mathbf{Q}_1, \ldots, \mathbf{Q}_K)$ that achieve any point on the boundary of the MIMO MAC capacity region. Since the MAC capacity region is convex, it is well known from convex theory that the boundary of the capacity region can be fully characterized by maximizing the function $\mu_1 R_1 + \cdots + \mu_K R_K$ over all rate vectors in the capacity region and for all nonnegative priorities $(\mu_1, \dots, \mu_K)$ such that $\sum_{i=1}^{K} \mu_i = 1$. For a fixed set of priorities, this is equivalent to maximizing the point on the capacity region boundary that is tangent to a line whose slope is defined by the priorities. See the tangent line in Fig. 4 for an example. The structure of the MAC capacity region implies that all boundary points of the capacity region are corner points of polyhedrons corresponding to different sets of covariance matrices. Furthermore, the corner point should correspond to decoding orders in order of *increasing* priority, i.e. the user with the largest priority should be decoded last and therefore sees no interference. Thus, the problem of finding that boundary point on the capacity region associated with priorities $\mu_1, \dots, \mu_K$ assumed to be in descending order (users can be arbitrarily re-numbered to satisfy this condition) can be written as

$$\max_{\mathbf{Q}_1, \cdots, \mathbf{Q}_K} \sum_{i=1}^{K-1} (\mu_i - \mu_{i+1}) \log |\mathbf{I} + \sum_{l=1}^{i} \mathbf{H}_l^{\dagger} \mathbf{Q}_l \mathbf{H}_l| + \mu_K \log |\mathbf{I} + \sum_{l=1}^{K} \mathbf{H}_l^{\dagger} \mathbf{Q}_l \mathbf{H}_l|$$

subject to power constraints on the traces of the covariance matrices. Note that the covariances that maximize the function above are the *optimal* covariances. The most interesting and useful feature of the optimization problem above is the concavity of the objective function in the covariance matrices. Thus, efficient convex optimization tools exist that solve this problem numerically [8]. A more efficient numerical technique to find the sum-rate maximizing (i.e. $\mu_1 = \cdots = \mu_K$) covariance matrices, called iterative waterfilling, was developed by Yu, et. al [81]. This technique is based on the Karush Kuhn Tucker (KKT) optimality conditions for the sum-rate maximization problem. These conditions indicate that the sum-rate maximizing covariance matrix of any user in the system should be the single-user water-filling covariance matrix of its own channel with noise equal to the actual noise plus the interference from the other $K - 1$ transmitters.

B.2 Fading Channels

Note that channel capacity for fading channels has many more parameters than for constant channels due to the time variation of the channel. As in the single-user case, the capacity of the MIMO MAC where the channel is time-varying depends on the definition of capacity and the availability of channel knowledge at the transmitters and the receiver. The capacity with perfect CSIT and CSIR is very well studied, as is the problem of capacity with perfect CSIR only. However, little is known about the capacity of the MIMO MAC (or the MIMO BC for that matter) with partial CSI at either the transmitter or receiver.

With perfect CSIR and perfect CSIT, the system can be viewed as a set of *parallel* non interfering MIMO MACs (one for each fading state) sharing a common power constraint. Thus, the ergodic capacity region can be obtained as an average of these parallel MIMO MAC capacity regions [60], where the averaging is done with respect to the channel statistics.The iterative waterfilling algorithm of [81] easily extends to this case, with joint space and time waterfilling.

The capacity region of a MAC with perfect CSIR but no CSIT was found in [25, 65]. In this case Gaussian inputs are optimal, and the ergodic capacity region is equal to the time average of the capacity obtained at each fading instant with a constant transmit policy (i.e. a constant covariance matrix for each user). Thus, the ergodic capacity region is given by

$$\bigcup_{\{\mathbf{Q}_i \geq 0,\ \mathrm{Tr}(\mathbf{Q}_i) \leq P_i \ \ \forall i\}} \left\{ (R_1, \ldots, R_K) : \sum_{i \in S} R_i \leq \frac{1}{2} \mathbb{E}_{\mathbf{H}} \left[\log |\mathbf{I} + \sum_{i \in S} \mathbf{H}_i^{\dagger} \mathbf{Q}_i \mathbf{H}_i| \right] \ \ \forall S \subseteq \{1, \ldots, K\} \right\}.$$

If the channel matrices $\mathbf{H}_i$ have i.i.d. complex Gaussian entries and each user has the same power constraint, then the optimal covariances are scaled versions of the identity matrix [70].

There has also been some work on capacity with perfect CSIR and partial CSIT [42]. In this paper, Jafar and Goldsmith determine the optimal covariance matrices when there is transmit antenna correlation and CCIT. This topic has yet to be fully investigated.

Asymptotic results on the sum capacity of MIMO MAC channels with the number of receive antennas and the number of transmitters increasing to infinity were obtained by Telatar [70] and by Viswanath, Tse, and Anantharam [77]. MIMO MAC sum capacity with only CSIR is found to grow *linearly* with $\min(M, NK)$ [70]. Sum capacity with perfect CSIR and CSIT also scales linearly with $\min(M, NK)$, but CSIT is of decreasing value as the number of receive antennas increases [34, 77]. Furthermore, the limiting distribution of the sum capacity with perfect CSIR and CSIT was found to be Gaussian by Hochwald and Vishwanath [34].

C. MIMO Broadcast Channel

In this section we summarize capacity results on the multiple-antenna broadcast channel (BC). When the transmitter has only one antenna, the Gaussian broadcast channel is a *degraded* broadcast channel, for which the capacity region is known [19]. However, when the transmitter has more than one antenna, the Gaussian broadcast channel is generally non-degraded. The capacity region of general non-degraded broadcast channels is unknown, but the seminal work of Caire and Shamai [10] and subsequent research on this problem have shed a great deal of light on this channel and the sum capacity of the MIMO BC has been found. In subsequent sections we focus mainly on the constant channel, but we do briefly discuss the fading channel as well which is still an open problem. Note that the BC transmitter has M antennas and each receiver has N antennas, as described in Section III-A.

C.1 Dirty Paper Coding Achievable Rate Region

An achievable region for the MIMO BC was first obtained for the $N = 1$ case by Caire and Shamai [10] and later extended to the multiple-receive antenna case by Yu and Cioffi [79] using the idea of *dirty paper coding* (DPC) [18]. The basic premise of DPC is as follows. If the transmitter (but not the receiver) has perfect, non-causal knowledge of additive Gaussian interference in the channel, then the capacity of the channel is the same as if there was no additive interference, or equivalently as if the receiver also had knowledge of the interference. Dirty-paper coding is a technique that allows non-causally known interference to be "pre-subtracted" at the transmitter, but in such a way that the transmit power is not increased. A more practical (and more general) technique to perform this pre-subtraction is the cancelling for known interference technique found by Erez, Shamai, and Zamir in [21].

In the MIMO BC, dirty-paper coding can be applied at the transmitter when choosing codewords for different receivers. The transmitter first picks a codeword for receiver 1. The transmitter then chooses a codeword for receiver 2 with full (non-causal) knowledge of the codeword intended for receiver 1. Therefore the codeword of User 1 can be pre-subtracted such that receiver 2 does not see the codeword intended for receiver 1 as interference. Similarly, the codeword for receiver 3 is chosen such that receiver 3 does not see the signals intended for receivers 1 and 2 as interference. This process continues for all K receivers. If User $\pi(1)$ is encoded first, followed by User $\pi(2)$, etc., the following is an achievable rate vector:

$$R_{\pi(i)} = \frac{1}{2} \log \frac{|\mathbf{I} + \mathbf{H}_{\pi(i)}(\sum_{j \geq i} \mathbf{\Sigma}_{\pi(j)})\mathbf{H}_{\pi(i)}^{\dagger}|}{|\mathbf{I} + \mathbf{H}_{\pi(i)}(\sum_{j > i} \mathbf{\Sigma}_{\pi(j)})\mathbf{H}_{\pi(i)}^{\dagger}|} \quad i = 1, \ldots, K. \tag{11}$$

The dirty-paper region $\mathcal{C}_{\text{DPC}}(P, \mathbf{H})$ is defined as the convex hull of the union of all such rates vectors over all positive semi-definite covariance matrices $\mathbf{\Sigma}_1, \ldots, \mathbf{\Sigma}_K$ such that $\text{Tr}(\mathbf{\Sigma}_1 + \ldots \mathbf{\Sigma}_K) = \text{Tr}(\mathbf{\Sigma}_x) \leq P$ and over all permutations $(\pi(1), \ldots, \pi(K))$:

$$\mathcal{C}_{\text{DPC}}(P, \mathbf{H}) \triangleq Co \left(\bigcup_{\pi, \mathbf{\Sigma}_i} \mathbf{R}(\pi, \mathbf{\Sigma}_i) \right) \tag{12}$$

where $\mathbf{R}(\pi, \mathbf{\Sigma}_i)$ is given by (11). The transmitted signal is $\mathbf{x} = \mathbf{x_1} + \ldots + \mathbf{x_K}$ and the input covariance matrices are of the form $\mathbf{\Sigma}_i = \mathbb{E}[\mathbf{x_i}\mathbf{x_i}^{\dagger}]$. From the dirty paper result we find that $\mathbf{x}_1, \ldots, \mathbf{x}_K$ are uncorrelated, which implies $\mathbf{\Sigma}_x = \mathbf{\Sigma}_1 + \ldots \mathbf{\Sigma}_K$.

One important feature to notice about the dirty paper rate equations in (11) is that the rate equations are neither a concave nor convex function of the covariance matrices. This makes finding the dirty paper region very difficult, because generally the entire space of covariance matrices that meet the power constraint must be searched over.

Note that dirty paper coding and successive decoding are completely equivalent capacity-wise for scalar channels, but this equivalence does not hold for MIMO channels. It has been shown [37] that the achievable region with successive decoding is contained within the dirty paper coding region.

C.2 MAC-BC Duality

In [73], Vishwanath, Jindal, and Goldsmith developed the following key result relating the MIMO BC dirty paper region and the MIMO MAC capacity region:

Theorem 1: The dirty paper rate region of the multi-antenna BC with power constraint P is equal to the union of capacity regions of the dual MAC, where the union is taken over all individual power constraints that sum to P:

$$\mathcal{C}_{\text{DPC}}(P, \mathbf{H}) = \bigcup_{\mathbf{P}: \sum_{i=1}^{K} P_i = P} \mathcal{C}_{MAC}(P_1, \ldots, P_K, \mathbf{H}^{\dagger}).$$

This is the multiple-antenna extension of the previously established duality between the scalar Gaussian broadcast and multiple-access channels [45]. In addition to the relationship between the two rate regions, for any set of covariance matrices in the MAC/BC (and the corresponding rate vector), [73] provides an explicit set of transformations to find covariance matrices in the BC/MAC that achieve the same rates. The union of MAC capacity regions is easily seen to be the same expression as in (10) but with the constraint $\sum_{i=1}^{K} \text{Tr}(\mathbf{Q}_i) \leq P$ instead of $\text{Tr}(\mathbf{Q}_i) \leq P_i \;\; \forall i$ (i.e. a sum constraint instead of individual constraints).

The MAC-BC duality is very useful from a numerical standpoint because the dirty paper region leads to non-concave rate functions of the covariances, whereas the rates in the dual MAC are *concave* functions of the covariance matrices. Thus, the optimal MAC covariances can be found using standard convex optimization techniques and then transformed to the corresponding optimal BC covariances using the MAC-BC transformations given in [73]. A specific algorithm that finds the sum rate optimal covariances is found in [44]. The algorithm is based on the iterative waterfilling algorithm for the MIMO MAC [81].

The dirty paper rate region is shown in Fig. 5 for a channel with 2 users, $M = 2$, and $N = 1$. The region is formed as a union of MAC regions. Since $N = 1$, each of the MAC regions is a pentagon, as discussed in Section III-B.1. Similar to the MAC capacity region, the boundary of the DPC region is curved, except at the sum-rate maximizing portion of the boundary. For the $N = 1$ case, duality also indicates that rank-one covariance matrices (i.e. beamforming) are optimal for dirty-paper coding. This

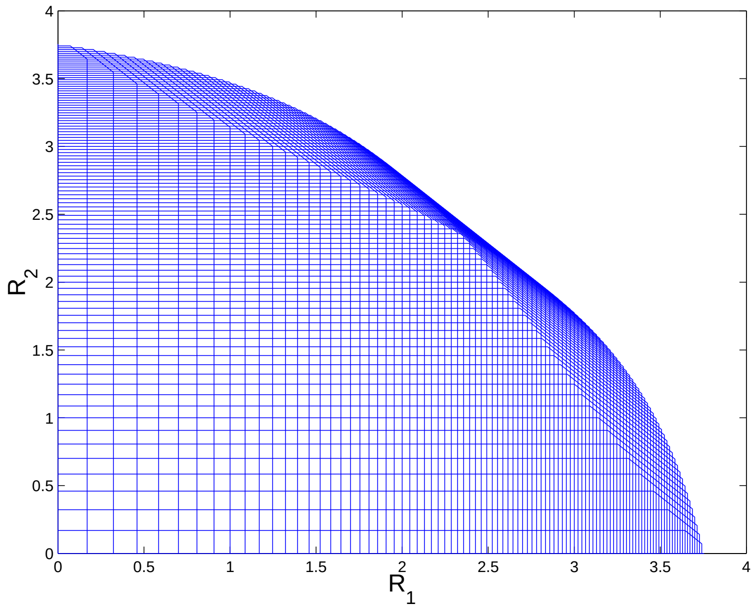

Fig. 5. Dirty paper rate region

fact is not obvious from the dirty paper rate equations, but follows from the transformations of [73] which find BC covariances that achieve the same rates as a set of MAC covariance matrices (which are scalars in the $N = 1$ case).

Duality also allows the MIMO MAC capacity region to be expressed in terms of the dual dirty paper BC rate region [73], as specified in the following theorem:

Theorem 2: The capacity region of a MIMO MAC is the intersection of the scaled dirty paper regions of the MIMO BC:

$$\mathcal{C}_{\text{MAC}}(P_1, \dots, P_K, \mathbf{H}^\dagger) = \bigcap_{\boldsymbol{\alpha} > 0} \mathcal{C}_{\text{DPC}}\Big(\sum_{i=1}^{K} \frac{P_i}{\alpha_i}; [\sqrt{\alpha_1}\mathbf{H}_1^T \cdots \sqrt{\alpha_K}\mathbf{H}_K^T]^T\Big). \tag{13}$$

C.3 Sum Rate Optimality of DPC

Dirty paper coding was first shown to achieve the sum-rate capacity of the MIMO BC for the 2-user, $M = 2$, $N = 1$ case by Caire and Shamai [10]. This was shown by using the Sato upper bound [64] on the broadcast channel, where the upper bound was shown to meet the achievable sum rate using DPC. The sum rate optimality of DPC was extended to the multi-user channel with $N = 1$ by Viswanath and Tse [76] and to the more general $N > 1$ case by Vishwanath, Jindal, and Goldsmith [73].

C.4 Fading Channels

Most of the capacity problems for fading MIMO BCs are still open, with the exception of sum rate capacity with perfect CSIR and perfect CSIT. In this case, as for the MIMO MAC, the MIMO BC can be split into parallel channels with an overall power constraint (see Li and Goldsmith [49] for a treatment of the scalar case). Again, the sum power waterfilling algorithm can be used in this case with combined time and space waterfilling.

Asymptotic results for the sum rate capacity of the MIMO BC (for $N = 1$) can also be obtained by combining the asymptotic results for the sum rate capacity of the MIMO MAC with duality [34]. Thus, the role of transmitter side information reduces with the growth in the number of transmit antennas, and hence, the sum capacity of the MIMO BC with K users and M transmit antennas tends to the sum capacity of a single user system with only receiver CSI and M receive antennas and K transmit antennas, which is given by

$$\log|\mathbf{I} + \frac{P}{K}\mathbf{H}\mathbf{H}^{\dagger}|.$$

Thus, the asymptotic growth with or without side information is linear as $\mathfrak{C}\min(M, K)$, and the growth rate constant $\mathfrak{C}$ can be found in [34].

D. Open Problems in Multiuser MIMO

Multiuser MIMO has been the primary focus of research in recent years, mainly due to the large number of open problems in this area. Some of these are:

1. No CSI: The capacity region of MIMO MAC and BC systems when neither the transmitter(s) nor the receiver(s) know the channel.
2. BC with receiver CSI alone: The Broadcast channel capacity is only known when both the transmitter and the receivers have perfect knowledge of the channel.
3. Imperfect CSI: Since perfect CSI is never possible, a study of capacity with imperfect CSI for both MAC and BC is of great practical relevance.
4. Non-DPC techniques for BC: Dirty-paper coding is a very powerful capacity-achieving scheme, but it appears quite difficult to implement in practice. Thus, non-DPC multi-user transmissions schemes for the downlink (such as downlink beamforming [62]) are also of practical relevance.

IV. MULTICELL MIMO

The MAC and the BC are information theoretic abstractions of the uplink and the downlink of a single cell in a cellular system. However a cellular system, by definition, consists of many cells. Due to the fundamental nature of wireless propagation, transmissions in a cell are not limited to within that cell. Users and base stations in adjacent cells experience interference from each other. Also, since the base stations are typically not mobile themselves there is the possibility for the base stations to communicate through a high speed reliable connection, possibly consisting of optical fiber links capable of very high data rates. This opens up the opportunity for base stations to cooperate in the way they process different users' signals. Analysis of the capacity of the cellular network, explicitly taking into account the presence of multiple cells, multiple users, and multiple antennas and the possibilities of cooperation between base stations is inevitably a hard problem and runs into several long-standing unsolved problems in network information theory. However, such an analysis is also of utmost importance because it defines a common benchmark that can be used to gauge the efficiency of any practical scheme, in the same way that the capacity of a single user link serves as a measure of the performance of practical schemes. There has been some recent research in this area that extend the single cell MAC and BC results to multiple cells. In this section we summarize some of these results.

The key to the extension of single cell results to multiple cell systems is the assumption of perfect cooperation between base stations. Conceptually, this allows the multiple base stations to be treated as physically distributed antennas of one composite base station. Specifically, consider a group of B coordinated cells, each with M antennas, and K mobiles, each with N antennas. If we define $\mathbf{H}_{i,j} \in \mathcal{C}^{N \times M}$ to be the *downlink* channel of User i from base station j, then the *composite* downlink channel of User i is $\mathbf{F}_i = [\mathbf{H}_{i,1} \cdots \mathbf{H}_{i,B}]$ and the composite uplink channel is $\mathbf{F}_i^\dagger$. The received signal of User i can then be written as $\mathbf{Y}_i = \mathbf{F}_i \mathbf{W} + \mathbf{n}_i$, where $\mathbf{W}$ is the composite transmitted signal defined as $\mathbf{W}^T = [\mathbf{W}_1^T \cdots \mathbf{W}_B^T]$. Here we let $\mathbf{W}_j$ represent the transmit signal from base j.

First let us consider the uplink. As pointed out by Jafar, Foschini and Goldsmith [37] the single cell MIMO MAC capacity region results apply to this system in a straightforward way. Thus, by assuming perfect data cooperation between the base stations, the multiple cell uplink is easily seen to be equal to the MAC capacity region of the composite channel, defined as $\mathcal{C}_{MAC}(\mathbf{F}_1^\dagger, \ldots, \mathbf{F}_K^\dagger; P_1, \ldots, P_K)$ in (10) where the power constraints of the i-th mobile is P_i.

On the downlink, since the base stations can cooperate perfectly, dirty-paper coding can be used over the entire transmitted signal (i.e. across base stations) in a straightforward manner. The application of dirty paper coding to a multiple cell environment with cooperation between base stations is pioneered in recent work by Shamai and Zaidel [66]. For one antenna at each user and each base station, they show that a relatively simple application of dirty paper coding can enhance the capacity of the cellular downlink. While capacity computations are not the focus of [66], they do show that their scheme is asymptotically optimal at high SNRs.

The MIMO downlink capacity is explored by Jafar and Goldsmith in [40]. Note that the multicell downlink can be solved in a similar way as the uplink. But this requires perfect data and *power cooperation* between the base stations. If we let $\mathbf{X}_{i,j}$ represent the transmit vector for User i from base station j, the composite transmit vector intended for User i is $\mathbf{X}_i^T = [\mathbf{X}_{i,1}^T \cdots \mathbf{X}_{i,B}^T]$. Thus, the composite covariance of User i is defined as $\Sigma_i = E[\mathbf{X}_i \mathbf{X}_i^T]$. The covariance matrix of the entire transmitted signal is $\Sigma_x = \sum_{i=1}^K \Sigma_i$. Assuming perfect data cooperation between the base stations, dirty-paper coding can be applied to the composite vectors intended for different users. Thus, the dirty paper region described in Section III-C.1, Equation (12), can be achieved in the multi-cell downlink.

While data cooperation is a justifiable assumption for capacity computations in the sense that it captures the possibility of base stations cooperating among themselves as described earlier in this section, in practice each base station has its own power constraint. The per base power constraint can be expressed as $E[\mathbf{W}_j^\dagger \mathbf{X}_j] \leq P_j$, where P_j is the power constraint at base j. Thus power cooperation, or pooling the transmit power for all the base stations to have one overall transmit power constraint is not as realistic. Note that on the uplink the base stations are only receiving signals, and therefore no power cooperation is required. The per base power constraints restrict consideration to covariance matrices such that $\sum_{i=1}^K \text{Tr}(E[\mathbf{X}_{i,j} \mathbf{X}_{i,j}^\dagger] \leq P_j \; j = 1, \ldots, B$. This is equivalent to a constraint of P_1 on the sum of the first M diagonal entries of Σ_x, a constraint of P_2 on the sum on the next M diagonal entries of Σ_x, etc. These constraints are considerably stricter than a constraint on the trace of Σ_x as in the single-cell case.

Though dirty paper coding yields an achievable region, DPC has not been shown to achieve the capacity region or even the sum rate capacity with per-base power constraints. Additionally, the MAC-BC duality (Section III-C.2) which greatly simplified calculation of the dirty paper region does not apply under per base power constraints. Thus, even generating numerical results for the multi-cell downlink is quite challenging.

However, data and power cooperation does give a simple upperbound on the capacity of the network. Based on numerical comparisons between this upperbound and a lower bound on capacity derived in [40], Jafar and Goldsmith find that the simple upper bound with power and data cooperation is also a

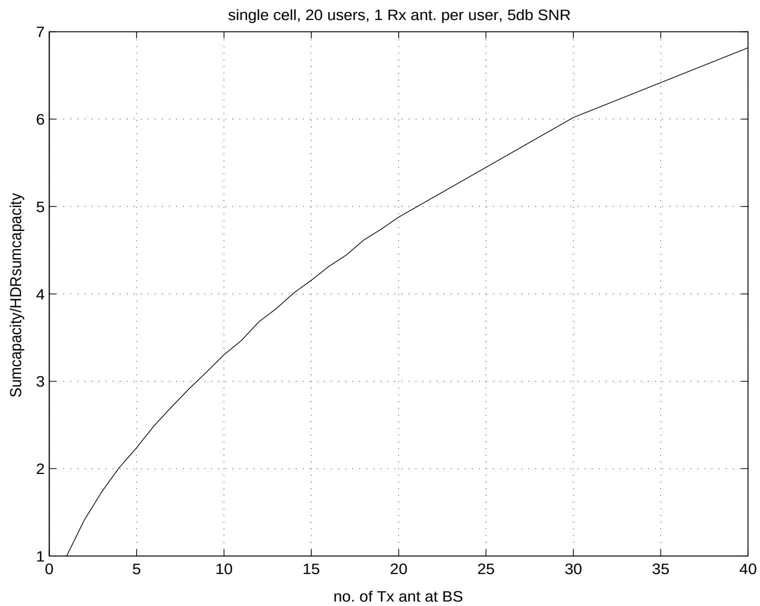

Fig. 6. Optimal Sum Rate relative to HDR

good measure of the capacity with data cooperation alone.

Note that current wireless systems use the high data rate (HDR) protocol and transmit to only one user at a time on the downlink. Dirty paper coding however allows the base station to transmit to many users simultaneously. This is particularly advantageous when the number of transmit antennas at the base station is much larger than the number of receive antennas at each user - a common scenario in current cellular systems. To illustrate the advantages of dirty paper coding, even for a single cell, the relative gains of optimal dirty paper coding over a strategy that serves only the best user at any time are shown in Figure 6. Note that this single cell model is equivalent to the multi-cell system with no cooperation between base stations so that the interference from other cells is treated as noise. With cooperation between base stations the gains are expected to be even more significant as dirty paper coding reduces the overall interference by making some users invisible to others.

The capacity results described in this section address just a few out of many interesting questions in the design of a cellular system with multiple antennas. Multiple antennas can be used not only to enhance the capacity of the system but also to drive down the probability of error through diversity combining. Recent work by Zheng and Tse [82] unravels a fundamental diversity versus multiplexing tradeoff. Also, instead of using isotropic transmit antennas on the downlink and transmitting to many users, it may be simpler to use directional antennas to divide the cell into sectors and transmit to one user within sector. The relative impact of partial CSI on each of these schemes is not fully understood. Although in this paper we focus on the physical layer, smart schemes to handle partial CSI can also be found at higher layers. An interesting example is the idea of opportunistic beamforming [75]. In the absence of CSIT, the transmitter randomly chooses the beamforming weights. With enough users in the system, it becomes very likely that these weights will be nearly optimal for one of the users. In other words, a random beam selected by the transmitter is very likely to be pointed towards a user if there are enough users in the system. Instead of feeding back the channel coefficients to the transmitter the users simply feedback the SNRs they see with the current choice of beamforming weights. This significantly reduces the amount of feedback required. By randomly changing the weights frequently, the scheme also treats all users fairly.

V. Conclusions

We have summarized recent results on the capacity of MIMO channels for both single-user and multiuser systems. The great capacity gains predicted for such systems can be realized in some cases, but realistic assumptions about channel knowledge and the underlying channel model can significantly mitigate these gains. For single user systems the capacity under perfect CSI at the transmitter and receiver is relatively straightforward and predicts that capacity grows linearly with the number of antennas. Backing off from the perfect CSI assumption makes the capacity calculation much more difficult, and the capacity gains are highly dependent on the nature of the CSIR/CSIT, the channel SNR, and the antenna element correlations. Specifically, assuming perfect CSIR, CSIT provides significant capacity gain at low SNRs but not much at high SNRs. The insight here is that at low SNRs it is important to put power into the appropriate eigenmodes of the system. In addition, channel correlation reduces capacity under perfect CSIR and no CSIT by a constant multiplier, and it increases capacity under imperfect CSIR and CSIT. This result indicates that channel correlation helps to compensate for capacity loss due to imperfect channel knowledge at the receiver. Finally, with no CSIR or CSIT at high SNRs capacity only grows relative to the double log of the SNR with the number of antennas as a constant additive term. This rather poor capacity gain would not typically justify adding more antennas. However, for block fading channels at moderate SNRs the growth without CSIT or CSIR is less pessimistic.

We also examined the capacity of MIMO broadcast and multiple access channels. The ergodic capacity region for the MIMO MAC is well-known and a convex optimization problem. Duality allows the dirty paper coding achievable region for the MIMO BC, a nonconvex region, to be computed from the MIMO MAC capacity region. These capacity and achievable regions are only known for ergodic capacity under perfect CSIT and CSIR. Relatively little is known about the MIMO MAC and BC regions under more realistic partial CSI assumptions. A multicell system with base station cooperation can be modeled as a MIMO BC (downlink) or MIMO MAC (uplink) where the antennas associated with each base station are pooled by the system. Exploiting this antenna structure leads to significant capacity gains over HDR transmission strategies.

There are many open problems in this area. For single-user systems the problems are mainly associated with partial CSI at either the transmitter or receiver. Most capacity regions associated with multiuser MIMO channels remain unsolved, especially ergodic capacity and capacity versus outage for the MIMO BC under perfect receiver CSI only. There are no existing results for partial CSI for any multiuser MIMO channel. Finally, the capacity of cellular systems with multiple antennas remains a relatively open area, in part because the single-cell problem is mostly unsolved, and in part because the Shannon capacity of a cellular system is not well-defined and depends heavily on frequency assumptions and propagation models. Other fundamental tradeoffs in MIMO cellular designs such as whether antennas should be used for sectorization, capacity gain, or diversity are not well understood. In short, we have only scratched the surface in understanding the fundamental capacity limits of systems with multiple transmitter and receiver antennas, and this area of research is likely to remain timely, important, and fruitful for many years to come.

REFERENCES

[1] Channel and propagation models for wireless system design i and ii. *IEEE JSAC*, 20(3,6), April/Aug 2002.

[2] A. Abdi and M. Kaveh. A space-time correlation model for multielement antenna systems in mobile fading channels. *IEEE JSAC*, 20(3):550–561, April 2002.

[3] N. Al-Dhahir. Overview and comparison of equalization schemes for space-time-coded signals with application to edge. *IEEE Trans. Sign. Proc.*, 50(10):2477–2488, Oct 2002.

[4] N. Al-Dhahir, C. Fragouli, A. Stamoulis, W. Younis, and R. Calderbank. Space-time processing for broadband wireless access. *IEEE Commun. Mag.*, pages 136–142, Sept. 2002.

[5] E. Biglieri, J. Proakis, and S. Shamai (Shitz). Fading channels: Information theoretic and communication aspects. *IEEE Trans. Inform. Theory*, 44(6):2619–2692, Oct. 1998.

[6] H. Bolcskei, D. Gesbert, and A. J. Paulraj. On the capacity of OFDM-based spatial multiplexing systems. *IEEE Trans. on Commun.*, 50(2):225–234, Feb 2002.

[7] S. Borst and P Whiting. The use of diversity antennas in high-speed wireless systems: Capacity gains, fairness issues, multi-user scheduling. *Bell Labs. Tech. Mem.*, 2001. Download available at http://mars.bell-labs.com.

[8] S. Boyd and L. Vandenberghe. *Introduction to convex optimization with engineering applications*. Course Reader, 2001.

[9] G. Caire and S. Shamai. On the capacity of some channels with channel state information. *IEEE Trans. Inform. Theory*, 45(6):2007–2019, Sept 1999.

[10] G. Caire and S. Shamai. On achievable rates in a multi-antenna broadcast downlink. In *38th Annual Allerton Conference on Commun., Control and Computing*, pages 1188–1193, Oct. 2000.

[11] S. Catreux, V. Erceg, D. Gesbert, and R. W. Heath. Adaptive modulation and MIMO coding for broadband wireless data networks. *IEEE Commun. Mag.*, pages 108–115, June 2002.

[12] R.S. Cheng and S. Verdu. Gaussian multiaccess channels with ISI: Capacity region and multiuser water-filling. *IEEE Trans. Inform. Theory*, 39(3):773–785, May 1993.

[13] D. Chizhik, G.J. Foschini, M.J. Gans, and R.A. Valenzuela. Keyholes, correlations, and capacities of multielement transmit and receive antennas. *IEEE Trans. Wireless Commun.*, 1(2):361–368, April 2002.

[14] D. Chizhik, J. Ling, P. Wolniansky, R. Valenzuela, N. Costa, and K. Huber. Multiple input multiple output measurements and modeling in manhattan. In *Proceedings of IEEE VTC*, 2002.

[15] Chong and L. Milstein. The performance of a space-time spreading cdma system system with channel estimation errors. In *Proceedings of Intl. Commun. Conf.*, pages 1793–1797, April 2002.

[16] C. Chuah, D. Tse, J. Kahn, and R. Valenzuela. Capacity scaling in MIMO wireless systems under correlated fading. *IEEE Trans. Inform. Theory*, 48(3):637–650, March 2002.

[17] C-N. Chuah, D. N. Tse, J. Kahn, and R. A. Valenzuela. Capacity scaling in MIMO wireless systems under correlated fading. *IEEE Trans. on Inform. Theory*, 48(3):637–650, March 2002.

[18] M. Costa. Writing on dirty paper. *IEEE Trans. Inform. Theory*, 29(3):439–441, May 1983.

[19] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley, 1991.

[20] S. N. Diggavi and T. M. Cover. The worst additive noise under a covariance constraint. *IEEE Trans. Inform. Theory*, 47(7):3072–3081, 2001.

[21] U. Erez, S. Shamai, and R. Zamir. Capacity and lattice strategies for cancelling known interference. In *International Symposium on Information Theory and its Applications*, pages 681–684, Nov. 2000.

[22] G. J. Foschini. Layered space-time architecture for wireless communication in fading environments when using multi-element antennas. *Bell Labs Techn. J.*, pages 41–59, Autumn 1996.

[23] G. J. Foschini, D. Chizhik, M.J. Gans, C. Papadias, and R. A. Valenzuela. Some basic layered spacetime architectures and their performance. *IEEE JSAC Special Issue on MIMO Systems*.

[24] G. J. Foschini and M. J. Gans. On limits of wireless communications in a fading environment when using multiple antennas. *Wireless Personal Commun. : Kluwer Academic Press*, (6):311–335, 1998.

[25] R. G. Gallager. An inequality on the capacity region of multiaccess fading channels. *Communication and Cryptography - Two Sides of one Tapestry*. Boston, MA: Kluwer, 1994, pp. 129-139.

[26] M. J. Gans, N. Amitay, Y. S. Yeh, H. Xu, T.C. Damen, R. A. Valenzuela, T. Sizer, R. Storz, D. Taylor, W. M. MacDonald, C. Tran, and A. Adamiecki. Outdoor BLAST measurement system at 2.44 GHz: Calibration and initial results. *IEEE JSAC*, 20(3):570–581, April 2002.

[27] D. Gesbert, M. Shafi, D. S. Shiu, P. Smith, and A. Naguib. From theory to practice: an overview of MIMO space-time coded wireless systems. *IEEE JSAC Special Issue on MIMO System*.

[28] A. Goldsmith and P. Varaiya. Capacity of fading channels with channel side information. *IEEE Trans. Inform. Theory*, 43:1986–1992, Nov. 1997.

[29] T. Guess and M. K. Varanasi. Multiuser decision-feedback receivers for the general gaussian multiple-access channel. In *Proceedings of Allerton Conf. on Commun., Control, and Comput.*, Oct 1996.

[30] S. Hanly and D.N. Tse. Multiaccess fading channels–Part II:Delay-limited capacities. *IEEE Trans. Inform. Theory*, 44(7):2816–2831, November 1998.

[31] B. Hassibi and B. Hochwald. How much training is needed in multiple-antenna wireless links? Submitted to IEEE Trans. on Info. Theory. Preprint available at http://mars.bell-labs.com.

[32] B. Hassibi and B. Hochwald. Cayley differential unitary space-time codes. *IEEE Trans. Inform. Theory*, 48(6):1485–1503, June 2002.

[33] B. Hochwald, T. L. Marzetta, and V. Tarokh. Multi-antenna channel-hardening and its implications for rate feedback and scheduling. *Submitted to IEEE Trans. Inform. Theory*, 2002. Available at http://mars.bell-labs.com.

[34] B. Hochwald and S. Vishwanath. Space-time multiple access: Linear growth in sum rate. In *Proceedings of 40th Annual Allerton Conference on Commun., Control and Computing*, Oct. 2002.

[35] M. Hochwald, T. L. Marzetta, T. J. Richardson, W. Sweldens, and R. Urbanke. Systematic design of unitary space-time constellations. *IEEE Trans. Inform. Theory*, 46:1962–1973, Sept. 2000.

[36] H. Huang, H. Viswanathan, and G. J. Foschini. Multiple antennas in cellular cdma systems: transmission, detection, and spectral efficiency. *IEEE Trans. Wireless Commun.*, 1(3):383–392, July 2002.

[37] S. Jafar, G. Foschini, and A. Goldsmith. Phantomnet: Exploring optimal multicellular multiple antenna systems. In *Proceedings of VTC*, 2002.

[38] S. Jafar and A. Goldsmith. Multiple-antenna capacity in correlated rayleigh fading with no side information. Preprint. See http://wsl.stanford.edu/Publications.html.

[39] S. Jafar and A. Goldsmith. Transmitter optimization and optimality of beamforming for multiple antenna systems with imperfect feedback. Submitted to IEEE Trans. on Wireless Communications. See http://wsl.stanford.edu/Publications.html.

[40] S. A. Jafar and A. Goldsmith. Transmitter optimization for multiple antenna cellular systems. In *Proceedings of Int. Symp. Inf. Theory*, page 50, June 2002.

[41] S. A. Jafar and A. J. Goldsmith. On optimality of beamforming for multiple antenna systems with imperfect feedback. In *Proceedings of Int. Symp. Inf. Theory*, page 321, June 2001.

[42] S. A. Jafar and A. J. Goldsmith. Vector MAC capacity region with covariance feedback. In *Proceedings of Int. Symp. Inf. Theory*, page 321, June 2001.

[43] S. A. Jafar, S. Vishwanath, and A. J. Goldsmith. Channel capacity and beamforming for multiple transmit and receive antennas with covariance feedback. In *Proceedings of IEEE International Conference on Communications*, volume 7, pages 2266–2270, 2001.

[44] N. Jindal, S. Jafar, S. Vishwanath, and A. Goldsmith. Sum power iterative water-filling for multi-antenna Gaussian broadcast channels. 2002. To appear, Proc. Asilomar Conf.

[45] N. Jindal, S. Vishwanath, and A. Goldsmith. On the duality of Gaussian multiple-access and broadcast channels. In *Proceedings of Int. Symp. Inf. Theory*, page 500, June 2002.

[46] E. Jorswieck and H. Boche. Channel capacity and capacity-range of beamforming in MIMO wireless systems under correlated fading with covariance feedback. Submitted to IEEE JSAC.

[47] E. Jorswieck and H. Boche. Optimal transmission with imperfect channel state information at the transmit antenna array. Submitted to Wireless Personal Communications.

[48] A. Lapidoth and S. M. Moser. Capacity bounds via duality with applications to multi-antenna systems on flat fading

channels. Submitted to IEEE Trans. Inform. Theory. Seehttp://www.isi.ee.ethz.ch/ lapidoth.

[49] L. Li and A. Goldsmith. Capacity and optimal resource allocation for fading broadcast channels–Part I: Ergodic capacity. *IEEE Trans. Inform. Theory*, 47(3):1083–1102, March 2001.

[50] L. Li and A. Goldsmith. Capacity and optimal resource allocation for fading broadcast channels–Part II: Outage capacity. *IEEE Trans. Inform. Theory*, 47(3):1103–1127, March 2001.

[51] Y. Li, J. Winters, and N. Sollenberger. MIMO-OFDM for wireless communication: Signal detection with enhanced channel estimation. *IEEE Trans. on Commun.*, pages 1471–1477, Sept. 2002.

[52] A. Lozano and C. Papadias. Layered space-time receivers for frequency-selective wireless channels. *IEEE Trans. Commun.*, 50(1):65–73, Jan. 2002.

[53] B. Lu, X. Wang, and Y. Li. Iterative receivers for space-time block-coded OFDM systems in dispersive fading channels. *IEEE Trans. on Wireless Commun.*, pages 213–225, April 2002.

[54] T. Marzetta and B. Hochwald. Capacity of a mobile multiple-antenna communication link in rayleigh flat fading. *IEEE Trans. Inform. Theory*, 45(1):139–157, Jan 1999.

[55] T. Marzetta and B. Hochwald. Unitary space-time modulation for multiple-antenna communications in rayleigh flat fading. *IEEE Trans. Inform. Theory*, 46(2):543–564, March 2000.

[56] A.F. Molisch, M. Stienbauer, M. Toeltsch, E. Bonek, and R. S. Thoma. Capacity of MIMO systems based on measured wireless channels. *IEEE JSAC*, 20(3):561–569, April 2002.

[57] A. Moustakas and S. Simon. Optimizing multi-transmitter single-receiver (miso) antenna systems with partial channel knowledge. Preprint. See http://mars.bell-labs.com.

[58] A. Narula, , M. Trott, and G. Wornel. Performance limits of coded diversity methods for transmitter antenna arrays. *IEEE Trans. Inform. Theory*, 45(7), Nov. 1999.

[59] A. Narula, M. J. Lopez, M. D. Trott, and G. W. Wornell. Efficient use of side information in multiple antenna data transmission over fading channels. *IEEE JSAC*, 16(8), Oct. 1998.

[60] Optimal power control in multiple access fading channels with multiple antennas. Writing on colored paper. In *Proceedings of Int. Conf. Commun.*, pages 575–579, 2001.

[61] G.G. Raleigh and J. M. Cioffi. Spatio-temporal coding for wireless communication. *IEEE Trans. Commun.*, 46(3):357–366, March 1998.

[62] F. Rashid-Farrokhi, K.J. R. Liu, and L. Tassiulas. Transit beamforming and power control for cellular wireless systems. *IEEE Journal on Selected Areas in Commun.*, 16(8):1437–1450, Oct 1998.

[63] S. Boyd S. Vishwanath and A. J. Goldsmith. On the worst-case capacity of multi-antenna channels. 2002. Submitted to Int. Conf. Commun.

[64] H. Sato. An outer bound on the capacity region of the broadcast channel. *IEEE Trans. Inform. Theory*, 24:374–377, May 1978.

[65] S. Shamai and A. D. Wyner. Information-theoretic considerations for symmetric, cellular, multiple-access fading channels. *IEEE Trans. Inform. Theory*, 43:1877–1991, Nov. 1997.

[66] S. Shamai and B. M. Zaidel. Enhancing the cellular downlink capacity via co-processing at the transmitting end. In *Proceedings of IEEE Vehicular Technology Conference*, pages 1745–1749, May 2001-Spring.

[67] D. Shiu, G.J. Foschini, M.J. Gans, and J.M. Kahn. Fading correlation and its effect on the capacity of multi-element antenna systems. *IEEE Trans. Commun.*, 48(3):502–513, March 2000.

[68] S. Simon and A. Moustakas. Optimizing MIMO antenna systems with channel covariance feedback. Preprint. See http://mars.bell-labs.com.

[69] P. J. Smith and M. Shafi. On a gaussian approximation to the capacity of wireless MIMO systems. In *Proceedings of Int. Conf. Commun.*, pages 406–410, April 2002.

[70] E. Telatar. Capacity of multi-antenna Gaussian channels. *European Trans. on Telecomm. ETT*, 10(6):585–596, November 1999.

[71] D.N. Tse and S. Hanly. Multiaccess fading channels–Part I:Polymatroid structure, optimal resource allocation and throughput capacities. *IEEE Trans. Inform. Theory*, 44(7):2796–2815, November 1998.

[72] S. Verdu. Spectral efficiency in the wideband regime. *IEEE Trans. Inform. Theory*, 48(6):1319–1343, June 2002.

[73] S. Vishwanath, N. Jindal, and A. Goldsmith. On the capacity of multiple input multiple output broadcast channels. In *Proceedings of Int. Conf. Commun.*, pages 1444–1450, April 2002.

[74] E. Visotsky and U. Madhow. Space-time transmit precoding with imperfect feedback. *IEEE Trans. Inform. Theory*, 47:2632–2639, Sep. 2001.

[75] P. Viswanath, D. Tse, and R. Laroia. Opportunistic beamforming using dumb antennas. *IEEE Trans. Inform. Theory*, 48(6):1277–1294, June 2002.

[76] P. Viswanath and D. N. Tse. Sum capacity of the multiple antenna Gaussian broadcast channel. In *Proceedings of Int. Symp. Inf. Theory*, page 497, June 2002.

[77] P. Viswanath, D. N. Tse, and V. Anantharam. Asymptotically optimal water-filling in vector multiple-access channels. *IEEE Trans. Inf. Theory*, 47(1):241–267, Jan. 2001.

[78] Y. Xin and G. Giannakis. High-rate space-time layered OFDM. *IEEE Commun. Letters*, pages 187–189, May 2002.

[79] W. Yu and J. Cioffi. Trellis precoding for the broadcast channel. In *Proceedings of IEEE Globecom*, pages 1344–1348, Oct. 2001.

[80] W. Yu, G. Ginis, and J. Cioffi. An adaptive multiuser pwoer control algorithm for VDSL. In *Proceedings of IEEE Globecom*, pages 394–398, Oct 2001.

[81] W. Yu, W. Rhee, S. Boyd, and J. Cioffi. Iterative water-filling for vector multiple access channels. In *Proceedings of IEEE International Symposium on Information Theory*, page 322, 2001.

[82] L. Zheng and D. Tse. Optimal diversity-multiplexing tradeoff in multiple antenna channels. In *Proceedings of Allerton Conference*, 2001.

[83] L. Zheng and D. N. Tse. Packing spheres in the grassmann manifold: A geometric approach to the non-coherent multi-antenna channel,. Submitted to IEEE Trans. Inform. Theory.