

Space-Time Wireless Systems:
From Array Processing to MIMO Communications

Edited by

H. Bölcskei, D. Gesbert, C. Papadias, and A. J. van der Veen

1

Perspectives on the Diversity-Multiplexing Tradeoff in MIMO Systems

H. Yao, L. Zheng, and G. W. Wornell
Massachusetts Institute of Technology

1.1 Introduction

Multiple-element antenna arrays have an important role to play in both improving the robustness of wireless transmission through spatial diversity and increasing the throughput of wireless links through spatial multiplexing. Indeed, a wide variety of coding and signal processing techniques have been developed to realize one or the other of these benefits, or combinations of the two.

In recent years there has been an increasing appreciation that the two types of gains are inherently coupled. In particular, for a given antenna, large robustness gains preclude the possibility of large throughput gains, and vice-versa. While the detailed nature of the fundamental tradeoff between these gains is rather complicated, in the high signal-to-noise ratio regime, the scaling behavior of this diversity-multiplexing tradeoff takes a particularly simple form.

To facilitate the use of this tradeoff by communication engineers in system design, this chapter provides an intuitive development of the tradeoff. In addition, we illustrate how a number of systems make these tradeoffs efficiently in various regimes of interest. In the process, a variety of additional interpretations and perspectives on the underlying relationships are developed,

An outline of the chapter is as follows. We begin with a basic channel model and some definitions. We then proceed to develop and interpret the efficient frontier of diversity-multiplexing tradeoffs. The remainder of the chapter focusses on how to achieve different portions of the frontier, what it means to operate on the frontier, and on how performance degrades when operating away from the frontier.

1.2 The Diversity-Multiplexing Tradeoff

1.2.1 Channel Models and System Definitions

We consider a multiple-input multiple-output (MIMO) wireless link with N_t transmit and N_r receive antennas in a flat fading environment, so that the N_r -dimensional complex baseband received signal takes the form

$$\mathbf{y} = \sqrt{\frac{\text{SNR}}{N_t}} \cdot \mathbf{H}\mathbf{x} + \mathbf{w} \quad (1.1)$$

with \mathbf{x} denoting the N_t -dimensional transmitted signal, \mathbf{w} denoting the white Gaussian noise, and \mathbf{H} denoting the $N_t \times N_r$ -dimensional random channel matrix. Finally, we normalize the noise \mathbf{w} and the transmitted signal \mathbf{x} to have unit variance entries, so that SNR is the overall signal-to-noise ratio of the channel. We also use the notation $\rho = \text{SNR}/N_t$ to denote the SNR per transmit antenna for convenience.

In the scenario of interest, the channel matrix \mathbf{H} varies with time. And while there are a number of different models of such time-variation, we adopt the block fading model in this chapter. Specifically, \mathbf{H} is constant within a block of T symbol periods, and varies independently from block to block. We further assume that the blocks are long enough that the receiver has effectively perfect knowledge of each realization of the channel matrix. However, we restrict our attention to the case where there is no feedback, so the transmitter only knows the channel statistics.

Multiple-antenna channels provides *spatial diversity*, which can be used to improve the reliability of the link. In [Tarokh et al., 1998], it is shown that for the multiple antenna channel, the pairwise error probability can be made to decay as $\text{SNR}^{-N_r N_t}$ at high SNR. The improvement of the reliability is thus dictated by the SNR exponent of the error probability, which is referred to as the *diversity gain*. Intuitively, the diversity gain corresponds to the number of independent fading paths that a symbol passes through; in other words, the number of independent fading coefficients that can be averaged over to detect the symbol.

In addition to the spatial diversity gain, it is also well-known that using multiple antennas at both the transmitter and the receiver allows higher data rates to be supported [Telatar, 1995]. To see this, first note that with an input with equal power on each antenna element, and for a fixed realization of the channel, the mutual information between input and output is, in

b/s/Hz,

$$I(\mathbf{x}, \mathbf{y}|\mathbf{H}) = \log \det \left(I + \frac{\text{SNR}}{N_t} \mathbf{H}\mathbf{H}^\dagger \right) = \sum_{i=1}^{\min\{N_t, N_r\}} \log \left(1 + \frac{\text{SNR}}{N_t} \lambda_i \right), \quad (1.2)$$

where the second equality follows from the eigenvalue decomposition $\mathbf{H}\mathbf{H}^\dagger = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\dagger$ with $\mathbf{\Lambda} = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_{\min\{N_t, N_r\}})$ and orthogonal \mathbf{U} . From (1.2) we see immediately that the effect of having multiple antennas is to provide $\min\{N_t, N_r\}$ separate *spatial subchannels* over which to multiplex the data. The ergodic capacity, with the operational meaning of the maximum rate that can be reliably transmitted in a fast fading scenario, is given by

$$C(\text{SNR}) = E[I(\mathbf{x}, \mathbf{y}|\mathbf{H})] = \sum_{i=1}^{\min\{N_t, N_r\}} E \left[\log \left(1 + \frac{\text{SNR}}{N_t} \lambda_i \right) \right], \quad (1.3)$$

Thus, at high SNR, the ergodic capacity increases linearly with $\log \text{SNR}$ according to the number of antennas, i.e.,

$$\lim_{\text{SNR} \rightarrow \infty} \frac{C(\text{SNR})}{\log \text{SNR}} = \min\{N_t, N_r\}$$

When we restrict our attention to codes over a single block (channel realization), as will be our focus, we have to back off from this ergodic capacity (1.3) to achieve small error rates, but can still achieve significant gains. To explore this, we consider rates that are a fixed fraction of this ergodic capacity at high SNR, i.e, rates of the form $R = r \log \text{SNR}$ with $0 < r < \min\{N_t, N_r\}$, and evaluate the associated error behavior. The value r can therefore be viewed as the *spatial multiplexing gain*.

Using (M, SNR) to denote a code with M codewords for a single block of a multiple antenna channel having the specified SNR, and focussing on the high SNR regime, the following more formal definitions are rather natural.

Definition 1 A *spatial multiplexing gain* r is said to be achievable if there exists a sequence of $(2^{TR(\text{SNR})}, \text{SNR})$ codes, with

$$\lim_{\text{SNR} \rightarrow \infty} \frac{R(\text{SNR})}{\log \text{SNR}} = r \quad (1.4)$$

such that the probability of error tends to 0 as $\text{SNR} \rightarrow \infty$.

Definition 2 A *spatial multiplexing gain* r and a *diversity gain* d are said to be simultaneously achievable if there exists a sequence of $(2^{TR(\text{SNR})}, \text{SNR})$

codes, satisfying (1.4) and the error probability $P_e(\text{SNR})$ satisfying

$$\lim_{\text{SNR} \rightarrow \infty} \frac{\log P_e(\text{SNR})}{\log \text{SNR}} = -d \quad (1.5)$$

For future convenience, we rewrite (1.5) as $P_e(\text{SNR}) \doteq \text{SNR}^{-d}$ with \doteq denoting exponential equality.[†]

Before proceeding, note that the two gains concern inherently different aspects of the channel, and thus two channels with the same multiplexing gain may support vastly different diversity gains. As an illustration, consider a pair of parallel scalar ergodic fading channels, with channel gains h_1 and h_2 that are each equally likely to be 0 or 1. Such channels have an ergodic capacity of $\log \text{SNR}$ at high SNR, corresponding to unit multiplexing gain. However, if $h_2 = h_1$, i.e., the two channels fade in unison, the error probability is 1/2 when coding over a single block, corresponding to zero diversity gain. By contrast, if $h_0 = 1 - h_1$, a simple repetition scheme over the single block can achieve the Gaussian channel error probability, which decays exponentially in SNR, and therefore has an infinite diversity gain, according to Definition 2.

Our focus in this chapter is to study the ability of the multiple antenna channels to provide combinations of the diversity and the spatial multiplexing gains and explore the fundamental tradeoff between how much each type of gain one can achieve. As one extreme point, the maximum diversity gain one can achieve is $d_{\max} = N_t N_r$. This can be achieved when transmitting at a fixed data rate, corresponding to a multiplexing gain $r = 0$. On the other hand, from our earlier discussion the maximum spatial multiplexing gain is $r_{\max} = \min\{N_t, N_r\}$. Communicating at a rate so close to the ergodic capacity results in large error probability, corresponding to a diversity gain $d = 0$. In the sequel, we explore the achievable (r, d) pairs between these two extremes.

One can interpret the diversity-multiplexing tradeoff in the context of traditional rate-reliability tradeoffs in communication theory. In particular, while in traditional communication theory one examines achievable rate at fixed SNR in the limit of large block lengths [Cover and Thomas, 1991], in this case we are examining achievable rate at a fixed block length in the limit of high SNR. Accordingly, we refer to the sequence of codes in

[†] Formally, we write $f(\text{SNR}) \doteq \text{SNR}^b$ when

$$\lim_{\text{SNR} \rightarrow \infty} \frac{\log f(\text{SNR})}{\log \text{SNR}} = b.$$

The symbols \gtrsim and \lesssim are similarly defined.

the above definition, indexed by SNR, with $2^{Tr \log \text{SNR}}$ codeword in each code, as a *scheme*. By studying the performance of a scheme, we focus not on the design of specific codes, but rather on the structure of MIMO codes. Examples of such schemes include random codes with the number of codewords increasing with SNR, algebraic codes over QAM symbols with the size of the QAM constellation increasing with SNR, etc.

With this interpretation, the multiplexing gain can be viewed as the number of bits reliably conveyed per dB of the SNR, where reliable means the error probability can be driven to zero with increasing $\log \text{SNR}$ rather than the block length L . Thus, while traditional communication theory examines the exponential decay of error probability with blocklength, here we examine the exponential decay of error probability with $\log \text{SNR}$. Furthermore, the number of bits that can be conveyed is linear in $\log \text{SNR}$ just as it is in block length L .

In our setup, we focus on transmitting at a fixed fraction of the maximum multiplexing gain and examine the exponent in the decay of error probability with $\log \text{SNR}$ for a fixed block length. By comparison, the traditional error exponent analysis of Gallager focusses on transmitting at a fixed fraction of capacity and examines the exponent in the decay of error probability with blocklength L . Thus, there is a natural connection to traditional capacity and error exponent analysis in communication theory; further discussion can be found in, e.g., [Zheng and Tse, 2003].

Before proceeding, we should emphasize at this point that the measures of reliability and rate provided by the definitions are rather coarse in some respects. In particular, the reliability measure is invariant to changes in coding gain, which scales the SNR inside the logarithm, and other such constant factors. As such, two systems that achieve a particular diversity-multiplexing operating point may not be equally attractive when subjected to a finer scale analysis in which the effects of factors like coding gain manifest themselves.

1.2.2 System Outage Characteristics

When the transmitter does not know the channel matrix and is coding within the coherence time of the channel, two types of errors are possible at the receiver. One type is due to noise realization, and the other due to the channel realization. As we now develop, it is the latter that dominates the error behavior in our scenario of interest.

Errors due to the channel realization occur as follows. Because the transmitter must choose a transmission rate without knowledge of the channel




Fig. 1.1. (a) Family of outage probability curves as functions of SNR for various target rates R in the $N_t = N_r = 2$ case. (b) Tradeoff between increasing the data rate and decreasing the error probability

matrix, it may choose a rate that even with the best possible coding and long coherence times cannot be supported by the channel. This corresponds to the rate R being higher than the mutual information of the realized channel, and the corresponding error event, termed *outage*, is defined via

$$\{I(x; y|\mathbf{H}) < R\} \quad (1.6)$$

Since the channel mutual information is a random variable $I = I(x; y|\mathbf{H})$, there is a probability associated with this error event, which we denote using $P_{\text{out}}(R, \text{SNR})$.

The overall error probability can be expressed in terms of the outage probability according to

$$P_e(\text{SNR}) = P(\text{error}|\text{outage})P_{\text{out}}(R, \text{SNR}) \\ + P(\text{error}|\text{no outage})(1 - P_{\text{out}}(R, \text{SNR}))$$

Conditioned on the outage event, the error probability is necessarily close to one and thus the outage probability is an asymptotic lower bound on the error probability, i.e., $P_e(\text{SNR}) \geq P_{\text{out}}(R, \text{SNR})$. It is for this reason that traditional space-time code design focusses on approaching the outage behavior as closely as possible.

We can visualize the relationship between SNR, rate, and outage probability by plotting P_{out} as a function of SNR for various rates R . This is depicted in Figure 1.1(a) for the case of two transmit and two receive antennas. Each curve represents how outage probability decays with SNR for a fixed rate R .

Observe that for each curve, corresponding to a fixed rate, at sufficiently

high SNR the slope of the curve approaches 4, i.e., the outage probability decreases by 10^4 for every 10 dB increase in SNR. This corresponds to the maximum diversity gain of $d_{\max} = 4$ available from the channel. Following an individual curve corresponds to using increases in SNR exclusively for improving the reliability of the link while keeping the data rate fixed. On the other hand, for a fixed outage probability, increases in R move the curves toward higher SNR. At sufficiently high SNR, the gaps between the curves approaches 3 dB, i.e, the rate can increase by 2 b/s/Hz for every 3 dB increase in SNR. This corresponds to the maximum multiplexing gain of $r_{\max} = 2$ available from the channel. This horizontal transecting of the curves corresponds to using increases in SNR exclusively for increasing the data rate, while keeping the outage probability fixed.

More generally, we are interested in transecting these curves not horizontally, but with some downward slope. Indeed, a coding scheme in our parlance corresponds to the resulting sequence of points on these plots whereby as the SNR increases, the data rate increased and the error probability is decreased so as to achieve both diversity and multiplexing gains simultaneously. Figure 1.1(b) shows these cross-cutting curves of various downward slopes. As the figure reflects, there is a tradeoff between how much of each type of gain can be obtained: a higher multiplexing gain, corresponding to increase the data rate faster with SNR, will result in a lower diversity gain, and vice versa.

1.2.3 The Efficient Frontier

The optimal tradeoff between the diversity gain and the multiplexing gain is described by a function $d^*(\cdot)$ whereby $d^*(r)$ gives the maximum achievable diversity gain at each multiplexing gain r .

The optimal tradeoff in the case of independent, identically-distributed Rayleigh fading between each pair of antennas, and sufficiently large coherence times, i.e., $T \geq N_t + N_r - 1$, is developed in [Zheng and Tse, 2003]. For this case, the tradeoff is given by the convex hull of the points $(k, d^*(k))$, $k = 0, 1, \dots, \min(N_t, N_r)$, where

$$d^*(k) = (N_t - k)(N_r - k). \tag{1.7}$$

The function $d^*(r)$ is plotted in Figure 1.2(a).

Before proceeding, it is worth noting that it is also shown that at any r , the optimum achievable diversity gain $d^*(r)$ is the same as the SNR exponent of the outage probability, i.e., for any r ,

$$P_e(\text{SNR}) \doteq P_{\text{out}}(R = r \log \text{SNR}, \text{SNR}) \doteq \text{SNR}^{-d^*(r)}. \tag{1.8}$$

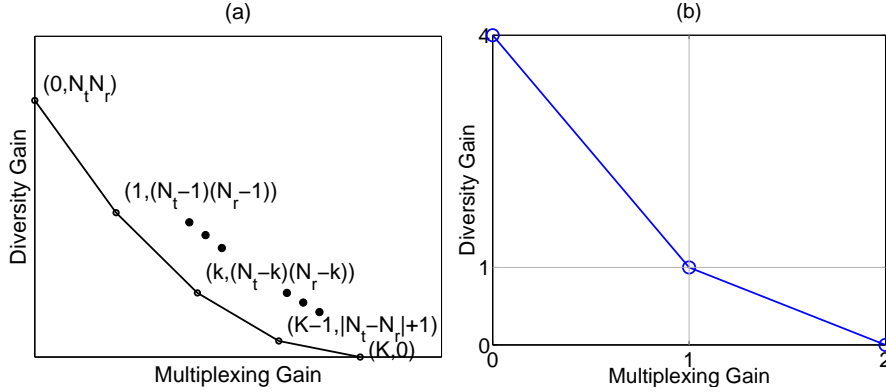


Fig. 1.2. (a) The efficient frontier $d^*(r)$ of diversity-multiplexing tradeoffs for a system with N_t transmit antennas and N_r receive antennas, where $K \stackrel{\text{def}}{=} \min(N_t, N_r)$. (b) The efficient frontier for $N_t = N_r = 2$.

This means that the upper limit of the performance given by the outage probability can actually be achieved. While this is not surprising in the limit of infinite coherence time since one can drive noise-induced errors to zero using good long codes, it is noteworthy that it also holds for finite coherence times. In essence, it tells us that when using good signaling schemes, the probability of error is asymptotically dominated by the outage probability, and that the typical way of making an error in the MIMO fading channel is from the channel being in a deep fade. In the sequel, we will develop some intuition for this result. For this purpose, it will be sufficient to focus on the 2×2 system, for which the optimal tradeoff curve is plotted in Figure 1.2(b).

We begin by relating the shape of the frontier to the dependence of error probability on rate and SNR for optimal systems. To this end, in Figure 1.3(a) we plot the outage probability (1.8) as a function of SNR for different values of r . To facilitate its interpretation, the curves of Figure 1.1(a) are overlaid as gray lines. Dashed lines with slopes $d^*(r)$ are also drawn to show the match between the asymptotic expression (1.8) and the exact outage probability.

The asymptotes corresponding to outage probability (1.8) are depicted in the Bode-style plot of Figure 1.3(b). Such a plot reveals two distinct regions of operation. The case where the rate is high relative to the SNR corresponds to a *heavily-loaded* system, and its performance is captured by the region of the plot above and to the right of the dashed line. Here the slopes of the outage probability curves are flatter. By contrast, performance

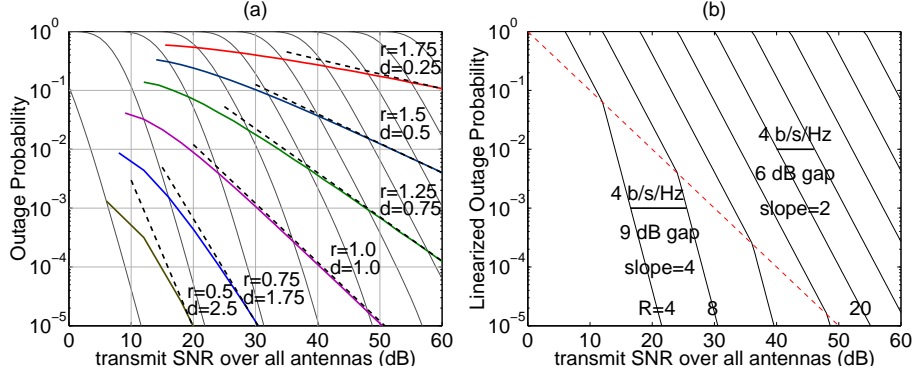


Fig. 1.3. (a) As rate grows with SNR, i.e., $R = r \log_2(\text{SNR})$, outage probability $P_{\text{out}}(R, \text{SNR})$ decays with SNR with slope $d^*(r)$. (b) Linearized approximation of Figure 1.1(a), which emphasizes the two distinct regions of the P_{out} -SNR space with differing curve slopes and horizontal spacings.

in the *lightly-loaded* regime is captured by the region below and to the left of the dashed line. Here the slopes of the outage curves are steeper. The dashed boundary is the line with $P_{\text{out}} = \text{SNR}^{-1}$, and corresponds to the (1,1) point of the frontier curve in Figure 1.2(b). More generally, there are $\min\{N_t, N_r\}$ distinct regions of operation, one corresponding to each piecewise linear segment of the frontier $d^*(r)$.

Note that associated with each of these regions is a set of inherently different *local* tradeoffs around each operating point in the region. To see this, observe from Figure 1.3(b) that in the heavily-loaded region, the lines have slope 2 and spacing of 1.5 dB for every 1 b/s/Hz increment in rate. Thus 3 dB of additional SNR can be used to either increase rate by 2 b/s/Hz or reduce error probability by a factor of $(1/2)^2$. More generally, one can achieve any linear combination of the two, i.e., for any $\alpha \in (0, 1)$, increase rate by 2α b/s/Hz while decreasing the error probability by $(1/2)^{2(1-\alpha)}$. In terms of Figure 1.2(b), this local tradeoff corresponds to a straight line connecting $(r, d) = (0, 2)$ and $(2, 0)$, which is an extension of the lower segment of the frontier down to $r = 0$. As we would expect, the maximum diversity gain of 4 is not achievable in this heavily-loaded region. Similarly, when we operate in the lightly-loaded region, the local tradeoff corresponds to an extension of the upper segment of the frontier down to $d = 0$, i.e., the straight line connecting $(0, 4)$ and $(4/3, 0)$. In this region, it is the maximum multiplexing gain of 2 that is not achieved.

To understand why the efficient frontier of diversity-multiplexing tradeoffs takes the particular form it does, it is convenient to rewrite the outage event

(1.6) in the form

$$\left\{ \prod_{i=1}^{\min\{N_t, N_r\}} (1 + \text{SNR}/N_t \cdot \lambda_i) < \text{SNR}^r \right\}. \quad (1.9)$$

If all the eigenvalues λ_i are large relative to the noise threshold, then each corresponding spatial subchannel is operating in the high-SNR regime, contributing a factor SNR to the lefthand side of (1.9), and no outage occurs. However, if some of the eigenvalues are small relative to the noise threshold, then those subchannels are operating in the low-SNR regime and do not contribute a factor of SNR . Thus, if there are not at least r large eigenvalues, outage occurs. From this perspective, the piecewise linear nature of the diversity-multiplexing frontier reflects that eigenvalues of $\mathbf{H}\mathbf{H}^\dagger$ are generally not all equal, i.e, that the corresponding spatial subchannels are not all equally strong.

This behavior underlies the distinct operating regions in Figure 1.3(b). In particular, in this case, when the eigenvalues are ordered according to $\lambda_1 \geq \lambda_2$, the smaller eigenvalue λ_2 takes values close to zero with a much larger probability than λ_1 does. Consequently, when the system is heavily-loaded ($r \geq 1$), the tradeoff curve is flatter since an outage occurs whenever λ_2 alone is small. By contrast, when the system is lightly-loaded ($r \leq 1$), both λ_1 and λ_2 must be small to get an outage, thus the tradeoff curve is steeper. While this can be quantified by deriving the joint distribution of ordered eigenvalues [Zheng and Tse, 2003], in the sequel we make use of a different decomposition of the channel instead, which also provides useful insights into code design.

For a 2×2 system, the mutual information (1.2) of the realized channel can be rewritten as

$$I(x; y | \mathbf{H}) = \log (1 + \rho \|\mathbf{H}\|^2 + \rho^2 |\det(\mathbf{H})|^2). \quad (1.10)$$

From this we see that $\det(\mathbf{H})$ being large relative to the noise is sufficient for both spatial subchannels to be in the high SNR regime (enabling a multiplexing gain of two), and that $\|\mathbf{H}\|$ being large relative to the noise is sufficient for at least one subchannel to be in the high SNR regime (enabling a multiplexing gain of one).

Using the factorization $\mathbf{H} = \mathbf{Q}\mathbf{R}$, where \mathbf{Q} is unitary and

$$\mathbf{R} = \begin{bmatrix} r_{11} & r_{12} \\ 0 & r_{22} \end{bmatrix}, \quad (1.11)$$

we see

$$\|\mathbf{H}\|^2 = r_{11}^2 + |r_{12}|^2 + r_{22}^2 \quad \text{and} \quad |\det(\mathbf{H})|^2 = r_{11}^2 r_{22}^2. \quad (1.12)$$

Letting \mathbf{h}_1 and \mathbf{h}_2 denote the columns of \mathbf{H} , and $\mathbf{h}_{2\parallel 1}$ and $\mathbf{h}_{2\perp 1}$ denote the components of \mathbf{h}_2 in the direction along and perpendicular to \mathbf{h}_1 , we see that

$$r_{11}^2 = \|\mathbf{h}_1\|^2 \quad \text{and} \quad |r_{12}|^2 = \|\mathbf{h}_{2\parallel 1}\|^2 \quad \text{and} \quad r_{22}^2 = \|\mathbf{h}_{2\perp 1}\|^2 \quad (1.13)$$

Thus for the case of interest where \mathbf{H} has independent identically distributed circularly symmetric complex Gaussian entries, r_{11}^2 , $|r_{12}|^2$, and r_{22}^2 are chi-squared random variables of order 4, 2, and 2, respectively. Since for $\alpha < 1$ we have $\Pr[\chi < \alpha] \doteq \alpha^{k/2}$ when χ is chi-squared of order k , having a small r_{11}^2 is less likely than having either small $|r_{12}|^2$ or r_{22}^2 . Therefore, all of r_{11}^2 , $|r_{12}|^2$, and r_{22}^2 need to be small in order for $\|\mathbf{H}\|^2$ to be small, while $|\det(\mathbf{H})|^2$ being small is most likely due to r_{22}^2 being small.

Now we are ready to observe the difference between the typical outage events for the cases $r > 1$ and $r < 1$, the heavily loaded and the lightly loaded regimes, respectively. For $r > 1$, the outage event

$$\left\{ \log \left(1 + \rho \|\mathbf{H}\|^2 + \rho^2 r_{11}^2 r_{22}^2 \right) \leq r \log \rho \right\}$$

occurs when $\rho^2 r_{11}^2 r_{22}^2 \leq \rho^r$. This happens typically when r_{22}^2 is as small as $\rho^{-(2-r)}$, with a probability of the order $\rho^{-(2-r)}$. Intuitively, when \mathbf{h}_1 and \mathbf{h}_2 are closely aligned with each other, i.e., $\|\mathbf{h}_{2\perp 1}\|^2 = r_{22}^2$ is abnormally small, the channel matrix \mathbf{H} becomes close to singular, and less rate is supported by the weaker spatial channel. In the extreme, when $r_{22}^2 \doteq \rho^{-1}$, with a probability ρ^{-1} , the weaker spatial channel supports zero multiplexing gain, and can be viewed as completely shut off. For even lower rate $r < 1$, the outage event occurs only if both $\rho \|\mathbf{H}\|^2$ and $\rho^2 r_{11}^2 r_{22}^2$ are less than ρ^r . Typically, this happens when the weaker spatial channel is shut off, $r_{22}^2 \doteq \rho^{-1}$, and the stronger spatial channel can barely support the rate, $r_{11}^2 + |r_{12}|^2 \doteq \rho^{-(1-r)}$. Since the stronger channel is better protected, there is a steeper slope in the tradeoff curve when $r < 1$.

For a general $N_t \times N_r$ channel, the tradeoff result suggests that k spatial subchannels are operating in the high SNR regime, supporting a multiplexing gain of $r = k$, with probability $\text{SNR}^{-(N_r - k)(N_t - k)}$, corresponding to the typical outage event that the $\min\{N_t, N_r\} - k$ smallest eigenvalues of $\mathbf{H}\mathbf{H}^\dagger$ are all of the same order as the noise.

1.3 Achieving Diversity-Multiplexing Frontier

To develop additional insights into diversity-multiplexing tradeoffs, we now turn our attention to developing MIMO coding schemes that achieve different portions of the efficient frontier. In the process we see how the structure of the associated schemes determines their diversity-multiplexing characteristics. While this section focuses on the two-transmit two-receive antenna case, many of the design principles apply to larger systems.

1.3.1 Achieving Full Diversity Gain or Full Multiplexing Gain

This section examines two well-known low-complexity MIMO coding schemes, each achieving exactly one point of the diversity-multiplexing frontier. One achieves the full multiplexing gain, while the other achieves full diversity.

1.3.1.1 Achieving Full Diversity Gain

The Orthogonal Space Time Block Code (OSTBC) [Alamouti, 1998] employs a smart repetition to achieve the full diversity gain. It only applies to systems with two transmit antennas. It encodes two independent information symbols, s_1 and s_2 , into the transmit signal matrix \mathbf{X} using,

$$\mathbf{X} = \begin{bmatrix} s_1 & -s_2^* \\ s_2 & s_1^* \end{bmatrix}, \quad (1.14)$$

where $(\cdot)^*$ indicates conjugation. The OSTBC effectively transforms a 2×2 MIMO channel to a scalar channels with channel gain $\|\mathbf{H}\|$. When OSTBC is combined with capacity-achieving error correction codes (ECC) designed for scalar channels, it is possible to communicate at rates

$$R < \log_2 (1 + \rho \|\mathbf{H}\|^2). \quad (1.15)$$

The associated diversity-multiplexing tradeoff can be derived using

$$\begin{aligned} P_e &= \Pr [\log_2 (1 + \rho \|\mathbf{H}\|^2) < R = r \log_2 \rho] \\ &\doteq \Pr [\|\mathbf{H}\|^2 < \rho^{r-1}] = \rho^{4(r-1)} = \rho^{-d(r)}. \end{aligned} \quad (1.16)$$

Therefore, the tradeoff achieved is $d(r) = 4(1-r)$, i.e., a straight line between $(0, 4)$ and $(1, 0)$, which is plotted in Fig. 1.4(a). It shows that only the full diversity point of the frontier is achieved. The lesson is that by transmitting each information symbol twice using different antennas at different times, each symbol interacts with all entries of the channel matrix \mathbf{H} . As a result, the dominant error event is when all entries of \mathbf{H} are small. Thus, the full diversity gain is achieved. This is true even when there are more than two

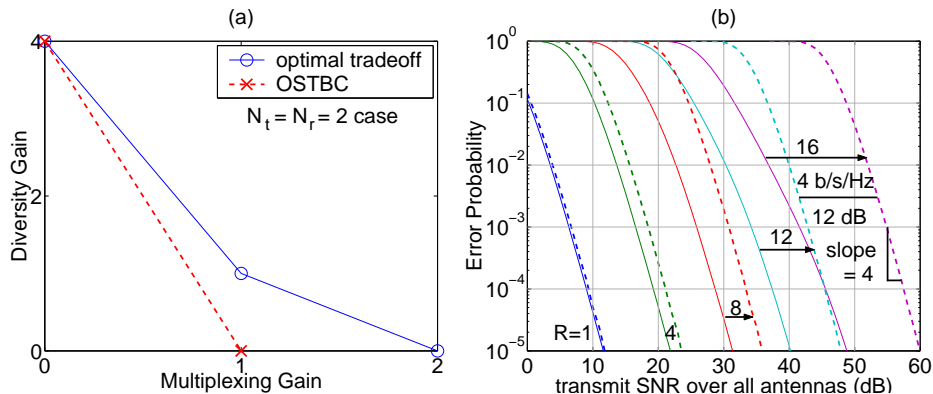


Fig. 1.4. (a) Diversity-multiplexing tradeoff achieved by OSTBC. (b) Comparison of the OSTBC error probability curves (dashed) and the channel outage probability curves (thin solid) as functions of SNR for rates 1, 4, 8, 12, and 16 b/s/Hz.

receive antennas. The drawback of OSTBC is that the maximum multiplexing gain achieved is only 1 since only one symbol is effectively transmitted at a time due to the repetition. It is only in the two-transmit one-receive antenna case that OSTBC achieves the optimal diversity-multiplexing tradeoff [Zheng and Tse, 2003].

The error probability performance of OSTBC can be visualized by plotting $\Pr[\log_2(1 + \rho\|\mathbf{H}\|^2) < R]$ as functions of SNR for various R , via Monte-Carlo simulation, as shown in Fig. 1.4(b) with dashed lines. These curves approximately form a set of parallel lines with slopes of 4 and horizontal gaps of 12 dB per 4 b/s/Hz, i.e., 1 b/s/Hz for every 3 dB. The slopes and gaps are consistent with the tradeoff points (0, 4) and (1, 0).

The channel outage probability curves are also plotted in Fig. 1.4(b) as thin solid lines for comparison. Below $R = 4$ b/s/Hz, the OSTBC curves are actually very close to the outage performance limit. This is because (1.15) differs from the channel capacity expression in (1.10) only by missing the second order term, which is negligible at low SNR. However, as rate increases, the performance gaps increase indefinitely.

1.3.1.2 Achieving Full Multiplexing Gain

The vertical Bell LABs LAYered Space Time (V-BLAST) code transmits two independently encoded codewords simultaneously using the two antennas as depicted in Fig. 1.5. By transmitting two independent symbols per channel use, it achieves the full multiplexing gain. Unlike OSTBC, V-BLAST applies to systems with any number of transmit antennas.

The basic form of V-BLAST scheme employs a low-complexity nulling




Fig. 1.5. V-BLAST, where coding is restricted to one row of the transmitted signal matrix.

and cancellation decoding. We refer to this scheme as V-BLAST-nulling. The received vector y is multiplied with the \mathbf{Q}^\dagger matrix in (1.11), making the effective channel matrix $\mathbf{Q}^\dagger \mathbf{H} = \mathbf{R}$ upper triangular. Consequently, codeword 2 experiences no interference and can be decoded first with channel gain r_{22}^2 . Assuming codeword 2 is correctly decoded, its interference on codeword 1 can then be canceled, and codeword 1 can be decoded with channel gain r_{11}^2 . Therefore, V-BLAST-nulling transforms a 2×2 MIMO channel to two parallel scalar channels with channel gains r_{11}^2 and r_{22}^2 , respectively.

Assuming the data rate R is split evenly between the two codewords and the codes are capacity achieving, the necessary and sufficient condition for successful decoding is that both subchannels can support rate $R/2$, i.e.,

$$\frac{R}{2} < \log_2(1 + \rho r_{22}^2) \quad \text{and} \quad \frac{R}{2} < \log_2(1 + \rho r_{11}^2). \quad (1.17)$$

The associated diversity-multiplexing tradeoff can be derived using

$$\begin{aligned} P_e &= \Pr \left[\log_2(\rho r_{22}^2 + 1) < \frac{R}{2} \text{ or } \log_2(\rho r_{11}^2 + 1) < \frac{R}{2} \right] \\ &\doteq \Pr \left[r_{22}^2 < \rho^{\frac{R}{2}-1} \text{ or } r_{11}^2 < \rho^{\frac{R}{2}-1} \right] = \rho^{\frac{R}{2}-1} = \rho^{-d(r)}. \end{aligned} \quad (1.18)$$

Therefore, the tradeoff achieved is $d(r) = 1 - r/2$, i.e., a straight line between $(0, 1)$ and $(2, 0)$, which is plotted in Fig. 1.6(a). It shows that only the full multiplexing gain point of the frontier is achieved. The drawback of V-BLAST-nulling is that the maximum diversity gain achieved is only 1, because r_{22}^2 is distributed like the energy of one entry of \mathbf{H} . The dominant error event is having a small r_{22}^2 and failing to decode the first codeword.

The error probability performance of V-BLAST-nulling can be visualized by plotting $\Pr \left[R > 2 \log_2(1 + \rho r_{22}^2) \text{ or } R > 2 \log_2(1 + \rho r_{11}^2) \right]$ as functions of SNR for various R , via Monte-Carlo simulation, as shown in Fig. 1.6(b) with dashed lines. These curves approximately form a set of parallel lines with slopes of 1 and horizontal gaps of 6 dB per 4 b/s/Hz, i.e., 2 b/s/Hz for every 3 dB. The slopes and gaps are consistent with the tradeoff points $(0, 1)$ and $(2, 0)$. Comparing to the channel outage probability curves plotted with thin

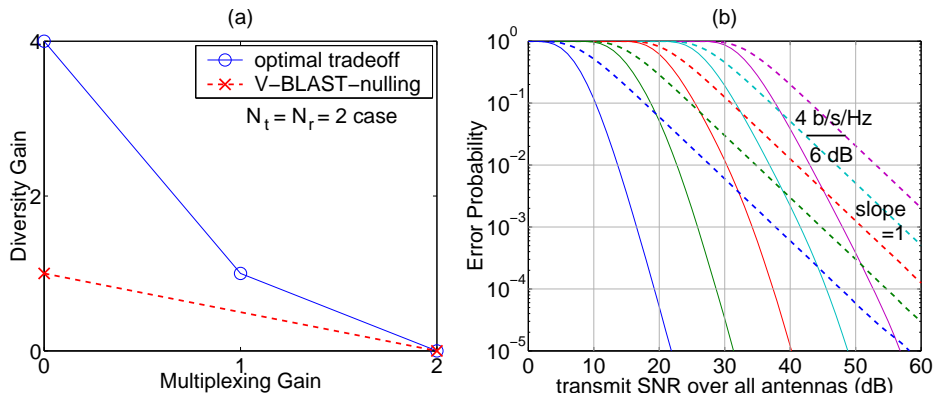


Fig. 1.6. (a) Diversity-multiplexing tradeoff achieved by V-BLAST encoding with nulling and cancellation decoding. (b) Comparison of the V-BLAST-nulling error probability curves (dashed) and the channel outage probability curves (thin solid) as functions of SNR for rates 4, 8, 12, 16, and 20 b/s/Hz.

solid lines in Fig. 1.6(b), the performance gaps in SNR increase quickly as the target error rate decreases due to the difference in slopes.

V-BLAST decoding can be improved by choosing the decoding order to optimize the effective channel gains for each codeword and employing MMSE decoding rather than nulling. However, these variations do not significantly affect the diversity-multiplexing tradeoff achieved.

1.3.2 Achieving the High-Multiplexing-Gain or High-Diversity-Gain Frontiers

This section discusses two coding schemes each achieving one segment of the diversity-multiplexing frontier in Fig. 1.2(b). We refer to the two segments as the high-diversity-gain frontier and the high-multiplexing gain frontier. The high-diversity-gain frontier corresponds to lightly-loaded systems where the multiplexing gain is less than one and rate is relatively small compared to SNR. The high-multiplexing gain frontier corresponds to heavily-loaded systems. We note that while the coding schemes discussed in this section achieve more points on the frontier than OSTBC and V-BLAST-nulling, they require higher complexity joint decoding.

1.3.2.1 Achieving the High-Multiplexing-Gain Frontier

The V-BLAST encoding scheme presented in Section 1.3.1.2 when combined with joint decoding can achieve the high-multiplexing-gain segment of the diversity-multiplexing frontier by optimally handling the interference be-




Fig. 1.7. (a) Diversity-multiplexing tradeoffs achieved by V-BLAST encoding with joint decoding. (b) Comparison of V-BLAST-joint error probability curves (dashed) and the channel outage probability curves (thin solid) as functions of SNR for rates 4, 8, 12, 16, and 20 b/s/Hz.

tween the codewords. We refer to this scheme as V-BLAST-joint. Again assuming the data rate R is split evenly between the two separately-encoded codewords and the codes are capacity achieving, with joint decoding, the necessary and sufficient condition for successful decoding is that the total rate can be supported by the channel and that the individual codewords can be decoded in the absence of interference from the other codeword, i.e.,

$$R < \log_2(1 + \rho\|\mathbf{H}\|^2 + \rho^2|\det(\mathbf{H})|^2),$$

$$\frac{R}{2} < \log_2(1 + \rho\|\mathbf{h}_1\|^2), \text{ and } \frac{R}{2} < \log_2(1 + \rho\|\mathbf{h}_2\|^2). \quad (1.19)$$

with \mathbf{h}_1 and \mathbf{h}_2 continuing to denote the columns of \mathbf{H} . Eq. (1.19) shows that the dominant error event for V-BLAST-joint is either $\|\mathbf{h}_1\|^2$ or $\|\mathbf{h}_2\|^2$ being small. The associated diversity-multiplexing tradeoff can be derived using

$$P_e \doteq \Pr \left[\|\mathbf{h}_1\|^2 < \rho^{\frac{r}{2}-1} \text{ or } \|\mathbf{h}_2\|^2 < \rho^{\frac{r}{2}-1} \right] = \rho^{2(\frac{r}{2}-1)} = \rho^{-d(r)}. \quad (1.20)$$

Therefore, the tradeoff achieved is $d(r) = 2 - r$, i.e., a straight line between (0, 2) and (2, 0), which is plotted in Fig. 1.7(a). It is clear that the high-multiplexing-gain segment of the frontier is achieved. The key idea is that with V-BLAST encoding, each transmitted symbol interacts with one column of the channel matrix \mathbf{H} . By fully taking advantage of this with joint decoding, entire columns of \mathbf{H} has to be small for decoding to fail. Since each column of \mathbf{H} has two entries, the maximum diversity gain achieved is 2.

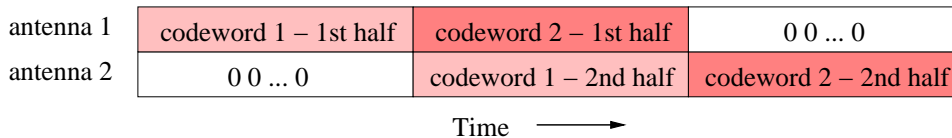


Fig. 1.8. D-BLAST encoding with two layers. Two halves of two codewords are transmitted by a certain antenna at a certain time as indicated. 0 indicates silence.

The error probability performance of V-BLAST-joint can be visualized by plotting the probability that the condition in (1.19) fails as functions of SNR for various R , via Monte-Carlo simulation, as shown in Fig. 1.7(b) with dashed lines. These curves approximately form a set of parallel lines with slopes of 2 and horizontal gaps of 6 dB per 4 b/s/Hz, i.e., 2 b/s/Hz for every 3 dB. The slopes and gaps are consistent with the tradeoff points (0, 2) and (2, 0). Comparing to the channel outage probability curves plotted with thin solid lines in Fig. 1.7(b), V-BLAST-joint performs quite well in the region above $P_e = \text{SNR}^{-1}$. The reason is this region corresponds to the high-multiplexing segment of the tradeoff curve.

1.3.2.2 Achieving the High-Diversity-Gain Frontier

The high-diversity-gain segment of the diversity-multiplexing frontier can be achieved with a two-layer diagonal-BLAST (D-BLAST) encoding scheme with joint decoding. We refer to it as 2L-D-BLAST-joint. While V-BLAST places separately encoded codewords on different antennas, D-BLAST transmits each codeword using different antennas at different times as illustrated in Fig. 1.8 for the two-layer case. Note that the antennas are sometimes silent. Similar to the V-BLAST case, the codewords are assumed to be capacity achieving and have equal rate. Although decoding can be done one codeword at a time via successive cancellation, we consider joint decoding in this section, since successive-cancellation decoding does not achieve the high-diversity-gain segment of the frontier.

The necessary and sufficient condition for successful joint decoding is that the total rate can be supported by the channel and the individual codeword can be decoded in the absence of interference from the other codeword, i.e.,

$$3R < \log_2(1 + \rho \|\mathbf{h}_1\|^2) + \log_2(1 + \rho \|\mathbf{H}\|^2 + \rho^2 |\det(\mathbf{H})|^2) + \log_2(1 + \rho \|\mathbf{h}_2\|^2) \quad (1.21)$$

$$\text{and } \frac{3R}{2} < \log_2(1 + \rho \|\mathbf{h}_1\|^2) + \log_2(1 + \rho \|\mathbf{h}_2\|^2). \quad (1.22)$$

Since $\log_2(1 + \rho \|\mathbf{H}\|^2 + \rho^2 |\det(\mathbf{H})|^2) < \log_2(1 + \rho \|\mathbf{h}_1\|^2) + \log_2(1 + \rho \|\mathbf{h}_2\|^2)$, (1.22) is always satisfied when (1.21) is. Therefore, (1.21) by itself is the necessary and sufficient condition for successful decoding.




Fig. 1.9. (a) Diversity-multiplexing tradeoffs achieved by two-layer D-BLAST encoding with joint decoding. (b) Comparison of 2L-D-BLAST-joint error probability curves (dashed) and the channel outage probability curves (thin solid) as functions of SNR for rates 4, 8, 12, 16, and 20 b/s/Hz.

The dominant event for the condition in (1.21) to fail is $\|\mathbf{H}\|^2 < \rho^{\frac{3}{4}r-1}$, which implies that all of $\|\mathbf{h}_1\|^2$, $\|\mathbf{h}_2\|^2$, and $|\det(\mathbf{H})|$ are just as small. The associated diversity-multiplexing tradeoff can be derived using

$$P_e \doteq \Pr \left[\|\mathbf{H}\|^2 < \rho^{\frac{3}{4}r-1} \right] = \rho^{4(\frac{3}{4}r-1)} = \rho^{-d(r)}. \quad (1.23)$$

Therefore, the tradeoff achieved is $d(r) = 4 - 3r$, i.e., a straight line between $(0, 4)$ and $(4/3, 0)$, which is plotted in Fig. 1.9(a). It is clear that the high-diversity-gain segment of the frontier is achieved. The 2L-D-BLAST-joint scheme achieves the full diversity gain because each codeword interacts with all entries of the channel matrix \mathbf{H} , just like in OSTBC. Furthermore, the encoding structure is such that four information symbols are transmitted in three channel uses, which results in the maximum multiplexing gain of $4/3$.

The error probability performance of 2L-D-BLAST-joint can be visualized by plotting the probability that the condition in (1.21) fails as functions of SNR for various R , via Monte-Carlo simulation, as shown in Fig. 1.9(b) with dashed lines. These curves approximately form a set of parallel lines with slopes of 4 and horizontal gaps of 9 dB per 4 b/s/Hz, i.e., $4/3$ b/s/Hz per 3 dB. The slopes and gaps are consistent with the tradeoff points $(0, 4)$ and $(4/3, 0)$. Compared to the channel outage probability curves plotted with thin solid lines in Fig. 1.9(b), 2L-D-BLAST-joint performs quite well in the region below $P_e = \text{SNR}^{-1}$. The reason is this region corresponds to the high-diversity segment of the tradeoff curve as discussed in Section 1.2.3.

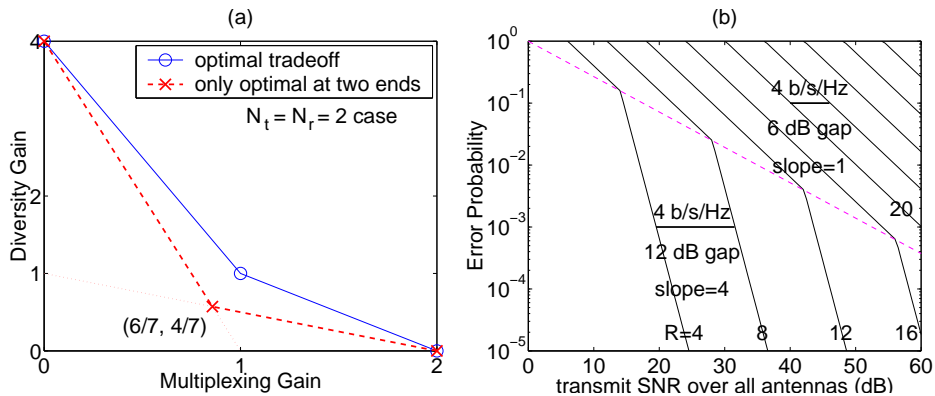


Fig. 1.10. (a) A hypothetical diversity-multiplexing tradeoff that only meets the optimal tradeoff at the end points, $d(r) = \max(4 - 4r, 1 - r/2)$. (b) A family of error probability curves consistent with the tradeoff for rates 4, 8, \dots 36 b/s/Hz.

1.3.3 Full-Diversity Full-Multiplexing Schemes

We next examine the implications of code constructions where both the full diversity gain and the full multiplexing gain are achieved, but not the rest of the frontier. Our focus in this section is not on the description of schemes with this property. Rather, it is to demonstrate how performance can degrade rather severely relative to schemes that achieve the entire frontier, to help code and system designers understand the underlying issues.

Let us consider the tradeoff $d(r) = \max(4 - 4r, 1 - r/2)$, $0 \leq r \leq 2$, which is suboptimal everywhere except at the end points, as shown in Fig. 1.10(a). The error probability function that is consistent with this tradeoff curve is $P_e(R, \text{SNR}) = \text{SNR}^{-d(R/\log_2(\text{SNR}))}$, which is plotted in Fig. 1.10(b) for a range of R and SNR, similar to Fig. 1.3(b) for the optimal tradeoff.

Corresponding to the two segments of the tradeoff curve, there are two regions in the $P_e(R, \text{SNR})$ plot. The boundary $P_e = \text{SNR}^{-4/7}$, marked with a dashed line, corresponds to the knee in the tradeoff. The slopes of and the gaps between the curves in each region are labeled in the figure. They agree with the intercepts of the two tradeoff segments with the axes. Compared to Fig. 1.3(b) for the optimal tradeoff, the slopes in the upper region are less and the gaps in the lower region are wider. These deficiencies lead to increasing performance gaps as SNR increases.

1.3.4 Achieving the Entire Frontier Using Rotation-Based Codes

In this section, we identify code properties that enable the full efficient frontier of diversity-multiplexing tradeoffs to be achieved. Then, as an illustra-

tion, we construct a particular *tilted-QAM* code that meet these conditions and examine its performance characteristics more generally to develop additional insights.

Tilted-QAM codes are code designs based on applying unitary transformations to multidimensional QAM constellations. Such designs have a rich history, going back at least as far as [Lang, 1963]. In the more recent wireless literature, such codes were used in, e.g., [Boulle and Belfiore, 1992; Wornell, 1995; Boutros et al., 1996; Giraud et al., 1997; Bayer-Fluckiger et al., 2004] for single-antenna communication over multiple fades. More recently still they have been considered as space-time code candidates for multiple-antenna communication within a single fade. For example, in [Sethuraman et al., 2003; Sharma and Papadias, 2004] such codes are used to obtain the maximum diversity subject to a multiplexing gain constraint of unity. As another set of examples, [Damen et al., 2002; Ma and Giannakis, 2003] focus on using such codes to achieve the two end-points of the diversity-multiplexing frontier. Finally, [Yao and Wornell, 2003; Dayal and Varanasi, 2003; Belfiore et al., 2004; Elia et al., 2004] develop the role of such codes in achieving the entire tradeoff frontier, as is of interest in this section, particularly in the two-transmit two-receive antenna case.

We begin with a natural sufficient condition for achieving the entire frontier. This condition is encapsulated in the following theorem [Yao, 2003]:

Theorem 1 *For a system with two transmit and at least two receive antennas and code-length $T \geq 2$, consider a family of codebooks indexed by rate R that is scaled such that the peak and average codeword energy grow with R as $E_s \doteq \max_{\mathbf{x}} \|\mathbf{x}\|^2 \doteq 2^{R/2}$. Then a sufficient condition for achieving the diversity-multiplexing frontier is*

$$\min_{\mathbf{x}_1 \neq \mathbf{x}_2} |\det(\mathbf{X}_1 - \mathbf{X}_2)| \stackrel{\cdot}{\geq} 1. \quad (1.24)$$

Eq. (1.24) means that either the worst-case codeword-difference determinants do not decay to zero with rate or decay at most sub-exponentially.

As we now develop, a tilted-QAM code can indeed achieve a constant worst-case determinant. A construction is as follows. Given a transmission rate $R = r \log_2(\text{SNR})$, a constellation of size $M^2 = 2^{R/2} = \text{SNR}^{r/2}$ is carved from $\mathbb{Z} + \mathbb{Z}j$. Four information symbols are chosen from this constellation

and encoded into a codeword matrix $\mathbf{X} = \begin{bmatrix} x_{11} & x_{12} \\ x_{21} & x_{22} \end{bmatrix}$ via two rotations,

$$\begin{aligned} \begin{bmatrix} x_{11} \\ x_{22} \end{bmatrix} &= \begin{bmatrix} \cos(\theta_1) & -\sin(\theta_1) \\ \sin(\theta_1) & \cos(\theta_1) \end{bmatrix} \begin{bmatrix} s_{11} \\ s_{22} \end{bmatrix}, \\ \begin{bmatrix} x_{21} \\ x_{12} \end{bmatrix} &= \begin{bmatrix} \cos(\theta_2) & -\sin(\theta_2) \\ \sin(\theta_2) & \cos(\theta_2) \end{bmatrix} \begin{bmatrix} s_{21} \\ s_{12} \end{bmatrix}, \end{aligned} \tag{1.25}$$

where $(\theta_1, \theta_2) = (\frac{1}{2} \arctan(\frac{1}{2}), \frac{1}{2} \arctan(2))$. Like OSTBC, each information symbol s_{ij} effectively appears in both rows and columns of the codeword matrix \mathbf{X} , which is essential for achieving the full diversity gain. Unlike OSTBC, with rotation instead of repetition, two information symbols are transmitted per channel use, so there is no sacrifice of multiplexing gain. It is shown in [Yao, 2003] that this tilted-QAM code satisfies Theorem 1. In particular, the worst-case determinant is $1/(2\sqrt{5})$ for arbitrarily large rates.

To illustrate some intuition behind why this rotation leads to a constant worst-case determinant, let us set $s_{12} = s_{21} = 0$ and limit s_{11} and s_{22} to real numbers. In this case, $\det(\mathbf{X}) = x_{11}x_{22}$. The rotation of s_{11} and s_{22} to obtain x_{11} and x_{22} is shown in Fig. 1.11. We first note that since $\sin(\theta_1)$ and $\cos(\theta_1)$ are irrational numbers, all points except the origin are kept off the x_{11} and x_{22} axes. So the determinant is always non-zero. While a formal proof of the lower bound on the worst-case determinant appears in [Yao, 2003], the intuition comes from focusing on a series of points circled in Fig. 1.11. We see that while the points get closer and closer to the x_{11} axis, they also get further and further away from the x_{22} axis. As the result, the product $|x_{11}x_{22}|$ actually remains constant. More generally, when all four symbols s_{ij} are complex integers, the worst-case determinant is still lower bounded.

Numerical simulations with this tilted-QAM encoding and maximum likelihood (ML) decoding, implemented using sphere decoding, are performed for rates $R = 4, 8, 12, \dots, 32$ b/s/Hz. The resulting family of block error rate curves are plotted in Fig. 1.12 with dashed lines. The channel outage probability curves for those rates are plotted with thin solid lines for comparison. We see that the tilted-QAM block error rate curves match the channel outage probability curves closely, especially at higher rates. Since diversity-multiplexing tradeoff is a high SNR characteristic, it is possible for two systems with the same tradeoff to have different low SNR performances. The slopes of and the horizontal gaps between the tilted-QAM block error rate curves above and below the $P_{\text{out}} = \text{SNR}^{-1}$ line are labeled in Fig. 1.12. They agree with the slopes and gaps labeled in Fig. 1.3(b), which




Fig. 1.11. Rotate (s_{11}, s_{22}) to obtain (x_{11}, x_{22}) . All points except the origin are off the x_{11} and x_{22} axes. Points circled have constant $|x_{11}x_{22}|$ values.

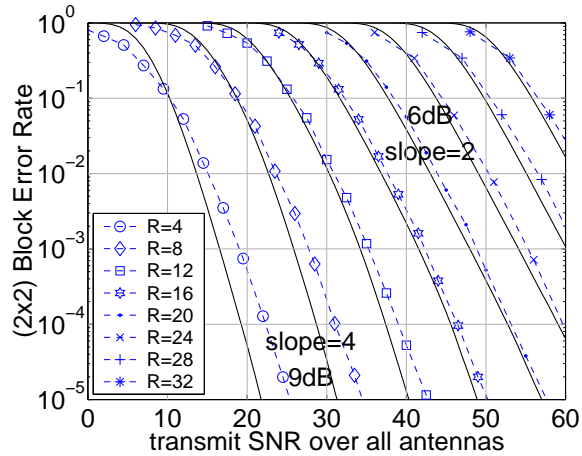


Fig. 1.12. Comparison of the titled-QAM block error rate curves (dashed) and the channel outage probability curves (thin solid) as functions of SNR for rates 4, 8, 12, \dots 32 b/s/Hz. They match quite well.

corresponds to the optimal diversity-multiplexing tradeoff. This shows that the tilted-QAM encoding defined in (1.25) together with ML decoding can achieve the entire diversity-multiplexing frontier.

1.4 Summary

The focus of this chapter was on an intuitive development of the diversity-multiplexing tradeoff inherent in the use of multiple-element antenna arrays for wireless links. The tradeoff is depicted and interpreted in a variety of ways that can be used by system designers in the engineering of communication links, and can guide their selection of various classes of coding and signal processing algorithms for use in conjunction with such arrays.

References

- S. M. Alamouti. A simple transmit diversity technique for wireless communications. *IEEE Journal on Selected Areas in Communications*, 16(8):1451–1458, October 1998.
- E. Bayer-Fluckiger, F. Oggier, and E. Viterbo. New algebraic constructions of rotated \mathbf{Z}^n -lattice constellations for the Rayleigh fading channel. *IEEE Transactions on Information Theory*, 50(4):702–714, April 2004.
- J.-C. Belfiore, G. Rekaya, and E. Viterbo. The golden code: a 2x2 full rate space-time code with non-vanishing determinants. *Proc. International IEEE Symposium on Information Theory*, page 308, June 2004.
- K. Boulle and J.-C. Belfiore. Modulation schemes designed for the Rayleigh fading channel. *Proc. Conf. Information Science and Systems*, pages 288–293, March 1992.
- J. Boutros, E. Viterbo, C. Rastello, and J.-C. Belfiore. Good lattice constellations for both Rayleigh and Gaussian channels. *IEEE Transactions on Information Theory*, 42(2):502–518, March 1996.
- T. M. Cover and J. A. Thomas. *Elements of information theory, Wiley Series in Telecommunications*. Wiley, New York, 1991.
- M. O. Damen, A. Tewfik, and J.-C. Belfiore. A construction of a space-time code based on number theory. *IEEE Transactions on Information Theory*, 48(3):753–760, March 2002.
- P. Dayal and M. Varanasi. An optimal two transmit antenna space-time code and its stacked extension. *Proc. Asilomar Conference on Signals, Systems and Computers*, November 2003.
- P. Elia, K. R. Kumar, S. A. Pawar, P. V. Kumar, and H. F. Lu. Explicit construction of space-time block codes: Achieving the diversity-multiplexing gain tradeoff. *submitted to IEEE Transactions on Information Theory*, September 2004.
- X. Giraud, E. Boutillon, and J. C. Belfiore. Algebraic tools to build modulation schemes for fading channels. *IEEE Transactions on Information Theory*, 43(3):938–952, May 1997.
- G. R. Lang. Rotational transformation of signals. *IEEE Transactions on Information Theory*, 9(3):191–198, July 1963.
- X. Ma and G. B. Giannakis. Full-diversity full-rate complex-field space-time coding. *IEEE Transactions on Signal Processing*, pages 2917–2930, November 2003.
- B. A. Sethuraman, B. S. Rajan, and V. Shashidhar. Full-diversity, high rate space-time block codes from division algebras. *IEEE Transactions on Information Theory*, 49:2596–2616, October 2003.
- N. Sharma and C. B. Papadias. Full-rate full-diversity linear quasi-orthogonal space-time codes for any number of transmit antennas. *EURASIP Journal on*

- Applied Signal Processing (special issue on Advances in Smart Antennas)*, 9: 1246–1256, August 2004.
- V. Tarokh, N. Seshadri, and A. R. Calderbank. Space-time codes for high data rate wireless communication: performance criterion and code construction. *IEEE Transactions on Information Theory*, 44(2):744–765, March 1998.
- E. Telatar. Capacity of multi-antenna Gaussian channels. *AT&T Bell Labs Internal Tech. Memo.*, June 1995.
- G. W. Wornell. Spread-signature CDMA: Efficient multiuser communication in the presence of fading. *IEEE Transactions on Information Theory*, 41(5):1418–1438, September 1995.
- H. Yao. *Efficient Signal, Code, and Receiver Designs for MIMO Communication Systems*. PhD thesis, Massachusetts Institute of Technology, June 2003.
- H. Yao and G. W. Wornell. Achieving the full MIMO diversity-multiplexing frontier with rotation based space-time codes. *Proc. of Allerton Conference on Communication, Control, and Computing*, October 2003.
- L. Zheng and D. N.C. Tse. Diversity and multiplexing: a fundamental tradeoff in multiple antenna channels. *IEEE Transactions on Information Theory*, 49(5): 1073–1096, May 2003.