Two Antennas are Better than One: A Measurement Study of 802.11n

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

ABSTRACT

We present a measurement study of how well 802.11n works in practice. Our goal is to understand the performance impact of the multiple antenna (spatial) techniques that are the key new capability of 802.11n. We experiment with a 10 node indoor testbed equipped with commodity Intel 802.11n NICs that have three antennas, i.e., 3x3 MIMO, and low-level instrumentation to let us to measure RF channel conditions. We find 802.11n to be highly effective at improving link rates. Two or three antennas improves the median rate by 2× and 2.2×, respectively. This realizes most of the predicted MIMO gains. However, we find that spatial diversity is key to realize these gains beyond 2 streams, with the 2x3 configuration outperforming the 3x3 configuration for at least three-quarters of our links. By analyzing the RF channel conditions, we conclude that additional diversity produces especially large gains because it mitigates frequency-selective fading. We believe this is the first measurement study of 802.11n with commodity wireless NICs that relates high-level performance to low-level channel effects.

1. INTRODUCTION

To meet the demand for high performance local-area wireless, the 802.11n standard [13] defines the next generation of the popular 802.11a/g technology. Surpassing earlier technologies, it is based on a multiple antenna physical layer that supports multiple input multiple output (MIMO) techniques. These techniques have generated much excitement since the late 1990s [7, 8, 27] due to predictions of large capacity gains. By using MIMO, for instance, 802.11n defines rates of up to 600 Mbps. NICs that implement the 802.11n draft standard have become widely commercially available in the past year and are now shipped on most high-end laptops. We can expect many future wireless systems to be based on 802.11n technology.

However, while many papers report on the performance of 802.11a/b/g wireless systems [14, 15] and multiple radio based systems [17, 4] in practice, there is no comparable literature in the research community on the performance of 802.11n systems. There are many MIMO techniques and a wealth of accompanying theoretical literature on capacity [10], but this work is based on idealized assumptions (e.g., rich scattering) and mostly concerns upper bounds. Much experimental work focuses on measurements of RF channels [25, 18] or the design of MIMO-OFDM receivers [30]; fewer papers report on the performance of MIMO systems over real channels [23]. We are aware of only one recent short paper [24] that studies the performance of commodity 802.11n NICs in real settings.

In this paper, we investigate how well the 802.11n physical layer works in practice, and present the results of a measurement study on a 10 node indoor testbed. Each node has an Intel 802.11n NIC with three antennas and can support up to 3x3 MIMO. The NIC is further instrumented to measure detailed RF channel conditions, i.e., OFDM subcarrier channel gains. We focus on the base spatial diversity and spatial multiplexing mechanisms[1]. These features differentiate 802.11n from 802.11a/g. Spatial diversity combines the signals from multiple antennas to improve the reliability, translating to a throughput increase, or range of wireless transmissions. Spatial multiplexing sends different signals simultaneously using multiple antennas to increase aggregate throughput. A simple textbook view is that MIMO improves performance by a factor of N for N antennas per node, diversity is most valuable at low SNR, and multiplexing is most valuable at high SNR.

We find that 802.11n is highly effective at boosting performance even with the base 802.11n MIMO techniques that use the simplest equal power and equal rate transmit streams. When we make use of all diversity and multiplexing configurations up to 3x3, the median link rate in our testbed increases from 52 Mbps to 117 Mbps. This increase of 2.2× for 3x3 MIMO captures the bulk (≈74%) of the theoretical linear scaling with the number of node antennas, and 2x2 MIMO does even better with an increase of 2×. By analyzing detailed channel information we find that *diversity* is the key factor in these increases beyond two streams. For example, more than three-quarters of our links run faster in a 2x3 configuration (with two transmit and three receive antennas)

[1]There are many optional features in 802.11n, and commonly unimplemented. We leave them to future work.

than a 3x3 configuration. To support high rates, we find it more useful to add diversity than to add multiplexing unless the link is already near the highest 802.11n per-stream rate. From our channel level analysis, we identify the underlying reason for these large diversity gains as *frequency-selective fading*. Diversity reduces the median power gap between the strongest and weakest subcarriers in our testbed by 7 dB as we shift from one to three antennas. This flatter channel makes better use of the transmitted power, and also makes SNR a better indicator of single stream performance.

We hope that our study will be of interest to the research community in several respects. We characterize how multiple antennas stand to improve 802.11 performance in practical environments. We assess how well the simplest 802.11n mechanisms realize MIMO benefits and hence the need to explore more complicated, optional 802.11n mechanisms. While our present work concerns the physical layer, we expect it to inform studies of the MAC and other higher layer considerations, e.g., rate adaptation. Our study highlights that 802.11n has many implications for these topics.

The rest of this paper is organized as follows. In Sections 2 and 3 we detail the goals of our study and provide an overview of the relevant 802.11n MIMO techniques. We describe our indoor testbed and methodology in Section 4. Sections 5, 6 and 7 investigate spatial diversity, spatial multiplexing, and overall network effects. We then describe related work in Section 8 and conclude in Section 9.

2. PROBLEM AND GOALS

Our goal is to experimentally understand the performance of the 802.11n physical layer when it operates in indoor environments and with stationary clients, as it is a common setting for wireless LANs. This will depend on both the characteristics of real RF paths and the specific multiple antenna schemes used. We explore basic 802.11n spatial diversity and spatial multiplexing, as they are the two main ways the 802.11n physical layer is improved. Our primary performance is throughput, as measured by link rate, but we also consider coverage for a given rate and reduced power requirements (or, equivalently, extended range). We tackle the following questions:

- Given the 802.11 base of coded OFDM, how much do multiple antennas improve performance over single antenna systems? This is our proxy for assessing the gains of 802.11n over 802.11a/g given that they are otherwise similar.

- How does this improvement compare to the gains predicted in theory? We would like to know if 802.11n and our implementation of it are able to realize the potential gains of MIMO.

- For spatial diversity, how does sophisticated receive diversity techniques compare with the simple antenna selection? This will highlight if the added complexity of multiple RF receive chains in 802.11n is warranted.

- For spatial multiplexing, how much do per-antenna streams improve performance? This will indicate whether the simplest MIMO schemes are sufficient to obtain good gains.

- When is it better to use diversity and when is it better to use multiplexing to improve performance? We want a head-to-head comparison.

- What channel characteristics impact performance? Can we predict performance from channel measurements? We want to relate high-level performance to low-level channel conditions and vice versa.

The answers to these questions will shed light on expected 802.11n performance in the short term and lead to improved methods for operating 802.11n networks in the long term.

3. 802.11N MIMO OVERVIEW

As background, we first briefly introduce 802.11n multiple antenna techniques. Spatial diversity and multiplexing both leverage the effect that antennas which are separated by more than roughly half a wavelength (≈2.5 cm at 5.4 GHz) in *rich scattering* environments have channels whose subcarriers fade independently.

3.1 Coded OFDM

802.11n adds multiple antennas to a physical layer based on coded Orthogonal Frequency Division Multiplexing (OFDM); the single antenna (SISO) physical layer is otherwise quite similar to that of 802.11a/g. Each 20 MHz channel is divided into 52 data subcarriers, 4 more than in 802.11a/g, and 4 pilot subcarriers. Each data subcarrier is modulated equally using BPSK, QPSK, 16-QAM, or 64-QAM, according to the desired bit rate. By the nature of OFDM, each subcarrier is a separate narrowband channel, and MIMO techniques are applied separately to each subcarrier. Typically, different subcarriers between a given transmit-receive antenna pair will experience different amounts of multipath fading due to their different frequencies; the overall channel is then *frequency-selective*. To tolerate a small number of deeply faded subcarriers, the transmitted data is interleaved across subcarriers and coded for error correction at rates of 1/2, 2/3, 3/4 or 5/6, the first three of which also available in 802.11a/g.

3.2 Spatial Diversity

It is unlikely that all of the channels between transmit and receive antenna pairs will experience a deep fade on the same subcarrier simultaneously. The gain obtained from using multiple separately faded channels is called *diversity gain*. There is also an added benefit for multiple receive antennas when multiple copies of the signal are received. This provides a power or *array gain*.

SEL: Selection Combining. The simplest way to use multiple receive antennas is to select for each packet the antenna with the highest SNR and ignore the others; this is commonly implemented by picking the strongest receive signal. We refer to this scheme as SEL. The same antenna is selected for all of the subcarriers so that only a single antenna is used at a time. Most 802.11a/b/g NICs support antenna selection with two antennas, though it is often used on access points but not on clients.

MRC: Maximal Ratio Combining. A better way to harness the useful power from all antennas is to add the signals from different antennas in a coherent manner. To do so, the receiver needs an estimate of the channel gains for each subcarrier, which it obtains by using training fields in the preamble. *Maximal-ratio combining* (MRC) delays signals from different antennas so that they have the same phase, weights them proportionally to their SNR, and adds them. This requires multiple RF chains (as well as multiple antennas).

Intuitively, SEL uses the least faded signal, whereas MRC also uses signals in deeper fades to improve on the best signal. Though MRC is not part of the 802.11n standard, it is closely tied to MIMO signal decoding and is likely to be available in any 802.11n NIC. So far as we know, it is not implemented in 802.11a/b/g NICs.

Transmit Diversity. There are transmit-side equivalents of both SEL and MRC: a sending node can select the best antenna to transmit a packet; or with optional *transmit beamforming* a transmitter can send delayed and amplified copies of the signal out different antennas so that they combine coherently at the receiver. The disadvantage of transmit diversity is complexity, because the transmitter must have knowledge of the current channel, typically relying on receiver feedback.

Other. Space-Time Block Codes (STBC) exploits transmit diversity without knowledge of the channel by coding symbols over time and multiple antennas. Extra redundancy is needed in such codes, and the achievable overall rate is lower than transmit diversity with channel information.

In our experiments, we compare SEL and MRC for up to three receive antennas, and estimate the value of transmit selection. The optional 802.11n diversity features of transmit beamforming and STBCs are left for future work; neither is supported by our NIC.

3.3 Spatial Multiplexing

Multiple antennas can also be used to send independent streams of information at the same time in the same frequency channel. This is known as *spatial multiplexing*. It achieves a *multiplexing gain* with multiple streams of data at moderate SNR rather than one signal at high SNR to improve the overall rate without extra transmit power. Theoretically the throughput scales linearly with the number of spatial streams [8].

Direct-mapped MIMO. The simplest way to multiplex is to transmit one data stream out each antenna. In a rich scattering environment where m transmit streams are received by n antennas, each receive antenna will measure an independent linear combination of the m signals. This is decodable when $n \geq m$ so that there are more measurements (n) than unknowns (m). Extra measurements ($n > m$) add diversity gain. All streams use the same rate and same power.

Advanced transmit techniques. There are many enhancements to direct-mapped, equal-power and equal-rate MIMO. Transmit selection can choose antennas with better spatial paths; transmit beamforming can increase SNR and reduce interference between streams; power and rate can be allocated unequally to optimize individual transmit streams that necessarily traverse different spatial paths with different channel gains. As with transmit beamforming, these techniques require transmitter knowledge of the MIMO channel.

Diversity-multiplexing tradeoff. Multiple antennas also can be used for combinations of diversity and multiplexing rather than one or the other.

We experiment with up to three direct-mapped streams and equal power and equal rates per stream. We assume a zero-forcing MIMO receiver [19], which solves the MIMO linear system in a straightforward way. We leave the advanced techniques such as unequal rates across the streams for future work. They are optional in 802.11n and neither is supported by our NIC.

4. TESTBED & METHODOLOGY

We describe our indoor testbed and the experimental methodology in this section.

4.1 Testbed & Nodes

Figure 1 shows the layout of our testbed. It consists of 10 nodes spread over 8 100 square feet in an indoor office environment. Each testbed node is a stationary desktop equipped with an Intel Wifi Link 5300 (`iwl5300`) a/b/g/n wireless network adapter. These NICs have three antennas, and can transmit and receive up to three spatial streams with equal modulation. In the most aggressive configuration, i.e., three spatial streams on a 40 MHz bandwidth channel, each stream using a short 400 ns inter-symbol guard interval with QAM-64 coded at a 5/6 rate for all streams, these NICs can send data at 450 Mbps. The NIC does not currently support op-

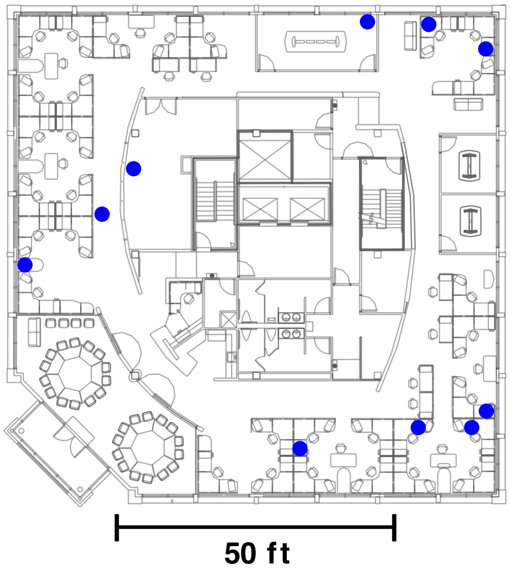

Figure 1: Our 802.11n testbed consists of 10 nodes spread over 8 100 square feet in an indoor office environment. The nodes are placed such that we have a large number of links between them, a variety of distance between nodes and diverse scattering

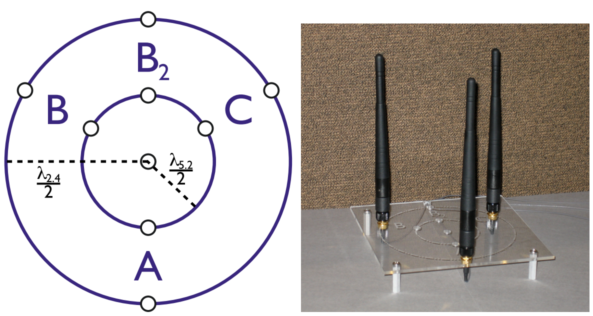

Figure 2: The antenna stand we use for consistent spatial geometry supports circular and linear arrays of 2 or 3 antennas with the correct $\lambda/2$ separation at either 2.4 GHz or 5 GHz.

tional features of 802.11n, such as unequal modulation for spatial multiplexing, Space-Time Block Coding and transmit beamforming.

As the antenna geometry is important for spatial diversity, we mount the three antennas per node on custom stands. These stands (Figure 2) allow for a circular array (using ports ABC) or a linear array of two or three (AB_2 and optionally center) antennas, with antenna separations of half the wavelength for either 2.4 GHz (channel 6, 2.437 GHz, $\lambda/2 = 6.15$ cm) or 5 GHz (channel 48, 5.240 GHz, $\lambda/2 = 2.86$ cm). We use the circular 3-antenna configuration for 2.4 GHz channel. It is robust and suited to dual-band NICs that need the wider 2.4 GHz antenna separation. Each antenna achieves 5 dBi gain for the 2.4 GHz band, and 3 dBi for the 5 GHz band.

Each node runs the Linux 2.6.26-rc6 kernel with a modified version of the `iwlagn` driver for the Intel WiFi Link 5300 network adapters [1]. Since the release driver only operates the adapter as a client in the 802.11n mode, and there were no commercially available 802.11n access points that support 3x3 MIMO, we modified the driver, the `mac80211` subsystem in the kernel, and `hostap` to run some testbed nodes as access points.

4.2 Measurement Tools

Our nodes allow us to send and receive packets in a variety of 802.11n configurations, and to observe the usual physical layer indications across the overall channel such as RSSI for received packets. However, we would like to observe detailed information about OFDM subcarriers in the wireless channel. This will let us look below the RSSI and better explain the packet level effects that we observe. To obtain this information, we make use of the 802.11n management action channel state information (CSI) frame. This is used during channel sounding to report the channel state from the receiver of a frame back to the transmitter. It is typically used for calibration or transmit beamforming. Instead, we configure the NIC to compute this feedback packet for every received frame, rather than just during sounding, and send it up to the driver instead of back to the transmitter.

The `iwl5300` provides channel state information in a format that reports the channel matrices for 30 subcarrier groups, which is about one group for every 2 subcarriers at 20 MHz. Each channel matrix entry is a complex number with 8-bit resolution giving the aggregate gain and phase of the spatial path between a single transmit-receive antenna pair. We export this channel state information, as well as other per-packet information, such as RSSI, AGC and noise via the driver to user-level programs for analysis.

4.3 Experimental Methodology

We study spatial diversity and multiplexing in two steps. First, we compute packet reception rates under different transmit and receive configurations to describe high-level performance. This alone allows us to estimate various diversity, array and multiplexing gains. Second, we use the low-level channel gains that were logged with each packet reception to look more deeply at the performance effects that we observe. This enables us to attribute different effects to their likely causes.

For each experiment, we generate and monitor traffic for each possible source-destination pair using the diversity or multiplexing setups specified (Sections 5 and 6). We use `iperf` to generate 5 second CBR traffic from the source and log received packets at the destination. The transmission bit rate is fixed during each run, and we iterate through all rates to study their performance. We also disable MAC layer retransmissions. This provides us with a simple and well-understood workload with which to observe packet delivery and channel gains.

We name antennas at each node as A, B, and C, and when using fewer than all three antennas at either transmitter or receiver we consistently use A for a single antenna and A and B for two antennas. When using multiple antennas or spatial streams, the emitted power is divided across the antennas, i.e., reduced by 3 dB for each of two streams and 4.77 dB (implemented as 4.5 dB) for each of three streams. The maximum aggregate transmit power we use is 40 mW (16 dBm), the FCC-imposed limit for our band. We consider lower power levels down to -10 dBm. We always send packets with the standard 800 ns guard interval; this corresponds to 802.11n rates 6.5, 13, 19.5, 26, 39, 52, 58.5, and 65 Mbps. Unless otherwise stated, we operate the testbed on channel 48 in the 5 GHz band, which is centered at 5.240 GHz. This is an otherwise unoccupied channel and there was no noticeable interference. We do not use 40MHz channels or MAC layer aggregation. The experiments are run at night to minimize variations.

4.4 Testbed characterization

To provide a baseline for improvements, we measured the single antenna performance of all 802.11n links in the testbed. The "A only" line of Figure 3(a) shows the maximum sustainable rate (> 90% packet delivery) using all 802.11n single stream rates (6.5–65 Mbps) and only antenna A from each transmitting or receiving node. We see a wide variety of links in the testbed and the testbed is fully connected only at the lowest rate of 6.5 Mbps. There were no significant differences between antennas A, B and C. We also observe that the maximum link rate is a good measure of throughput in our testbed. We also note that the links exhibit little variation over time.

5. SPATIAL DIVERSITY

In this section we investigate the practical gains of 802.11n spatial diversity techniques, focusing on transmitting a single spatial stream with receive diversity (MRC) using up to three antennas, i.e., 1x1, 1x2 and 1x3 SIMO. We also consider transmit diversity briefly. Our goal is to assess the diversity benefit in terms of improved coverage and rate as well as reduced transmit power, when compared to single antenna or antenna selection-based legacy 802.11 NICs. Furthermore, we analyze channel state information for a deeper understanding of packet level observations.

5.1 Experimental Setup

The techniques we compare are: single antenna 802.11n (SISO); receive antenna selection (SEL) for up to three receive antennas, i.e. 1x1, 1x2 and 1x3 SIMO; maximum-ratio combining (MRC) for up to three receive antennas; and transmit antenna selection for up to three transmit antennas with a three-antenna MRC receiver. We use a single transmit antenna for all receive diversity techniques.

These techniques form a progression that helps us to tease apart the value of increasingly more complex systems. SISO is a baseline that represents 802.11a/g performance for single antenna per NIC, as is typical for experiments reported in the literature, e.g., [17]; it is an upper bound since our 802.11n configuration includes four added subcarriers, hence slightly higher bit rates. SEL shows improvement over the baseline with added receive antennas while retaining a single RF receive chain. Selection with two antennas is common for commercial legacy 802.11a/g APs and NICs. MRC is optimal for SIMO performance and is a building block for MIMO reception; thus it is likely to be the diversity technique used in 802.11n NICs. MRC with transmit selection provides a further potential improvement at the cost of requiring the transmitter to track the channel.

Since our 802.11n NIC implements MRC directly but not SEL, we emulate the effects of SEL. In addition to the SIMO experiments above we collect three SISO traces using receive antennas A, B, and C in turn. We then use the SISO trace that corresponds to the antenna with the best median RSSI over all packets in the 1x3 SIMO trace. This simple method (picking the antenna per trace rather than per packet) is sufficient because we observed by manual inspection that our channels differ and are stable so that SEL would consistently select the same antenna at least 95% of the time.

5.2 Improvement in Coverage and Rate

We first consider how receive diversity improves the maximum sustainable rate of links and hence increases the coverage over our testbed for a given target rate. Figure 3(a) shows a breakdown for diversity techniques of the maximum link rates at the maximum transmit power (16 dBm) over the testbed. Recall that, in our definition, a link supports a rate if the packet reception rate exceeds 90%. For our testbed, maximum link rate translates into maximum throughput.

With a single antenna, 15% of links in our testbed operate below 39 Mbps and only 33% operate at the maximum supported rate of 65 Mbps. All receive diversity techniques improve the lower end of this distribution, with less than 10% of links forced to operate below 39 Mbps. SEL allows 40% of links to operate at the maximum rate, while MRC more than doubles this improvement, to 60%. It also halves the number of links operating below 58.5 Mbps compared to SEL.

To understand these gains, we separate the improvement due to the diversity of independent spatial paths and due to adding more power (from receiving copies of the signal on multiple antennas). We compare a trace using single fixed antenna at 11.5 dBm (16 dBm less

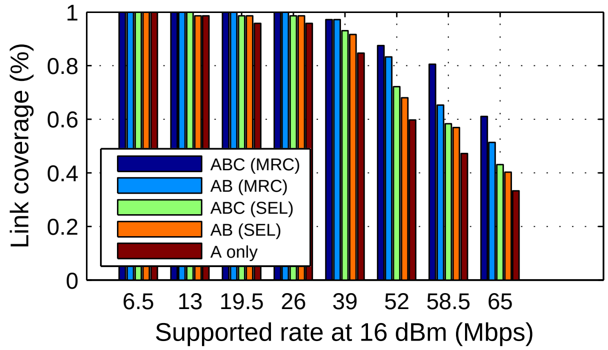

(a) Max supported rate (Mbps)

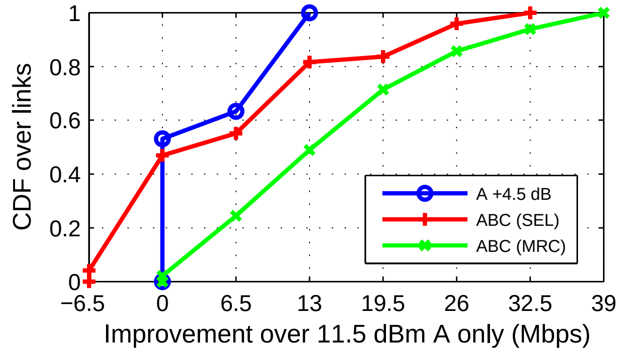

(b) Improvement in max rate (Mbps)

Figure 3: Improvements on maximum sustainable link rates due to receive diversity and power gain. Graph (a) shows a breakdown for diversity techniques of the maximum link rates at the maximum transmit power (16 dBm) over the testbed. Graph (b) shows the CDF of rate improvement for individual links for the three configurations compared to single antenna rate without added power.

the expected 4.5 dB gain from three antennas) to the higher power level due to array gain. Then we test 1x3 SEL and MRC configurations at the lower power level. This highlights improvements that on average include the expected array gain but also provide spatial diversity gains.

Figure 3(b) shows the results as a CDF of rate improvement for individual links for the three configurations (SISO with increased power, 1x3 SEL, and MRC), compared to the single antenna rate without added power. We omit those links which support the maximum 65 Mbps rate and hence cannot improve. Adding more power yields little performance improvement by itself — half the links see no improvement at all with added power, one third increase the rate by 13 Mbps and the remainder by 6.5 Mbps. Conversely the added spatial diversity significantly impacts rate improvement. SEL provides a rate increase of up to 32.5 Mbps, but half the links improve not at all and 4% even *drop* in performance as the larger RSSI link actually works worse. And using MRC improves link speed for 98% of links not already operating at the max rate, with a maximum change of 39 Mbps; half of the improvements due to MRC are 19.5 Mbps or larger.

This analysis highlights two points: (1) it is spatial diversity, not array gain that leads to improved performance using multiple antennas with OFDM, and (2) though SEL and MRC both exhibit the same theoretical diversity gain of three for a 1x3 link, MRC significantly outperforms SEL in practice.

5.3 Reduction in Power

As well as increasing the supported rate, spatial diversity enables links to operate at the same rate over larger distances or, alternately, at a lower transmitted power level for the same distance. We evaluate the latter for our testbed by measuring the minimum required RSSI on antenna A only to support the given link rate for single antenna and for MRC configurations. The change in RSSI is equal to that in transmit power, assuming the channel is stable between experimental runs. For this experiment, we need links that exhibit minimum RSSI transition points for the given link rate within our limited transmit power range (−10 dBm to +16 dBm), for single and multiple antenna configurations. This will exclude under- or over-powered links for the rate.

Figure 4(a) plots the reduction in transmit power needed to support the 52 Mbps rate using MRC as we increase from 1 to 2 and 3 antennas. 52 Mbps is chosen because it allows us to study the most links in the testbed (but still 61 of the 90 unidirectional links are excluded) and we expect similar graphs for other rates with small differences due to modulation details. We see a spread of values across links rather than a uniform ability to reduce transmit power. Going from a single antenna to all three reduces the required transmit power for a median link by 5 dB, and by as much as 11 dB. More pronounced gains come from shifting from one to two antennas (median 4 dB, max 9 dB), with a lesser gain moving from two to three (median 2 dB, max 5 dB).

We further compare the reduction in transmit power with the increase in received SNR due to MRC. For a single antenna channel there is a direct correspondence, e.g., adding 3 dB of receive power with a higher gain antenna allows a reduction of 3 dB in transmit power. However, this does not necessarily hold for multiple antenna techniques. Figure 4(b) is a scatter plot of the increase in SNR using MRC for the three scenarios (1→2 antennas, 2→3 antennas, and 1→3 antennas) versus the decrease in required transmit power to support the maximum sustainable link rate. The SNR for MRC scenarios is computed by adding the individual antenna

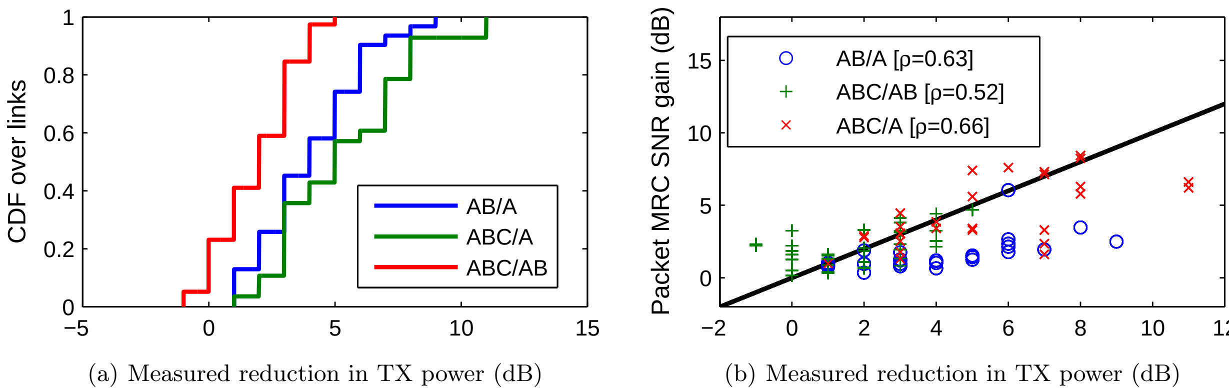

Figure 4: Measured reduction in transmit power needed to support a fixed link rate. Graph (a) plots the distribution of this improvement for 52 Mbps going from one to two and three antenna MRC configurations. It shows the most substantial reductions moving from one to two antennas. Graph (b) compares the reduction in transmit power with the gain in SNR. These two values are only weakly correlated.

SNRs in line with coherent combining. There is a fairly weak correlation, and in general reductions in transmit power are larger than the improvements in SNR. This is especially the case going from 1 to 2 antennas in which SNR gains are usually around 3 dB but transmit power reductions range from 1 to 9 dB. MRC is apparently effective through more than a simple increase in SNR — we explore the underlying effects next.

5.4 Frequency-Selective Channel Effects

We take a deeper look at the RF channel to understand the packet-level effects observed. Surprisingly, MRC substantially outperforms SEL, even though the expected power gains for the two techniques differ only by about 2 dB; for two antennas SEL is expected to provide a gain of 1.76 dB and MRC a gain of 3 dB, while for three antennas the SEL and MRC expected gains are 2.63 dB and 4.77 dB, respectively ([9], Ch 7). Further, our transmit power experiments tell us that MRC provides greater benefits than are anticipated from its improvements in SNR.

OFDM Channel State. Recall that we instrumented our NICs so that they can directly report the channel gains of the subcarriers on the OFDM channel (Section 4.1). This physical layer state sheds light well below the aggregate packet-level SNR. To the best of our knowledge it has not been exploited previously for commodity NICs due to lack of accessibility.

Example of a Link with SEL and MRC. In Figure 5(a) we show the subcarrier channel gains for an example link over the entire channel for individual antennas as well as for diversity techniques. Subcarrier power is measured in dB and normalized by the strongest subcarrier; there is at least 20 dB of variation. For each individual antenna, the gains vary and change slowly from one subcarrier to the next. For this link, antenna B has the largest RSSI (defined for OFDM as mean subcarrier power) and is consistently chosen by antenna selection. Thus, the three lines involving antenna B overlap completely. Even with SEL, fades of 15 dB (compared to 20 dB without it) are observed on some subcarriers. MRC, on the other hand, not only increases the total power on each subcarrier by combining subcarriers across antennas, but significantly reduces power differences (to ≈5 dB) in the process.

We use subcarrier strengths to generalize our example and characterize the difference between MRC and SEL across the links in our testbed. Figure 5(b) shows the gap between the power of the strongest and weakest subcarriers of each link. We plot this difference when using antenna A only and then also calculate the gap when using SEL or MRC with two or three antennas. We use the channel state information (CSI) averaged over the packets of our experimental runs to calculate these values. Since the CSI we have does not include per-subcarrier noise values, we assume a uniform and independent noise floor on different receive antennas.

This graph reveals a large distinction between SEL and MRC. Both improve SNR, by picking the best antenna or combining across antennas, but we know from (Figure 3(b)) that increased SNR (power) contributes the minority of the rate increase. It is a different story for the power distribution across subcarriers. SEL does not improve the power distribution over antenna A because it is simply selecting a different antenna. On the other hand, MRC does improve the power distribution substantially by making it flatter. It combats frequency-selective fading and narrows the median power differential between the strongest and weakest subcarriers from 15.5 dB (35×) with one antenna to 9.4 dB (9×) with two antennas and 8.2 dB (7×) with three. This is signif-

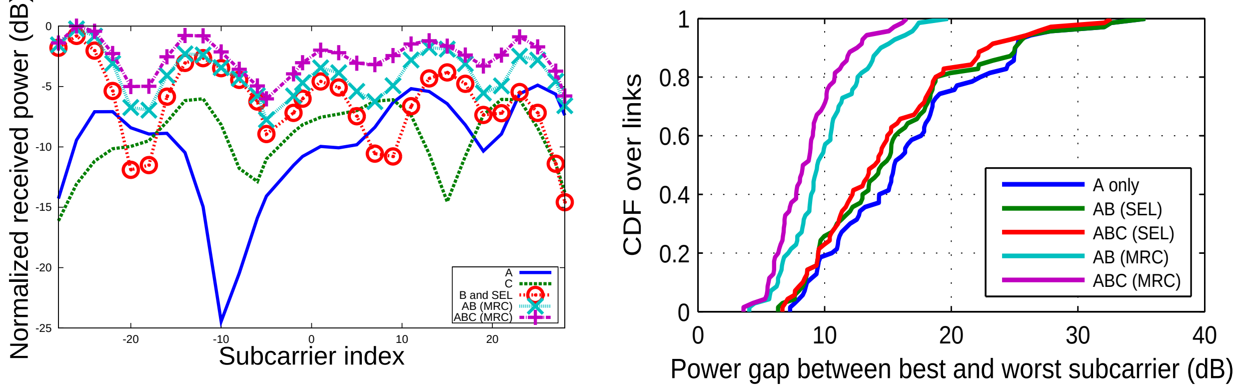

(a) Frequency-selective fading for an example link

(b) How diversity techniques affect selective fading

Figure 5: Frequency-selective fading over testbed links. (a) shows, for an example link, the received power on each subcarrier for individual antennas and under SEL and MRC. (b) shows how MRC combats frequency-selective fading and SEL does not, measured in terms of the gap between the strongest and weakest subcarriers. MRC narrows the power differential from 15 dB (32×) with one antenna to 7.7 dB (5.9×) with three antennas.

icant because the 802.11 style of OFDM modulates all subcarriers equally. Putting aside coding, which can tolerate a small fraction of highly faded subcarriers, flatter channels will correspond to a more efficient allocation of power and rate. Thus Figure 5(b) is suggestive of the advantages of MRC over SEL.

To better understand these gains, we relate packet reception to the subcarrier power distribution rather than the packet SNR. For 802.11 OFDM, the packet SNR roughly indicates the average per-subcarrier SNR. To tolerate faded subcarriers that are below this average, 802.11 uses forward error correction with rates of $1/2$, $2/3$, $3/4$, or $5/6$. Higher coding rates correspond to lower link rates. All else being equal, more skewed distributions will have more subcarriers that are faded below the average and require higher coding rates for successful reception. Unfortunately, however, these interactions are difficult to model [20]. Thus packet SNR is not necessarily a good indicator of maximum link rate.

As a simple model of subcarrier effects, we look for a percentile of the distribution that captures the expected performance improvement for a given coding rate. This approach is motivated because we know that the worst subcarrier will over-predict improvement (due to coding) and the best subcarrier will under-predict it (due to fading). We also know that a loose upper bound for rate r coding is that it can tolerate $(1-r)/2$ errors. However, these errors are unlikely to be restricted to the bottom $(1-r)/2$ subcarriers because even faded subcarriers are likely to lose only some of their (low order) bits. So we expect a higher percentile. Thus, we look at our experimental data to select the subcarrier percentile that best predicts the reduction in required transmit power.

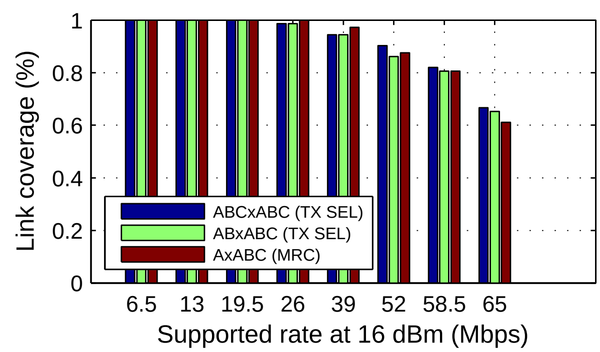

Figure 7: TX diversity adds little rate improvement over MRC.

For the 52 Mbps links analyzed in Figure 4(a), we find that the 25th percentile (22nd strongest of 30 subcarrier groups) subcarrier power is the best predictor of effective packet SNR. Figure 6(a) shows the gain in the representative subcarrier for different receive diversity configurations and Figure 6(b) shows the associated prediction error using this subcarrier. We see that this percentile is a good predictor on the whole, and that it predicts substantially larger performance improvements for MRC than for SEL. SEL does not result in any improvement the majority of the time. MRC provides a median improvement of 7.2 dB (5.4×) that is larger than the expected improvement of 4.77 dB (3×). This analysis, which captures the frequency-selective fading aspects, highlights why MRC is so much better than SEL.

5.5 Transmit Antenna Selection

Finally, we consider the performance improvement due to transmit antenna selection. By itself, transmit

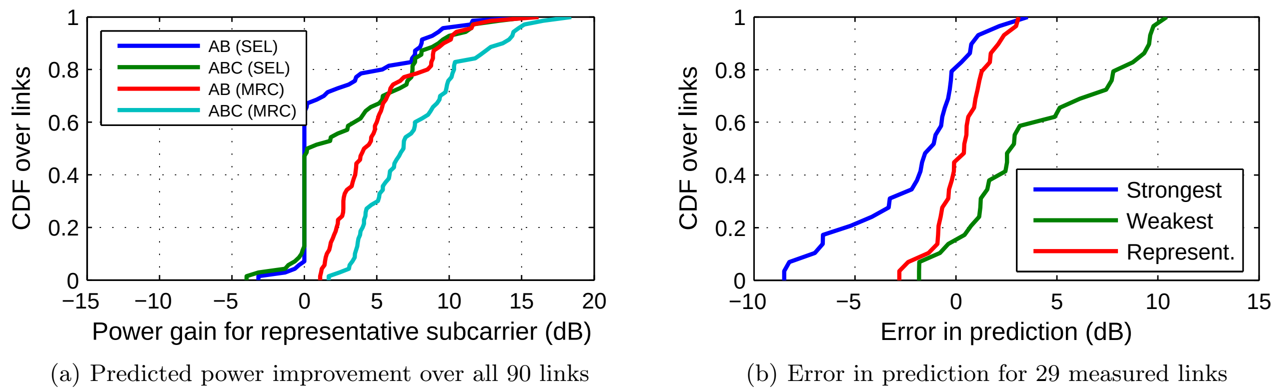

(a) Predicted power improvement over all 90 links

(b) Error in prediction for 29 measured links

Figure 6: Improvements predicted by subcarrier power. Graph (a) shows the power increase of the representative subcarrier (25th percentile for the 52 Mbps rate) for MRC and SEL. SEL often does not help while MRC boosts this power by a median of 7.2 dB (5.4×). Graph (b) shows the error predicting transmit power reduction as the change in power of the representative subcarrier. The error is usually within 1 dB and up to 3 dB (1.25–2×) whereas the weakest and strongest subcarrier consistently over- and under-predict, respectively.

antenna selection should match the gains of receive antenna selection, which we have already studied. Instead, we consider the benefits of adding transmit antenna selection to 1x3 MRC receive diversity. Figure 7 shows the difference between selecting antenna A, B or C for our links. There is little additional gain beyond MRC.

5.6 Summary

All the results in this section indicate that MRC is very effective, increasing the average link rate by 12 Mbps for our testbed (18 Mbps for links not already at the max rate), whereas SEL is less effective. The MRC gains exceed that which is expected for narrowband channel in theory. This is because MRC tends to flatten the frequency-selective overall channel and make better use of the power allocation.

6. SPATIAL MULTIPLEXING

We now investigate the gains achieved by spatial multiplexing in 802.11n. We begin our study by measuring the performance of the base SISO and MIMO spatial multiplexing cases that have equal numbers of transmit and receive antennas, i.e., 1x1, 2x2, and 3x3 antenna configurations. We then use channel measurements to understand these gains as we did in Section 5.

6.1 Multiplexing Gain

Figure 8 shows the distribution of the best supported rate on a link using one, two, and three streams in our testbed. The median single-stream rate of 58.5 Mbps improves to 104 Mbps with two streams (each stream at 52 Mbps) and to 117 Mbps with three streams (each stream at 39 Mbps) using these MIMO configurations. With MIMO, the median link rate compared to a SISO link increases by 78% going from one to two streams and doubles with three streams. However, the *average* rate increases from 53 Mbps to 101 Mbps (1.9×) with two streams and to 127 Mbps with three streams (2.4×). This is closer to the expected rate increases of 2 and 3 times.

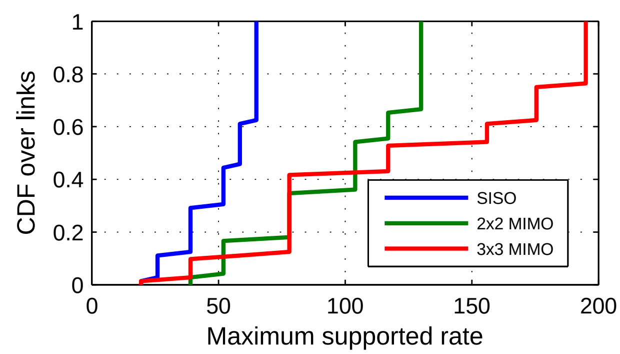

Figure 8: CDF of maximum supported rate for spatial multiplexing

Figure 9(a) shows the ratio of the best MIMO rate to the best SISO rate for each link. With two streams, multiplexing delivers the expected factor of two improvement for 60% of links while with 3x3 MIMO only 40% of links reach this goal. For nearly all links, two streams are better than one, and three is better than two for 65% of cases.

Some of the above gains are due to saturated links that operate at the maximum supported 802.11n with excess SNR. To evaluate multiplexing gains in a less well-connected testbed, we plot these same graphs after removing all links that operate at the maximum rate independent of the number of streams (Figure 9(b)). This dampens the improvement slightly, as now only 40% of

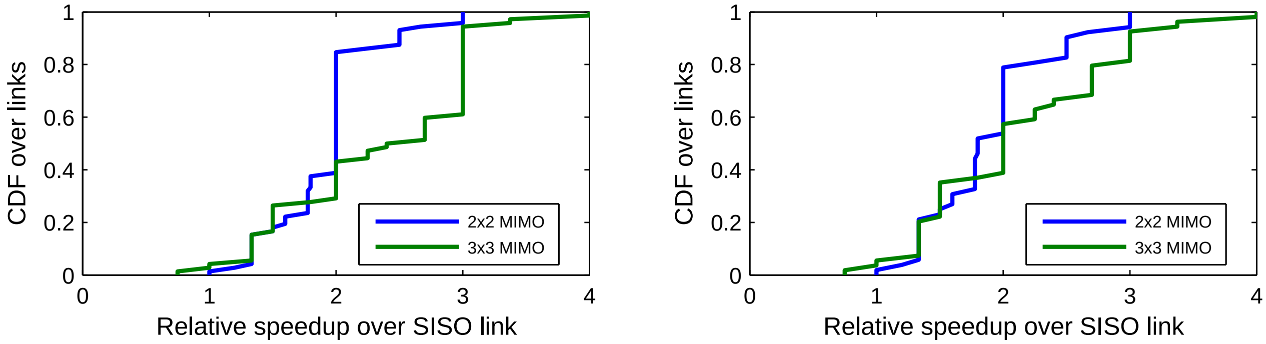

(a) With all links. (b) Removed links at the max rate

Figure 9: CDF of improvement in data rate for 2x2 (ABxAB) and 3x3 (ABCxABC) over 1x1 (AxA) configuration. Graph (a) shows the rate increase due to multiplexing across all links in the testbed. Graph (b) shows the rate increase after removing all links operating at maximum rate independent of the number of streams.

two stream and 20% of three stream MIMO links achieve the expected throughput, but MIMO still improves on SISO for nearly all links with 3x3 better than two about half the time.

Super-linear Speedup. One interesting feature in Figure 9 is that some links experience *more* than the expected factor of two or three speedup. Ignoring the SISO links at 65 Mbps, 21% of 2x2 and 7% of 3x3 MIMO links are better than twice or thrice as fast as when using SISO. The reason for this super-linear MIMO gain is added diversity. Each transmitted stream is received by multiple antennas which may increase the effective SNR of that stream, even after the transmitted power is divided over the multiple outgoing streams. Though the expected diversity gain of MIMO over SISO is 1, i.e., no improvement, performance in practice is not determined by average SNR but a lower percentile (Section 5.4) and some links see a sizable gain.

6.2 Channel Effects

We turn to channel measurements to explain why some links see large improvements and other links do not. Our analysis of frequency-selective fading for a SISO link (Section 5.4) found the 25th percentile of subcarrier SNR to be a reasonable predictor of the maximum supported rate. We now extend this model to multiple streams to predict rates for spatial multiplexing.

As defined in the 802.11n standard, the bits of a packet are separated into different streams after error correction but before interleaving. The consequence of this is that bit errors in one stream are isolated from other streams except at symbol boundaries. That is, the streams are nearly separated. Given that equal modulation is used for each of the streams, we expect the maximum supported MIMO rate to be limited by the stream with the minimum energy. To predict this maximum rate we proceed as follows. For each multiplexed link, we calculate the per-stream subcarrier SNRs using standard results for zero-forcing receivers [19]. We then estimate the effective SNR of each stream by finding the 25th percentile of the per-stream subcarrier SNRs. The lowest effective SNR predicts the maximum supported rate.

We plot the minimum effective SNR (in dB) against the measured MIMO rate in Figure 10. We observe a consistent relationship across the 1x1, 2x2 and 3x3 cases, although the SNR required to support a particular rate increases with the number of streams. Thus with a correction in the SNR based on the number of spatial streams, we can use the same model to predict rate given the SNR for MIMO as for we did for single streams in Section 5.4.

6.3 Summary

In our 802.11n testbed, an indoor office environment with good connectivity between links, high signal strengths and multipath scattering are conducive to the use of multiple spatial streams. In the median case, 2x2 MIMO achieves the expected 2× speedup for 2x2 MIMO, and 3x3 MIMO delivers 80% of the expected performance with a median rate increase of 2.4×. We see gains that are only slightly lower even if we exclude links that are capped by the maximum 802.11n rate to reflect less-connected networks. We see significant variability in gain across the links due to channel conditions. Fortunately, the simple model we developed can predict the maximum rate from the channel measurements with reasonable accuracy.

7. COMBINED BENEFITS

So far we have looked deeply into spatial diversity and spatial multiplexing, operating individually. In this section, we evaluate how well links operate when any

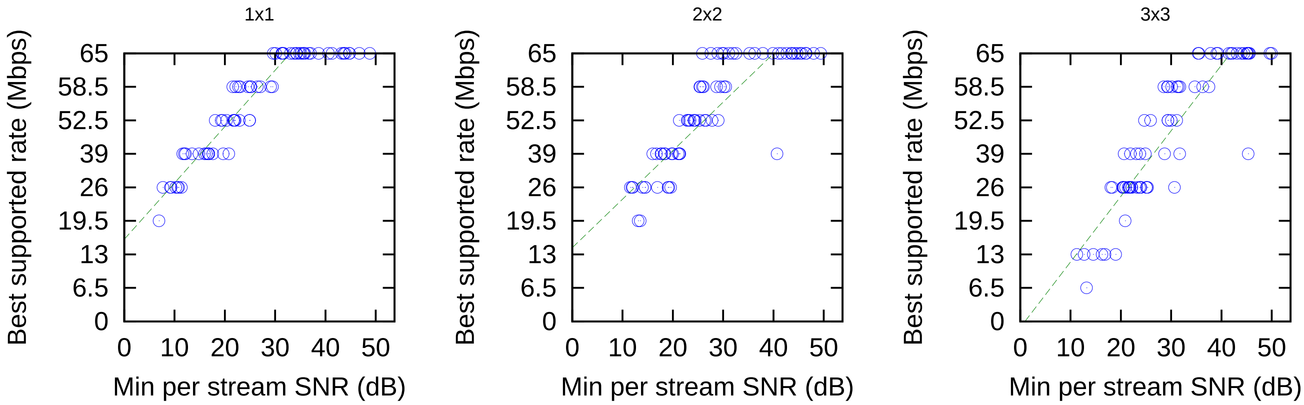

Figure 10: Scatter plot of the SNR of the 25th percentile subcarrier and maximum supported rate for 1x1, 2x2, 3x3 MIMO. The graph shows a linear relationship between SNR and rate.

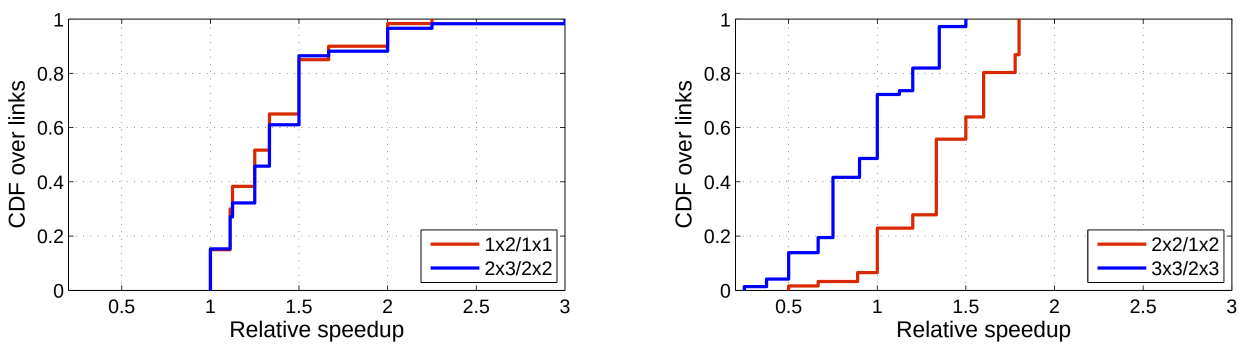

(a) Adding diversity to links

(b) Trading diversity for multiplexing

Figure 11: Interplay between multiplexing and diversity. Adding diversity yields comparable speedups whether there is one or two spatial streams. Adding multiplexing (i.e., another stream) helps substantially moving from one to two streams but less so moving from two to three streams.

combination of these techniques can be used. We also briefly discuss higher layer implications.

Relative Value of Techniques. In our testbed, the majority of the rate increase comes from spatial multiplexing: adding diversity to SISO and 2x2 MIMO links increases the median link speed by at most 10%, while 2x2 adds a 2× improvement over SISO and 3x3 improves to 2.7× as fast as SISO. However, this effect is likely biased by the well-connected, high SNR regime in which our testbed operates. Many links already run at the maximum per-stream rates supported by 802.11n, and cannot improve except by adding another stream.

To study this issue, we remove links that scale from the maximum 65 Mbps SISO rate to the maximum 2x2 or 3x3 MIMO rate. Figure Figure 11(a) and Figure 11(b) shows the relative speedup for the remaining links. We see that links that are not saturated benefit from added diversity most of the time (85%) irrespective of the number of streams. Conversely, the fraction of links that benefit from multiplexing drops rapidly with an increasing number of streams. 80% of links improve by adding a second stream while only 30% benefit from a third stream. We would expect this trend to worsen with more spatial streams because of equal rate stream constraint (i.e., the rate for all streams is determined by the weakest stream).

Overall Rate Improvement. Figure 12 shows how the overall distribution of link rates improves as we add each individual feature provided by 802.11n and links operate in the best configuration available in the cumulative feature set. The order in which we add configurations to SISO is: 1x2 MRC; 1x3 MRC; 2x2 MIMO; 2x3 MIMO; and 3x3 MIMO. That is, we add diversity first and then multiplexing when no further diversity can be added. Figure 13 shows the same six lines but now measures the distribution of the per-link speedup over SISO. In these plots we again remove saturated links with excess SNR that scale directly at 65 Mbps with 1, 2, or 3

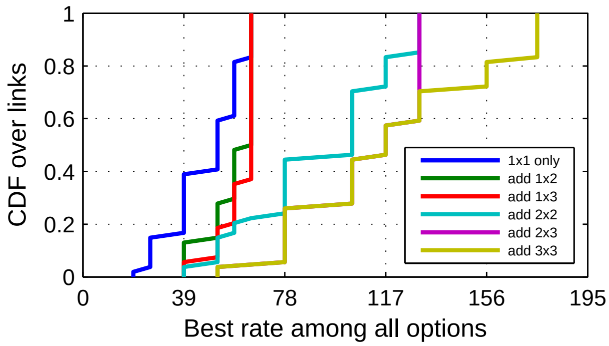

Figure 12: Improvement in the best supported rate as we add features. Each line shows a cumulative result.

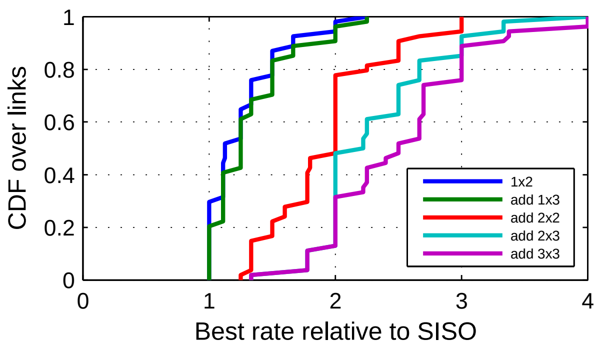

Figure 13: Relative improvement in the best supported rate as we add 802.11n features. Each line shows a cumulative result.

streams.

Examining the absolute rates (Figure 12) reveals that median link rate increases from 52 Mbps for SISO to a best case 117 Mbps, a 2.2× gain, while the average increases by 2.5× from 49 to 120 Mbps. For individual links (Figure 13), a few improve as little as 33%, some as much as 4×, and the majority (78%) of links improve between a factor of 2 and 3 with a median of 2.5×. We find that MIMO technology is largely delivering on its potential.

Predicting Rate from SNR. We also note one aspect of diversity that has implications for higher-level issues such as rate adaptation. Figure 14 plots the distribution of link SNR values for SISO and SIMO links against the maximum supported rate on that link. We see that with more receive antennas the spread of SNR values shrinks dramatically, as does the overlap in SNR values that support different rates. This means that SNR is now a good indicator of the maximum supported rate on a link, without the need to dig deeply into the per-subcarrier channel state.

In summary, we have drawn a few high level conclusions from our investigation of MIMO and 802.11n: diversity is more widely applicable than multiplexing as the number of streams grows; for the most part, 802.11n delivers on its promised gains; and with added diversity, SNR becomes a better predictor of rate. This last insight is simple and yet can simultaneously enhance and simplify a portion of the 802.11n rate selection algorithms. We expect that more discoveries in this vein will occur as we develop and deepen our understanding of 802.11n links and networks.

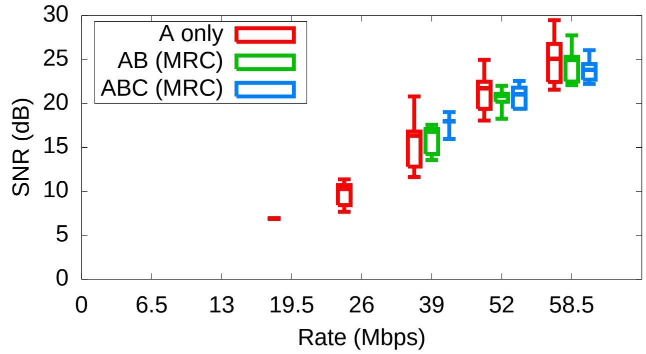

Figure 14: Diversity narrows the range of SNR values that support a given maximum link rate and makes the SNR to rate relationship more predictable.

8. RELATED WORK

The body of work related to this paper falls broadly into four classes: development of a theoretical understanding of MIMO and its bounds; algorithms for approaching these bounds; the simulation, implementation, and evaluation of MIMO systems; and finally research tailored to practical understanding and control of wireless networks in the context of IEEE 802.11.

There is a rich corpus of theoretical work inspired by the early foundational papers of Foschini [7], Foschini and Gans [8] and Telatar [27], and moving into areas as diverse as modeling wireless channels [5], information theory and coding [22, 12], and understanding the tradeoff between diversity and multiplexing [31]. Although this theoretical foundation helped guide MIMO transceiver architectures, it primarily provided an upper bound of achievable capacity realized by ideal rich scattering, time-varying channel models and algorithms with boundless complexity. Our focus is on the practical benefits of MIMO in the context of indoor 802.11n systems, for which many of the classic assumptions do not hold (e.g., channel models).

Algorithmic techniques have been developed to realize various points in the diversity–multiplexing tradeoff space. Diversity comes from such techniques as delay diversity, space-time block code (STBC) [3] and space time trellis codes (STTC) [26]. Increases in capacity were achieved through spatial multiplexing when channel state information is only available only at the receiver (CSIR) [28], and through resource allocation adapted per stream when channel state information is also known at the transmitter (CSIT) [11]. While many different such algorithms exist, some well understood in theory, simulation, or even prototype implementation, 802.11n is the first commodity, end-user MIMO wireless standard to see widespread (and increasing) adoption. Practical complexity, speed, and economic concerns influence the set of features adopted in hardware — for instance, the 802.11n NICs that we use do not implement any space-time codes and only support multiplexing based on CSIR.

Another body of practical work considers the design of MIMO-OFDM receivers with the aim of devising techniques that simplify processing yet achieve good capacity gains [23, 30, 21]. This work is followed with characterizations of MIMO channels measurements of the corresponding system performance in indoor and outdoor environments [25, 18, 16, 29]. These measurement studies are typically performed by custom hardware and are limited to measurements from a single link. Most studies find that real channels that are non-line-of-sight and have antenna separation of at least half a wavelength have the potential for large capacity increases that are a substantial fraction (70-80%) of that predicted by models. Our work corroborates these results and provides a concrete base in the context of practical 802.11n systems with commercially viable hardware.

Finally, a vast set of papers report on the performance of legacy 802.11 wireless systems in practice [14, 15, 6]. Wireless systems, like MRD [17], that exploit path diversity across distributed access points to improve loss resilience have demonstrated throughput gains of 2.3× over single radio schemes. However, these techniques operate above the physical layer and require coordination between distributed radios. The Roofnet group presents measurements of 802.11b [2] which attribute causes to the physical layer but without direct measurements thereof. To the best of our knowledge, ours is the first paper that relates the high-level performance of commodity 802.11n equipment to the low-level channel conditions; the one recent paper that takes a step in this direction does so without access to the low-level channel conditions [24].

9. CONCLUSION

Many future 802.11 systems will be based on 802.11n because of its superior performance. Given that 802.11n differs fundamentally from 802.11a/b/g, it is important to understand how well it works in real settings. The work we present in this paper is a first step in that direction. We experiment with commodity Intel 802.11n NICs that support 3x3 MIMO on a 10 node indoor testbed. We investigate how well the 802.11n spatial diversity and multiplexing techniques improve link rates, coverage and power requirements. By using custom, low-level instrumentation on our NICs, we are able to explain high-level 802.11n performance in terms of low-level RF channel measurements. No other papers do this as far as we are aware.

We find that 802.11n is highly effective at boosting performance with the simplest MIMO techniques that use equal power and equal rate transmit streams. When we make use of all diversity and multiplexing configurations up to 3x3, the median link rate in our testbed increases from 52 Mbps to 117 Mbps. This increase of 2.2× for 3x3 MIMO captures the bulk (≈74%) of the theoretical linear scaling with the number of node antennas, and 2x2 MIMO does even better with an increase of 2×. We find that diversity is consistently a large factor in these increases whereas multiplexing helps most when rates are otherwise capped or there are two streams. For instance, more than three-quarters of our links run faster in a 2x3 configuration (with two transmit and three receive antennas) than a 3x3 configuration. We use RF channel measurements to pinpoint frequency-selective fading as the reason for these large diversity gains. Diversity tends to flatten the overall channel, reducing the subcarrier power gap by 7 dB in our testbed, and this makes better use of the transmitted power. These and other results have many implications for topics such as rate adaptation. We plan to explore them in our future work.

10. REFERENCES

[1] Intel wireless WiFi Link drivers for Linux. http://intellinuxwireless.org/.

[2] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, volume 34, pages 121–132, New York, NY, USA, October 2004. ACM Press.

[3] S. Alamouti. A simple transmit diversity technique for wireless communications. *IEEE Journal on Selected Areas in Communications*, 49:1451–1458, 1998.

[4] P. Bahl, A. Adya, J. Padhye, and A. Wolman. Reconsidering wireless systems with multiple radios. *SIGCOMM Comput. Commun. Rev.*, 34(5):39–46, 2004.

[5] M. K. Çolakoğlu and M. Şafak. On the MIMO channel capacity predicted by Kronecker and Müller models. *Wirel. Pers. Commun.*, 47(1):91–100, 2008.

[6] S. Choi, K. Park, and C.-k. Kim. On the performance characteristics of wlans: revisited. In *SIGMETRICS '05: Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pages 97–108, New York, NY, USA, 2005. ACM.

[7] G. Foschini. Layered space-time architecture for wireless communication in a fading environment when using multi-element antennas. Bell Labs Technical Journal, Autumn 1996.

[8] 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.

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

[10] A. J. Goldsmith and P. P. Varaiya. Capacity of fading channels with channel side information. *IEEE Trans.*

Inform. Theory, 43:1986–1992, 1986.

[11] J. Ha, A. N. Mody, J. H. Sung, J. R. Barry, S. W. Mclaughlin, and G. L. Stüber. LDPC coded OFDM with Alamouti/SVD diversity technique. *Wirel. Pers. Commun.*, 23(1):183–194, 2002.

[12] B. Hassibi and B. M. Hochwald. How much training is needed in multiple-antenna wireless links. *IEEE Trans. Inform. Theory*, 49:951–963, 2003.

[13] IEEE. IEEE P802.11n/D7.0: Draft of 802.11 amendment 5: Enhancements for higher throughput. http://www.ieee802.org, 2008.

[14] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer. Understanding congestion in ieee 802.11b wireless networks. In *IMC '05: Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*, pages 25–25, Berkeley, CA, USA, 2005. USENIX Association.

[15] G. Judd, X. Wang, and P. Steenkiste. Efficient channel-aware rate adaptation in dynamic environments. In *MobiSys '08: Proceeding of the 6th international conference on Mobile systems, applications, and services*, pages 118–131, New York, NY, USA, 2008. ACM.

[16] P. Kyritsi, D. Cox, R. Valenzuela, and P. Wolniansky. Correlation analysis based on MIMO channel measurements in an indoor environment. *IEEE Journal on Selected Areas in Communications*, 21(5):713–720, June 2003.

[17] A. Miu, H. Balakrishnan, and C. E. Koskal. Improving loss resilience with multi-radio diversity in wireless networks. In *ACM MOBICOM*, pages 16–30, 2005.

[18] A. Molisch, M. Steinbauer, M. Toeltsch, E. Bonek, and R. Thoma. Capacity of MIMO systems based on measured wireless channels. *IEEE Journal on Selected Areas in Communications*, 20(3):561–569, Apr 2002.

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

[20] F. Peng, J. Zhang, and W. E. Ryan. Adaptive modulation and coding for IEEE 802.11n. In *IEEE WCNC*, pages 656–661, 2007.

[21] D. Perels, S. Haene, P. Luethi, A. Burg, N. Felber, W. Fichtner, and H. Bölcskei. ASIC implementation of a MIMO-OFDM transceiver for 192 Mbps WLANs. *European Solid-State Circuits Conference (ESSCIRC)*, pages 215–218, Sept. 2005.

[22] P. Rapajic and D. Popescu. Information capacity of a random signature multiple-input multiple-output channel. *IEEE Trans. Communications*, 48(8):1245–1248, 2000.

[23] H. Sampath, S. Talwar, J. Tellado, V. Erceg, and A. Paulraj. A fourth-generation MIMO-OFDM broadband wireless system: design, performance, and field trial results. *IEEE Communications Magazine*, 40(9):143–149, Sep 2002.

[24] V. Shrivastava, S. Rayanchu, J. Yoon, and S. Banerjee. 802.11n under the microscope. In *ACM IMC*, pages 105–110, 2008.

[25] R. Stridh, B. Ottersten, and P. Karlsson. MIMO channel capacity of a measured indoor radio channel at 5.8 GHz. *Asilomar Conference on Signals, Systems and Computers*, 1:733–737 vol.1, 2000.

[26] V. Tarokh, N. Seshadri, and A. Calderbank. Space-time codes for high data rate wireless communication: performance criterion and code construction. *IEEE Trans. Information Theory*, 44:744–765, 1998.

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

[28] P. Wolniansky, G. Foschini, G. Golden, and R. Valenzuela. V-BLAST: an architecture for realizing very high data rates over the rich-scattering wireless channel. *URSI International Symposium on Signals, Systems, and Electronics*, pages 295–300, 1998.

[29] K. Yu, M. Bengtsson, B. Ottersten, D. McNamara, P. Karlsson, and M. Beach. Second order statistics of NLOS indoor MIMO channels based on 5.2 GHz measurements. *IEEE GLOBECOM*, 1:156–160 vol.1, 2001.

[30] A. V. Zelst, T. C. W. Schenk, S. Member, and S. Member. Implementation of a MIMO OFDM-based wireless LAN system. *IEEE Trans. Sign. Proc*, 52:483–494, 2004.

[31] 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.