802.11 with Multiple Antennas for Dummies

Daniel Halperin*, Wenjun Hu*, Anmol Sheth†, and David Wetherall*†
University of Washington* and Intel Research Seattle†

ABSTRACT

The use of multiple antennas and MIMO techniques based on them is the key feature of 802.11n equipment that sets it apart from earlier 802.11a/g equipment. It is responsible for superior performance, reliability and range. In this tutorial, we provide a brief introduction to multiple antenna techniques. We describe the two main classes of those techniques, spatial diversity and spatial multiplexing. To ground our discussion, we explain how they work in 802.11n NICs in practice.

1. INTRODUCTION

The use of multiple antennas at the receiver and transmitter has revolutionized wireless communications over the past decade. It has long been known that multiple receive antennas can improve reception through the selection of the stronger signal or combination of individual signals at a receiver. In the mid 1990s, however, seminal research by Foschini, Gans [1] and Telatar [6] predicted large performance gains from using multiple antennas at *both* transmitter and receiver. This kind of system is called a MIMO (Multiple-Input Multiple-Output) system in contrast with a SISO (Single-Input Single-Output) system that uses one transmit antenna and one receive antenna. SIMO and MISO systems also exist, as we will see shortly.

The excitement around MIMO is that, for richly scattered wireless environments such as an indoor 2.4 GHz 802.11 LAN, the multiple antenna pairs can provide independent paths between the transmitter and receiver. This *spatial degree of freedom* changes the fundamental relationship between power and capacity per second per Hz. Shannon capacity increases by up to one bit/sec/Hz for every doubling of power. With N antennas at each end, however, capacity increases by up to N bits/sec/Hz for every doubling of power. That is, *simply adding antennas has the potential to linearly scale the capacity even though the antennas transmit and receive on the same frequency band at the same time.* This is a key result in the quest for speed in modern wireless systems, since available spectrum is scarce and added power yields diminishing returns. Over the past decade, MIMO techniques have proved that they can deliver this value in practice. Today most high-rate wireless systems use MIMO technologies, including 802.11n, 4G mobile phone technology under the name LTE, and WiMAX.

Our aim in this note is to introduce multiple antennas as they are used in 802.11n wireless LANs to networking researchers with little previous knowledge of wireless communications. We choose 802.11n to ground the discussion in a relevant technology, but most of our discussion applies broadly to MIMO wireless systems. 802.11n is an extension of the earlier 802.11a/g standard that adds the use of multiple antenna techniques at the physical layer. Strictly speaking, the 802.11n standard is in draft form, but the physical layer details have been finalized for years. Draft 802.11n hardware has been commercially available since 2007 and now ships standard in many laptops.

The way 802.11n uses multiple antennas is quite different than earlier 802.11a/g access points (APs) that have multiple antennas sticking out of the box. These APs typically choose the best antenna but still uses a single antenna at a given moment. In terms of wireless signal processing, they are still SISO systems. With 802.11n, multiple antennas at the transmitter and/or receiver are used at the same time (and on the same frequency band). For this to happen, the transmitters and receivers must have multiple RF processing chains to go with the multiple antennas; it is not only antennas that count. This is the hallmark of a MIMO system.

There are two basic classes of multiple antenna techniques that are described in textbooks and used in 802.11n. *Spatial diversity* techniques increase reliability and range by sending or receiving *redundant* streams of information in parallel along the different spatial paths between transmit and receive antennas. This helps with reliability because it is unlikely that all of the paths will be degraded at a given moment. Improved range, and some performance increase too, comes from the use of multiple antennas to gather a larger amount of signal at the receiver. *Spatial multiplexing* techniques increase performance by sending *independent* streams of information in parallel along the different spatial paths between transmit and receive antennas. This improves performance because, if we take care in how we construct and decode signals, adding an antenna and independent stream of information need not slow down the streams that are already being sent.

We describe basic techniques for both these classes that are compatible with 802.11n and used in commercial NICs to the best of our knowledge. The 802.11n standard does not give any of the techniques per se because, as a standard, it is concerned with interoperability rather than implementation. It also contains rather a lot of options and we have focused on those options that are most commonly used today.

To put the role of multiple antennas in 802.11n in context, consider that the highest data rate in 802.11a/g is 54 Mbps and the highest data rate in 802.11n is 600 Mbps. This is an increase of a factor of 11. Of this, a factor of four comes from the use of four antennas. This is the bulk of the increase and easily the largest single factor. Another factor of two comes from simply using double width channels of 40 MHz instead of 20 MHz. The remaining factor of around 1.4 comes from tweaking the OFDM and coding constants to shave overhead. In practice, many devices may not have four antennas. Up to three antennas are commonly supported by NICs, and it is expected that clients will tend to have fewer antennas for space and power reasons, while APs will tend to have more antennas for performance reasons.

The rest of this tutorial is organized as follows. We begin with a quick discussion of an 802.11 wireless link in the single antenna case. Here, fading wireless channels are the key difficulty that the physical layer overcomes through the use of diversity techniques. We then describe how spatial diversity schemes add to the picture, from the simple selection of antennas as can be done in a SISO system to combining that requires a SIMO (or MISO) system. Next, we describe spatial multiplexing schemes, from simple direct-mapped MIMO to the use of pre-coding to extract larger gains in practice. We conclude with pointers to more advanced techniques and other introductory material for the interested reader.

2. WIRELESS CHANNELS & SISO 802.11

We begin with background on indoor wireless channels at 2.4 GHz and 5 GHz, and how single antenna 802.11 systems send information over these channels at the physical layer.

2.1 Faded Wireless Channels

In wireless communications, the performance of a link is fundamentally determined by the Signal-to-Noise Ratio (SNR), which measures the received signal strength of a transmission relative to the thermal noise in the receiver hardware that distorts the received signal. Over a typical 802.11a link today, packets are transmitted with 50 mW of power, and for a strong link the received power might be as high as 50 nW, a million-fold loss (60 dB) of power. This received signal is still much greater than the noise floor, which for a 20 MHz 802.11 channel is 0.1 pW. Thus the high SNR ($10 \log_{10}(50/0.1) \approx 27$ dB) supports a fast bit rate.

The *attenuation* of the signal between transmitter and receiver comes from several effects. One effect is *path loss* as the radiated signal spreads out over a wider area and passes through different materials such as walls. Path loss causes the power to drop off at least as fast as the square of the distance traveled. Other *fading* effects cause the signal to be weakened beyond the path loss. For example, shadowing is the degradation of the signal due to large obstacles such as buildings that lie in the path. This causes *slow fading* in which the signal strength varies slowly over time as the receiver moves or the environment changes.

The most problematic kind of fading for 802.11 is due to *multi-path*. At 2.4 GHz and 5 GHz, RF signals bounce off metal and glass surfaces that are common indoors. This *scattering* leads to a situation in which many copies of the signal arrive at the receiver having traveled along many different paths. When these copies combine they may add constructively, giving a good overall signal, or destructively, mostly canceling the overall signal, all depending on the relative phases of the portions. Measurement studies of fading report signal variations as high as 15-20 dB [3].

Worse yet, small changes in path lengths can alter the situation from good to bad because the wavelength at 2.4 GHz and 5 GHz (over which the RF signals go through a complete phase) is only 12 cm and 6 cm, respectively. Statistical models tell us that multi-path fading effects are independent for locations separated by as little as half a wavelength. This means that multi-path causes rapid signal changes or *fast fading* as the receiver moves, or in the case of a stationary node the surrounding environment changes. Because multi-path effects depend on the phases of signals, they are strongly *frequency selective*. This means that some unlucky frequencies in a 20 MHz 802.11 channel may be wiped out while others are unaffected. We will see an example in the next section.

The net effect of multi-path fading is that the received wireless signal can vary significantly over time, frequency and space. This is a problem for good performance because at any given time there is a significant probability of a deep fade that will reduce the SNR of the channel below the level needed for a given communication scheme.

2.2 Single Antenna 802.11 OFDM

The main technique used in wireless systems such as 802.11 to cope with variable wireless channels is *diversity*. Diversity is the spreading of information with some redundancy across multiple independently faded channels. When this is done, it is unlikely that a deep fade on a single channel will prevent successful communication. The trick, however, is to find independently faded channels. These exist within the physical layer and come from harnessing the time, frequency and spatial resources of the wireless link.

The 802.11a/g/n physical layer is based on OFDM (Orthogonal Frequency Division Multiplexing). This technique divides the relatively wideband 20 MHz 802.11 channel into 64 subcarriers of 312.5 kHz each, such that each subcarrier can be thought of as its own narrowband channel. It is completely different than the spread spectrum technique used in older 802.11b equipment. There are many variations on OFDM, but in 802.11 data is sent on the subcarriers using the same modulation and transmit power for each subcarrier. This modulation ranges from BPSK, to QPSK, QAM-16 and QAM-64, with each higher modulation sending more bits per symbol and being used when there is a higher SNR. There are minor differences between 802.11a/g and 802.11n. In 802.11a/g there are 48 data subcarriers, 4 pilot tones for control, and 6 unused guard subcarriers at each edge of the channel. In 802.11n, there are only 2 guard subcarriers at each edge of the channel, and two adjacent 20 MHz channels can be used as a single 40 MHz channel.

The beauty of OFDM is that it divides the channel in a way that is both computationally and spectrally efficient. High aggregate data rates can be achieved, while the encoding and decoding on different subcarriers can use shared hardware components. More relevant to our point here, however, is that OFDM transforms a single large channel into many relatively independently faded channels. This is because multi-path fading is frequency selective, so the different subcarriers will experience different fades. Some adjacent subcarriers may be faded in a similar way, but the fading for more distant subcarriers is often uncorrelated. Dividing the channel also increases the symbol time per channel, since many slow symbols will be sent in parallel instead of many fast symbols in sequence. This adds time diversity because the channel is more likely to average out fades over a longer period of time.

802.11 makes use of the diversity provided by OFDM by

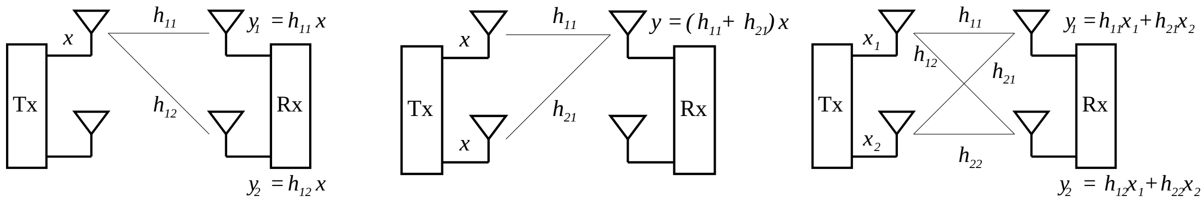

(a) Receive diversity (b) Transmit diversity (c) Spatial multiplexing

Figure 1: Using some of the transmit/receive antennas in an example 2x2 system to exploit diversity and multiplexing gain. x_i and y_i represent transmitted and received signals. h_{ij} is the channel gain between the ith transmit antenna and the jth receive antenna, indicating a signal's amplitude attenuation and phase shift as it traverses the channel.

coding across the data carried on the subcarriers. This uses a fraction of them for redundant information. It can later be used to correct errors that occur when fading reduces the SNR on some of the subcarriers. First, a convolutional code of rate 1/2 adds redundant information. It is then punctured by removing bits as needed to support coding rates of 2/3 and 3/4, plus 5/6 for 802.11n. At a rate of 3/4, for example, a quarter of the data on the subcarriers is redundant. An alternative LDPC (Low-Density Parity-Check) code with slightly better performance can also be used for 802.11n.

The net effect of OFDM plus coding is to provide consistently good 802.11 performance despite significant variability in the wireless signal due to multi-path fading.

3. SPATIAL DIVERSITY

In this section we look at spatial diversity techniques that can be applied at the receiver and at the transmitter. Adding multiple antennas to an 802.11n receiver or transmitter provides a new set of independently faded paths, even if the antennas are separated by only a few centimeters. This adds spatial diversity to the system, which can be exploited to improve resilience to fades. There is also a power gain from multiple receive antennas because, everything else being equal, two receive antennas will receive twice the signal. This increases performance at a given distance, and hence range.

3.1 Receiver diversity techniques

Consider the arrangement in Figure 1(a). One transmit antenna at a node is sending to two receive antennas at a second node. This is known as a 1x2 system. Real systems may have more than two receive antennas, but two will suffice for our explanation. With this setup, each receive antenna receives a copy of the transmitted signal modified by the channel between the transmitter and itself. Note that the channel differs for each subcarrier (because they are independently faded) as well as for each antenna. The question now is how to combine the two received signals to make best use of them.

We consider two combination methods to show the extremes. The simplest method is to select the antenna that has the strongest signal, hence the largest SNR, to receive the packet and ignore the others. We will call this method SEL, for selection combining. This is essentially what is done by 802.11a/g APs with multiple antennas. It helps with reliability, because both signals are unlikely to be bad, but it

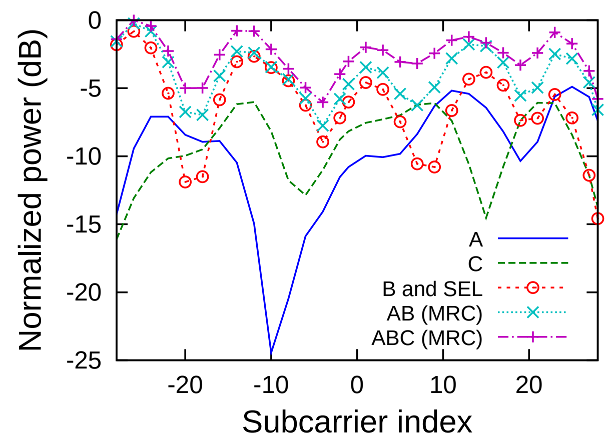

Figure 2: Frequency-selective fading over testbed links: the figure shows, for an example link, the received power on each subcarrier for individual antennas and under SEL and MRC, normalized to the strongest subcarrier power.

wastes perfectly good received power at the antenna that is not chosen.

The better method is to add the signals from the two antennas together. However, this cannot be done by simply superimposing their signals, or we will have just recreated the effects of multi-path fading. Rather, the subcarriers of the second signal should each be delayed until they are in-phase with the corresponding subcarriers of the first signal. Then, the power in the signals will add for each subcarrier. To do this, the receiver needs an estimate of the channel gains for each subcarrier. This is obtained by measuring training fields sent in the preamble. The receiver also needs a dedicated RF chain for each antenna to process the signals. This increases the hardware complexity and power consumption, but yields better performance.

As a twist in the above, the subcarrier signals are also weighted by their SNR. This gives less weight to a signal that has a larger fraction of noise, so that the effects of the noise are not amplified. The result is Maximal Ratio Combining, or MRC. It is known to be optimal and produce an SNR that is the sum of the component SNRs.

As an example of how MRC and SEL work for 802.11, consider Figure 2. This figure shows the strength of wireless signals received for each subcarrier of a 20 MHz channel at three antennas for one of the links in our indoor testbed. The subcarrier strengths are measured in decibels normalized to the strongest subcarrier strength. This gives a much more detailed view than metrics such as the RSSI (Received Signal Strength Indication) for a link, which gives only the sum of the signal strength over all subcarriers. For each antenna, A, B and C, the signal varies over the channel, changing slowly from one subcarrier to the next. It shows some deep fades due to (frequency-selective) multipath, particularly at antenna A which sees at least 20 dB (100x) of variation in subcarrier strengths. These deep fades will cause errors, since 802.11 uses the same modulation technique for all subcarriers. Coding across the subcarriers will have to repair these errors for successful reception.

Figure 2 also shows how selection (SEL) and MRC work on the received signals. SEL picks the antenna with the strongest overall signal, which is antenna B in this case. However, this signal will still vary over its subcarriers due to multi-path fading. In our example, SEL reduces fades to 15 dB (from > 20 dB without it). This means that SEL can avoid unlucky antennas that have pronounced fades, but does little to improve on antennas that already have reasonable signals.

In contrast, MRC adds the signals (weighted by their SNR) for each subcarrier. This produces the top line on the figure that is better than the individual signals at every point and significantly flatter over the channel. Now, the fading has been reduced to roughly 5 dB. This in turn means that coding will have to deal with fewer and less pronounced errors, which allows higher coding rates or higher modulation rates. MRC can produce significant diversity gains in practice that exceed the gains of antenna selection. Though receiver processing algorithms is not specified by the 802.11n standard, MRC is closely tied to MIMO signal decoding and is likely to be available in any 802.11n NIC.

3.2 Transmit Diversity Techniques

The receiver diversity techniques we have looked at use a single transmit and multiple receive antennas. There are also transmit-side equivalents of both SEL and MRC that use multiple transmit and single receive antennas. A 2x1 setup is shown in part (b) of Figure 1. This can be useful when the AP has more antennas than the client, so that it can use its multiple antennas to benefit a single antenna client.

The transmit-side equivalent of SEL is simply to select the single best antenna on which to transmit a packet; we do not consider this further. The transmit-side equivalent of MRC is a kind of *transmit beamforming* in which the transmitter *precodes* the signals that are sent out the transmit antennas. The signals then combine in the desired way as they pass over the wireless channel. To achieve the equivalent of MRC, the transmitter must delay each subcarrier signal sent by the second antenna so that it will have the same phase as that of the first antenna after both signals travel over the channel. The signals will then add rather than cancel each other. To complete the picture, the signals are further weighted at the transmit side with their relative SNRs at the receiver (since this cannot be done at the receiver).

The disadvantage of transmit diversity compared to receive diversity is that the transmitter must know the channel gains to know how to precode the signals. These channel gains are measured at the receiver (during the preamble) as part of its normal operation, but they are not normally known to the transmitter. In 802.11n, there is a channel state feedback packet that the receiver can use to send channel gains to the transmitter. Alternatively, since the properties of RF channels are reciprocal, the transmitter can learn the channel gains when it in turn receives a packet from the target receiver. In practice, calibration is needed to account for the differing properties of the NICs at each end. In both cases, regular updates are needed because the channel state changes over time, often very quickly due to multipath fading, and out-of-date channel gains make precoding less effective.

It is also worth noting that there are different beamforming techniques that use phased antenna arrays to direct the signal. These techniques are based on precise geometric antenna arrangements (circles or lines) and orient the signal in physical space with the same pattern for each subcarrier. The measurement-based beamforming described above has no particular physical interpretation and treats each subcarrier individually.

4. SPATIAL MULTIPLEXING

The real excitement around MIMO is that the independent paths between multiple antennas can be used to much greater effect than simply for diversity to boost the SNR. *Spatial multiplexing* takes advantage of the extra degrees of freedom provided by the independent spatial paths to send independent streams of information at the same time over the same frequencies. The streams will become combined as they pass across the channel, and the task at the receiver is to separate and decode them.

To get an idea of the potential benefits, we turn briefly to theory. For a single antenna at the transmitter and receiver, Shannon's classic formula gives the capacity as $C = B\log(1+\rho)$. Here, C is the capacity in bits/sec, B is the system bandwidth, and ρ is the SNR of the channel. Now consider the case where each node has N antennas and there are independent spatial paths between the pairs of transmit and receive antennas. There are N spatial degrees of freedom in the system, since the signal from each transmit antenna can change the received signals in a different manner. By using the antennas to divide the transmit power over these degrees of freedom, the transmitter can send N streams of data, each getting an SNR of ρ/N. This is a rough argument for the theoretical capacity for a MIMO system [1]:

$$C = BN\log\left(1+\frac{\rho}{N}\right)$$

At high SNR, this capacity scales nearly linearly with the number of antennas, even for a small number of antennas. That is a much larger performance improvement than simply sending a single stream at the aggregate SNR over all the receive antennas. At low SNR, however, the gain from receive antennas is the larger effect, with extra transmit antennas making little difference.

There are many ways to process signals at the transmitter and receiver to realize MIMO gains that have different tradeoffs. We will look at a basic MIMO scheme that is easy to implement in practice, and an improved scheme that comes closer to the MIMO capacity.

4.1 Direct Mapped MIMO

The simplest way to get spatial multiplexing benefits is to transmit multiple packets (or *spatial streams*) directly out each antenna. We will call this direct mapped MIMO, and Figure 1(c) shows a 2x2 system. For each subcarrier, x_i denotes the signal sent on each transmit antenna, y_j the signal received at each receive antenna, and h_{ij} the channel gain (i.e., attenuation and phase shift) between the ith transmit and jth receive antenna. Expressing terms in matrix form, we have $\hat{y} = H \cdot \hat{x} + \hat{n}$, where $\hat{n}$ is the vector of noise terms on each receive chain. For a fair comparison with SISO systems, it is conventional to fix the aggregate transmit power across all antennas.

The problem is how to decode the multiple streams at the receiver. One way to think about the MIMO system is that the signal received at each antenna is a linear combination (due to superimposition on the channel) of the transmitted signals. Because multi-path leads to varying channel gains, the linear combinations between different transmit-receive antenna pairs will be independent of one another with high probability. This is a linear system of equations. Each of the channel gains, h_{ij}, is known at the receiver because it is measured by the receiver during the preamble. The y_j are measured as received signals too; only the transmitted packets, the x_i, are unknown. The receiver can recover the x_i with processing that solves a system of linear equations. It can simply invert H and multiply it by $\hat{y}$. This will work as long as there are enough independent equations (from antennas) for the unknowns (from packets). It is not necessarily the case that H is invertible, but with independent fading between transmit-receive antenna pairs, H is invertible with high probability.

The above receiver processing is a simple way to receive multiple signals and is called Zero-Forcing (ZF). Geometrically, it is equivalent to recovering each stream by projecting the vector of receive signals in a direction that is orthogonal to the channel gains of the other, unwanted streams. That is, ZF recovers a signal by nulling the interfering signals, forcing them to zero. The difficulty with ZF, however, is that it lowers the SNR when the channel gains of the different streams are correlated. This is because more correlated channels cause more of the wanted signal energy to be lost during receiver processing. Since the noise is not reduced, the SNR falls. More sophisticated receiver processing can do a better job, for example by trying all combinations of transmitted signals and picking the best fit, but this quickly becomes computationally intensive.

As an example, we examine a direct mapped 3x3 MIMO link in our testbed with a Zero-Forcing (ZF) receiver. Figure 3 shows the SNR for each subcarrier for each of the three streams on this link after they are decoupled by ZF. The SNRs are normalized with respect to the aggregate SNR of the received signals on the strongest subcarrier. As before, there is a significant amount of variation across subcarriers, as much as 16 dB in this example. This variation comes from multi-path fading plus the different correlations of channel gains that vary the effectiveness of the ZF receiver. By comparing the sum SNR across the three streams to the aggregate SNR of the received signals, we can get a rough sense of the SNR that is lost by the ZF receiver.

In 802.11n with multiple spatial streams, the above processing happens for each subcarrier. The data that is sent is coded across the subcarriers of each stream using coding, as before, to provide resilience to some faded subcarriers.

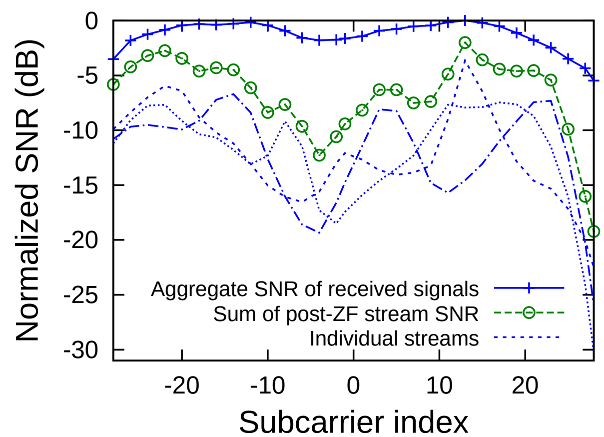

Figure 3: Subcarrier strength distribution for an example 3x3 link: The top two lines indicate the aggregate SNR of the received signals and the amount of SNR being utilized by the ZF receiver; the lower three lines show the SNR for each stream.

There is also a choice of what fraction of the total power to send out each antenna, and what modulation rate to use for each stream. Direct mapped MIMO is typically used in a setting in which the transmitter does not know the channel gains. Lacking this information, it makes sense to divide the power evenly across antennas and to modulate each stream at the same rate. This will not in general give the highest throughput because some streams may have better channels than others. Overall, expect equal power, equal rate, direct-mapped MIMO streams might deliver roughly 70-80% of their scaling potential in moving from one to three antennas.

4.2 Pre-coded MIMO

Direct mapped MIMO wastes capacity when the power is not matched to the channel and wastes SNR when ZF receivers imperfectly untangling streams. We can do better. In much the same manner as transmit diversity, we can benefit from knowledge and work at the transmitter. In this case, the transmitter can use the channel to its full potential and the receiver can decode efficiently without wasting SNR. The downside of this strategy is that, as with transmit diversity, the transmitter must know the channel gains and track them as the channel changes.

A standard construction to use the MIMO channel in this manner is based on the *singular value decomposition (SVD)* of the channel matrix H. From linear algebra, any matrix H can be factored into the form $H = U\Sigma V^H$, where U and V are unitary matrices, Σ is a diagonal matrix of singular values σ_i and H^H indicates the Hermitian or conjugate transpose operation of H. Now $\hat{y} = H \cdot \hat{x} + \hat{n}$ from the previous section becomes $\hat{y} = U\Sigma V^H \cdot \hat{x} + \hat{n}$. This is significant because it suggests that the channel consists of orthogonal paths (the diagonal matrix) but only when viewed in signal spaces that are rotations of the signal coordinates (the unitary matrices) at the transmitter and receiver.

We can access the orthogonal paths at the transmitter by *precoding* the signal by V and at the receiver by *shaping* the signal by U^H. If we let $\tilde{y} = U^H \cdot \hat{y}$, $\hat{x} = V \cdot \tilde{x}$, and $\tilde{n} = U^H \cdot \hat{n}$,

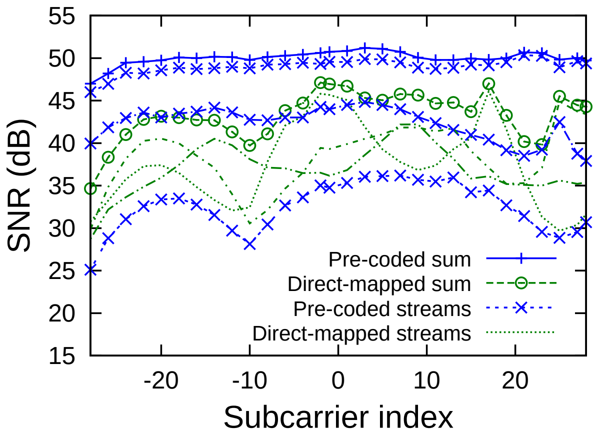

Figure 4: Distribution of per-stream SNR across subcarriers for an example link, using pre-coded vs direct mapped multiplexing. Also shown are the sum SNR across streams for the two methods. The relative strength between the pre-coded streams differs by 14-20 dB on many subcarriers. This suggests the link has correlated spatial paths that are not used effectively by direct mapped MIMO.

and rewrite the channel, then we have:

$$\tilde{y} = U^H \hat{y} = U^H U \Sigma V V^H \tilde{x} + U^H \hat{n} = \Sigma \tilde{x} + \tilde{n}$$

Physically, $\hat{x}$ and $\hat{y}$ are still the actual signals transmitted and received. However, when viewed in terms of $\tilde{x}$ and $\tilde{y}$, the effective channel between them is simply Σ. Since Σ is diagonal, the streams do not interfere at the receiver and are decoupled to the simple form of $y_i = \sigma i x_i + n_i$. Each of these signals can be independently and easily decoded and the original signals retrieved. U and V being unitary, the total power of the original signals, received signals or noise remains unchanged during precoding or shaping.

The singular values (the σ_i) give the capacity of each independent spatial path to carry information. They vary with the specifics of the channel. If the singular values are close to each other, the spatial paths have roughly equal capacities. Large multiplexing gains can then be obtained. If, on the other hand, the singular values differ markedly, then some of the spatial paths have relatively low capacity. This can happen when some of the paths are significantly correlated. It is often the case with line-of-sight links for which multiple antennas see the same dominant signal. In such cases, it is better to direct a larger fraction of the overall power to the high capacity paths and a smaller fraction of power to the low capacity ones. A well-known algorithm called *water-filling* gives the transmit power allocations to maximize the throughput of multi-stream systems as a function of the capacity of the individual streams.

To see the difference between direct-mapped and pre-coded multiplexing, we examine another testbed link. For both methods, Figure 4 shows the subcarrier distribution of per-stream SNRs and the sum of the SNR across the streams. The channel matrices for this link have widely varying singular values on most subcarriers, suggesting correlated spatial paths. This leads to large SNR differences across subcarriers for direct mapped MIMO, more than typical with fading alone. The pre-coded MIMO streams compensate for this situation by putting most of the transmit power into the best path, at the expense of weakening the other two paths. The much greater sum SNR for pre-coded MIMO over direct mapped MIMO suggests that pre-coded MIMO improves the overall situation.

To use pre-coded MIMO in 802.11n, we also need to choose modulation rates. 802.11n provides some unequal modulation rates in which high-rate streams are a multiple of four faster than low-rate streams. These are likely important to gain the benefits of pre-coding in which different power is allocated to different streams. However, the unequal rates are unlikely to be a close fit to the waterfilling power allocations, so some compromise will be needed. Pre-coded MIMO has the potential improve performance over direct mapped MIMO, but it is not widely used in 802.11 NICs yet to the best of our knowledge given the added complexity.

5. NEXT STEPS

MIMO technologies are rapdily being adopted in 802.11 and other wireless systems, despite their complexity over SISO systems, because of the significant benefits they can deliver in practice. The techniques we described in this tutorial are the tip of the iceberg. Multiple antennas can also be used for combinations of diversity and multiplexing rather than one or the other. For example, a 3x3 MIMO system might send one, two or three streams, with the extra antennas used for diversity benefits. There is a fundamental tradeoff between the performance from using diversity and multiplexing, and the optimal combination depends on the SNR of the channel and the performance goals [8].

Other advanced topics include multi-user MIMO, in which a node with multiple antennas communicates with multiple users simultaneously to improve performance, and space-time coding, in which information is coded across multiple antennas as well as time. For more information, we refer the interested reader to deeper introductory text [5], textbooks on wireless communications and MIMO [7, 4, 2], and the references therein.

6. REFERENCES

[1] G. J. Foschini and M. J. Gans. On limits of wireless communications in a fading environment when using multiple antennas. *Wireless Personal Communications*, 6:311–335, 1998.

[2] A. Goldsmith. *Wireless Communications*. Cambridge University Press, 2005.

[3] G. Judd, X. Wang, and P. Steenkiste. Efficient channel-aware rate adaptation in dynamic environments. In *ACM MobiSys*, pages 118–131, 2008.

[4] C. Oestges and B. Clerckx. *MIMO Wireless Communications: From Real-World Propagation to Space-Time Code Design*. Academic Press, 2007.

[5] Y. Palaskas, A. Ravi, S. Pellerano, and S. Sandhu. Design Consideration for Integrated MIMO Radio Transceivers. In *Wireless Technologies: Circuits, Systems, and Devices*, chapter 4. Taylor & Francis, Inc, 2007.

[6] E. Telatar. Capacity of multi-antenna Gaussian channels. *European Trans. Telecommunications*, 10(6):585–595, 1999.

[7] D. Tse and P. Viswanath. *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.

[8] L. Zheng and D. Tse. Diversity and multiplexing: a fundamental tradeoff in multiple-antenna channels. *IEEE Trans. Information Theory*, 49(5):1073–1096, May 2003.