

On the Impact of MIMO Diversity on Higher Layer Performance

Ece Gelal, Konstantinos Pelechrinis,
Ioannis Broustis, Srikanth V. Krishnamurthy
University of California, Riverside
Computer Science & Engineering
Riverside, CA, USA
{ecegelal,kpele,broustis,krish}@cs.ucr.edu

Saif Mohammed, A. Chockalingam
Indian Institute of Science
Electrical & Communication Engineering
Bangalore, India
{smohham, achockal}@ece.iisc.ernet.in

Sneha Kasera
University of Utah
School of Computing
Salt Lake City, Utah, USA
kasera@cs.utah.edu

Abstract—In this paper, we shed light on the cross-layer interactions between the PHY, link and routing layers in networks with MIMO links operating in the diversity mode. Many previous studies assume an overly simplistic PHY layer model that does not sufficiently capture these interactions. We show that the use of simplistic models can in fact lead to misleading conclusions with regards to the higher layer performance with MIMO diversity. Towards understanding the impact of various PHY layer features on MIMO diversity, we begin with a simple but widely-used model and progressively incorporate these features to create new models. We examine the goodness of these models by comparing the simulated performance results with each, with measurements on an indoor 802.11n testbed. Our work reveals several interesting cross-layer dependencies that affect the gains due to MIMO diversity. In particular, we observe that relative to SISO links: (a) PHY layer gains due to MIMO diversity do not always carry over to the higher layers, (b) the use of other PHY layer features such as FEC codes significantly influence the gains due to MIMO diversity, and (c) the choice of the routing metric can impact the gains possible with MIMO.

Index Terms—MIMO; diversity; modeling

I. INTRODUCTION

The use of multiple antennas to achieve spatial diversity (MIMO diversity) can lead to a significant improvement in data transmission reliability in comparison to the traditional single antenna or SISO (single-input single-output) systems. One may envision that this in turn will significantly increase the wireless network throughput. This artifact has motivated recent studies on designing higher layer protocols that would benefit from the MIMO physical layer.

Surprisingly, we find that the diversity benefits due to MIMO are not fully exported to the upper layers due to complex cross-layer interactions. In particular, due to these interactions the MIMO diversity benefits are not as high at lower transmission rates as they are at higher transmission rates. In addition, the choice of the routing metric has to account for the use of MIMO diversity at the physical layer; as we show later, different metrics may yield differing levels of improvements over SISO. Unfortunately, the existing upper layer protocol designs over the MIMO physical layer, have failed to carefully consider these cross-layer interactions. Moreover, these existing protocols have been evaluated using very common but simplistic MIMO physical layer models [1], [2], [3].

In this paper, our goal is to comprehensively understand the cross-layer interactions between the physical, link, and routing layers in wireless networks with MIMO links, and develop realistic models of the MIMO physical layer towards the design and evaluation of higher layer protocols. Towards this goal, we

first perform experiments with MIMO spatial diversity on an 802.11n testbed. We compare the measured application layer throughput with simulated results derived from the most popular physical layer models that have been used in the past. We find that these models of MIMO diversity project higher layer performance results that are significantly different from those obtained via real measurements. Astonishingly, the existing simplistic models overestimate the MIMO gains over SISO by a factor of as much as 180.

Understanding the intricacies of the MIMO layer in a real deployment is very hard because none of the physical layer details are exposed to the upper layers in any useful manner. Furthermore, theoretical models do not necessarily capture the interactions between MIMO and other system functions and therefore, its realism. In order to understand the experimentally observed behavior of higher layer protocols built over MIMO, we take a slightly unconventional route. We start our exploration by simulating a simple physical layer model of MIMO and successively build a more complex model in our simulation, by progressively adding physical layer features and idiosyncrasies. Many of the key features that we add are missing in existing models. For each newly constructed model, we directly compare application-throughput obtained through simulations with throughput measurements from our 802.11n testbed. This process reveals the impact of various cross-layer interactions on the performance gains possible with MIMO diversity. Our study demonstrates that the throughput benefits due to MIMO diversity are affected not only by other PHY layer features (such as FEC codes) but also by higher layer functions (such as the routing metric). Using our most complete model, we obtain throughput results that match our experimental testbed results reasonably well. In particular, we obtain a performance difference of at most 11% between the simulation model-based and the testbed-based results.

To the best of our knowledge, our work is the first to provide an understanding of how the spatial diversity of MIMO interacts with the different layers of the protocol stack. We present a summary of some of our findings below:

(a) The link layer performance gains due to MIMO diversity over SISO, depend on the transmission rate. For a desired packet error rate (PER), the signal to noise (SNR) requirement with MIMO decreases by only 1 dB at low rates as compared to SISO (as we see later in Figures 2 and 4 (c)). This is in stark contrast to what has been reported at the physical layer. Previous studies have shown that at the physical layer there is a significant reduction in the SNR requirement for a desired bit error rate (BER) at all rates [4]. We observe this in our simulations as well; as an example, for a BER of 10^{-3} , due to MIMO diversity, the SNR requirement is reduced by 10 dB or

This work is supported by the U.S. Army Research Office under the Multi-University Research Initiative (MURI) grant W-911NF-04-1-0224 and the Multi-University Research Initiative (MURI) grant W911NF-07-1-0318.

higher (as seen later in Figure 3). This contrast between the bit-level and packet-level behaviors is due to complex interactions between packetization, use of forward error correction (FEC) codes and the fading characteristics (discussed later).

(b) The use of FEC impacts the benefits from MIMO diversity. At low transmission rates, the use of FEC codes are effective in coping with channel induced errors and thus, MIMO diversity is not of much use. At higher transmission rates, MIMO diversity increases the effectiveness of FEC.

(c) The end-to-end performance gains with MIMO depend on the choice of the routing metric. While the gains are significant with shortest path routing, the use of a path metric that accounts for link quality, such as ETX [5] (that we use in this paper), diminishes the gains. Note that in absolute terms, ETX still provides performance benefits over shortest path routing; however, MIMO diversity provides diminished performance gains over SISO. In our experiments, we find that as compared to shortest path routing, the end-to-end throughput gain is lower by about 15% when ETX is used. We choose ETX in this paper because it is a well understood path metric. Several other path metrics, some shown to perform better than ETX, have also been proposed in existing work. Routing based on these other path metrics is expected to bring up similar issues that we address while studying ETX.

Scope: Although MIMO products are commercially available, the IEEE 802.11n specification is not a standard yet. Hence, different vendors may implement the protocol differently; future changes may also be forthcoming. Thus, the critical in depth understanding of how MIMO links operate and the understanding of cross-layer interactions that we obtain in this paper, will be useful in the design and evaluation of future MIMO-based systems.

The rest of this paper is organized as follows. In Section II we provide background on spatial diversity and discuss related efforts. We describe the measurements on our testbed in Section III. In Section IV, we describe the models of MIMO diversity and present the link level performance with each model. We present our higher layer evaluations in Section V. Our conclusions form Section VII.

II. BACKGROUND AND RELATED WORK

In this section, we provide background on MIMO diversity and the IEEE 802.11n draft. We also discuss related work.

Spatial Diversity with MIMO: Spatial diversity with MIMO is achieved using Space-Time Block Codes (STBC). With STBC [6], [4], [7], appropriately encoded data blocks are transmitted from the elements of an antenna array at different times. If the fades between pairs of transmit and receive antenna elements are independent, STBC offers a lower BER (bit-error-rate) for a given SNR, in comparison to a SISO system [4]. With STBC, the information bits are first mapped on to modulation symbols, which are then input to the STBC encoder. The outputs of the STBC encoder are space-time codewords that determine which symbol is transmitted from which antenna and at what times. The received symbols are combined and decoded using estimates of channel gains (called channel state information or CSI) between each antenna pair. The receiver typically obtains CSI from pilot tones transmitted from each antenna of the sender. Use of STBC reduces the variations in the SNR and thus, increases the probability of correct decoding. In our work we focus on *Alamouti Codes* which are STBCs used in 2x2 MIMO systems [6]. We discuss the generality of

our work in Section VI. More details on STBC are found in [4]. Space-time diversity or spatial diversity is simply referred to as MIMO diversity in this work.

IEEE 802.11n Draft: The IEEE 802.11n specification supports both STBC and space-division multiplexing (SDM) modes of MIMO [8]. 802.11n supports 24 different data rates using different numbers of spatial streams and with different modulation and coding schemes. It also supports channel bonding wherein adjacent channels (of 20 MHz each) can be combined to form a wider channel of 40 MHz. In our work, we focus on STBC. In order to compare SISO and MIMO performance on our testbed, we operate using 20 MHz bandwidth in the 2.4 GHz band. We consider the FEC code rate $R=3/4$ with 4-QAM, 16-QAM and 64-QAM modulation schemes; the corresponding data rates with 802.11n are 19.5 Mbps, 39 Mbps, and 58.5 Mbps, respectively. With 802.11g, the corresponding data rates are 18, 36 and 54 Mbps, due to a different channel spacing. These map onto rates of 1 bps/Hz, 2 bps/Hz and 3 bps/Hz with the two specifications.

Related Work: As discussed earlier, higher layer protocol evaluations with MIMO typically employ a model to represent the PHY layer (for example [1], [2], [3]). In [1], the authors design a cross-layer routing protocol that adapts its behavior based on feedback from the MAC layer. Their model characterizes gains due to MIMO diversity as follows: (i) if nodes are equipped with k antennas, each MIMO hop spans k SISO hops (due to diversity), or, (ii) diversity increases the packet delivery rate on a link by a constant factor. In [2], diversity is created via node cooperation (referred to as Virtual Multi-Input Single-Output or VMISO). The model used considers the performance gain due to diversity (diversity gain) to be a constant. A similar model is used in [3]. *As we show later, these representations overly simplify the MIMO PHY layer.*

In [9], Shrivastava et al., perform an extensive measurement-based study with 802.11n. In terms of space-time diversity, the authors focus on the performance improvements using maximum ratio combining. Their study does not consider the cross-layer dependencies examined in this work.

In [10], Toledo et al., examine TCP performance with the SDM and STBC MIMO modes. The interactions between diversity and different modulation schemes or error correction codes have not been considered. Furthermore, the implications of packetization and routing have not been examined.

SISO PHY layer models employed in ad hoc network simulations are studied in [11], [12], [13]. In [13], Takai *et al.*, examine the impact of PHY layer models on the evaluated performance of ad hoc routing protocols. In [11], Kotz *et al.*, review six ad hoc simulation assumptions and evaluate them with regards to a real system via outdoor experiments. Both of these studies find that higher layer protocols may differ in performance with different PHY layer models.

There are some physical layer efforts on building packet error estimators for MIMO-OFDM systems. Bjerke *et al.*, estimate the PER based on the SNR in MIMO-OFDM systems [14]; however, their models are only validated via simulations. In [15], Daniels *et al.*, perform measurements of signal waveforms to estimate the fading patterns. These measured channel variations are used in simulations to compute the PER. They do not perform protocol simulations.

III. MEASUREMENTS

In this section, we describe our measurements on an 802.11n indoor experimental network. Our measurements serve as benchmarks for evaluating the fidelity of the MIMO-STBC simulation models that are considered in this work.

Our testbed includes 7 Dell S530 desktop PCs with 2.5 GHz Intel dual-core processor and 1 GB RAM. The deployment is depicted in Fig. 1. Each node is equipped with an RT2860



Fig. 1. A snapshot of our 7-node 802.11n testbed. Nodes are represented by diamonds along with their IDs.

miniPCI card that supports 802.11a/g/n, and an EMP-8602 Atheros-based card that supports 802.11a/g. We use 5 dBi gain omnidirectional antennas and the default transmission power of 20 dBm. We carry out link-level measurements and examine the measured data in terms of both bit-level and packet-level statistics.

Packet Delivery Statistics: We collect packet statistics on both SISO and MIMO links. We relocate our MIMO nodes to create 33 different transceiver pairs (at different locations) and adjust their powers to get a rich set of channel properties and SNR values. The measurements consider both (a) line-of-sight (LoS) links where the node pair is not separated by walls or other large obstacles, and (b) non-line-of-sight links (NLoS) where the end nodes are located in different rooms. The same set of links is used in SISO and MIMO modes. We transmit 10000 packets on every link; we measure the PER and RSSI (Received Signal Strength Indicator) based on the reportings of our cards¹. The RSSI is an average of the signal strength measured at each antenna. We conduct our experiments late at night in order to avoid interference from co-located WLANs. For each link, we conduct the MIMO measurements right after the SISO measurements in order to reduce the possibility of significant variations in channel quality.

Fig. 2 presents the measured PER for a range of SNR values at different rates, using SISO and MIMO modes. Notice that a significant performance improvement with MIMO over SISO is only observed at the higher rate of 3 bps/Hz (unlike what has been assumed in [1] and [2]). For a desired PER, the SNR requirement only decreases by about 1 dB at transmission rates of 1 bps/Hz and 2 bps/Hz; the decrease is much more significant at 3 bps/Hz. We deliberate on the reasons for this observation in Section IV.

Bit Error Statistics: We also examine the burstiness of errors in our indoor setting by conducting bit-level measurements at the MAC layer. In these experiments we utilize our EMP-8602 6G cards. We modify the MadWifi driver (version 0.9.2) [17] to disable the CRC (cyclic redundancy check) in order to measure the BER at the MAC layer. We transmit pseudo-random payloads; this allows us to identify the bit errors at the receiver. Our measurements suggest that bit errors are likely

¹The cards operate at a constant noise floor level of -95 dBm, with an accuracy of 99.5%.

to occur in bursts of varying length. The burstiness of errors depends on the coherence time of the channel i.e., the duration for which the channel quality does not change significantly. Similar results have been reported in [18] for SISO links. Note that at the PHY layer (prior to decoding), the error patterns in coded bits may be different from that observed at the MAC layer. As we discuss in Section IV, the burstiness of bit errors is one of the contributing factors to the diminished gains with MIMO at the low rates.

IV. MODELING MIMO DIVERSITY

In this section we describe four different models of the space-time coded MIMO physical layer; each model reflects the relation between PER and SNR. We begin with a simple model, variants of which have been previously used in [1],[2] and [3]. Each of the subsequent models incorporates additional PHY layer components². We focus on rates of 1-3 bps/Hz; this is a representative set that brings to light the behavioral nuances of MIMO diversity (as we discuss in Section VI).

A. Model 1 (M-1): Constant Gain at a Given Rate

With M-1, a constant diversity gain is attributed to all (2x2) MIMO links at a certain transmission rate. If the SNR of a link is Γ_{SISO} dB and the corresponding BER in the SISO mode is α , then as per M-1, in the MIMO mode the same node pair achieves a BER $\beta < \alpha$. Here β corresponds to the BER that would be achieved with an SNR of $\Gamma_{SISO} + d$ dB on the SISO link. The difference, d , in the two SNR values (for a target BER of β) is referred to as the diversity gain. It is the “increase in SNR” needed on a SISO link in order to achieve the same BER as on the corresponding MIMO link. *As per M-1, the diversity gain d is a fixed value and is directly applied at the higher layers.* In order to achieve the same PER on a link, SISO needs an SNR of d dB more than MIMO.

Realizing M-1: In [3], d is quantified based on the bit-level simulations reported in [4]. In [4], the Alamouti coded MIMO system is used with BPSK, QPSK and 8-PSK modulation schemes, respectively, to achieve rates of 1, 2 and 3 bps/Hz. From the reported BER vs SNR curves we lookup the value of d for a target $\beta = 10^{-3}$ (as typically considered in prior work, e.g., [2]); these values are shown in Table I.

Data Rate	SISO SNR at BER= 10^{-3}	MIMO SNR at BER= 10^{-3}	Diversity Gain
1 bps/Hz	24 dB	7 dB	17 dB
2 bps/Hz	27.08 dB	14 dB	13.08 dB
3 bps/Hz	31.25 dB	18.75 dB	12.5 dB

TABLE I
DIVERSITY GAIN WITH M-1

B. Model 2 (M-2): Diversity Gain Depends on the SNR

M-1 considers d to be fixed at a given rate, irrespective of Γ_{SISO} . However, d depends on Γ_{SISO} [4]. In other words, the advantage due to MIMO is more pronounced in the higher SNR regimes, and the diversity gain is much lower when the reception SNR is low. In addition, for a given SNR, the BER increases if the link operates at a higher bit rate.

Unlike M-1, the model M-2 captures the aforementioned characteristics of the diversity gain. BER vs SNR results in [4] are quantized by creating bins of SNR intervals of 1 dB. The diversity gains for each of these SNR intervals are then obtained

²For every new model with an added PHY layer component, we also incorporate the new component in the the corresponding SISO system.




Fig. 2. A Comparison of PER, measured on SISO and MIMO links at different rates.




Fig. 3. Bit-Level Models.

from the BER vs SNR curves for each of the three chosen rates. We plot the values as continuous curves in Fig. 3(a); we observe that as the SNR increases, the difference between the achieved BER with MIMO relative to SISO increases. The curves show that with this model, for a target BER of 10^{-3} , a reduction of about 20 dB is possible with MIMO diversity at the low rates.

Realizing M-2: To carry over the diversity gain to the higher layers, we first map the SNR to a BER. When a frame is received with a certain SNR, random bit errors are placed within the frame according to a Bernoulli distribution with a probability equal to the BER. The reception is deemed successful only if the total number of bit errors is zero.

Transition to higher fidelity models: Although M-2 is more accurate than M-1 in depicting the behavior of MIMO diversity, it represents a PHY layer that is different from that described in the 802.11n draft [8] in two main ways:

(1) The modulation schemes used in 802.11n to generate rates of 1, 2 and 3 bps/Hz are 4 QAM, 16 QAM and 64 QAM (and not BPSK, QPSK and 8-PSK as assumed in [4]).

(2) The measured results correspond to a system where convolutional FEC codes are employed; the results in [4] do not assume any FEC codes.

Given this, we build a new model M-3, which has a higher fidelity with respect to the 802.11n system.

C. Model 3 (M-3): Improving Consistency with 802.11n

Towards constructing this new model, we implement the appropriate modulation and error correction schemes in a PHY layer simulator (we could not find previous PHY layer results that fulfill the above requirements). Prior to describing our model, we provide an overview of our PHY layer simulator.

Our PHY layer simulator: We use Matlab to carry out the PHY layer simulations. The functional modules incorporated are random bit generation, convolutional encoding, bit-to-QAM symbol mapping, space-time block coding (STBC), slow and fast Rayleigh fading generation, noise generation, STBC decoding, log-likelihood ratio (LLR) computation, and Viterbi decoding. Table II contains the simulation parameters.

Random information bits (0 or 1) are generated using the `randint()` function in Matlab. These information bits are orga-

nized into packets of size $N_b = 1500$ bytes (to conform with 802.11n). The information bits are then convolutionally encoded using a rate-1/2, constraint length 7 convolutional code using the `convenc()` function in Matlab. The rate-3/4 code is obtained by puncturing the rate-1/2 code. With reference to Table II, the constraint length of a code refers to the number of input (6 past and 1 present) data values that are used to generate the code; Generator polynomials are used to construct the coded output symbols. As per the IEEE standards, the generator polynomials are 171 and 133 (in octal). The output bits from the convolutional encoder are mapped to M -QAM-symbols. 4-QAM, 16-QAM and 64-QAM symbols with Gray mapping are used. These QAM symbols are then fed to the STBC encoder. We have implemented the 2×2 Alamouti code for space-time encoding. The output of the STBC encoder is multiplied with Rayleigh (complex Gaussian) fading samples, following which AWGN (Additive White Gaussian Noise) samples are added. In case of fast fading, we assume that the fade remains constant over one STBC block and varies independently from one STBC block to the other; this is widely referred to as the quasi-static assumption in the literature [4]. In the case of slow fading, we assume that the fade remains constant over the entire packet, and varies independently from one packet to the other. The SNR values in the simulations are varied by fixing the signal power and varying the noise variance. At the receiver side, we have implemented the maximum-likelihood (ML) decoder for STBC decoding. From the output of the STBC decoder, we compute the LLR values of the bits that form the QAM-symbols. These LLR values are then fed as soft inputs to the Viterbi decoder for decoding the convolutional code. The `vitdec()` function in Matlab is used. We calculate the BER and the PER at the output of the decoder.

The M-3 model: Using the simulator described above, we create M-3, a new model that incorporates the modulation and coding schemes that are used in 802.11n (Table II). We simulate M-3 with both fast fading and slow fading channels. These are two extreme cases. In reality the channel exhibits fade durations that are somewhere in between; we discuss the implications thereof later in this section. We refer to the two new models as M-3 (FF) and M-3 (SF), respectively.




Fig. 4. Packet-Level Models and Results from Experiments.

System Component	Specification
Error Correction Coding	Convolutional Encoding Viterbi Decoding
Generator Polynomials of the Encoder	$G_0=(171)_s$, $G_1=(133)_s$
Encoder Constraint Length	7
Code Rate	3/4 (with puncturing) for all modulations
Modulation schemes	QPSK, 16-QAM, 64-QAM
Packet Length (for M-3P)	1500 bytes

TABLE II
SPECIFICS OF OUR BIT-LEVEL SIMULATOR

The results from our PHY layer simulations depicting SISO and MIMO link behaviors of M-3 (FF) and M-3 (SF) appear in Fig. 3(b) and Fig. 3(c), respectively. As with M-2, we quantize the BER-SNR results (into bins of 1 dB each) and form tables based on these figures (the tables are not included here but creation of the quantized discrete levels is straightforward). From these figures we make the following observations.

a. The SISO BER behavior is worse with slow fading than with fast fading, for a given SNR: With fast fading, the bit errors are independent. FEC codes are able to better handle these errors. With slow fading, there is a high degree of correlation between bit errors (they occur in bursts). The duration of a fade (reflecting the size of the burst) is equal to the transmission duration of a packet. Here, FEC codes are not as successful as with the fast fading case. Thus, the SNR requirement for a target BER is higher with slow fading.

b. For a given SNR, the BER with MIMO is also worse with slow fading but to a lesser extent: The performance with slow fading is worse than with fast fading, again due to the reduced effectiveness of FEC codes when errors are correlated. However, space-time diversity reduces the possibility of encountering deep fades. Consequently, the errors are reduced and therefore the degradation is less significant.

Despite these additional levels of sophistication, the curves in Fig. 3(b) and 3(c) have a more gradual fall off with SNR compared to what is seen in the PER vs SNR curves in Fig. 4(c). We conjecture that this is because the measurements provide PER as opposed to BER statistics. In order to have a more meaningful comparison with the measured behavior on 802.11n MIMO links, we generate PER statistics using MATLAB simulations, as described earlier. In particular, we generate 1500-byte packets (as in our measurements), and discard a packet if any of its bits is in error. These PER statistics constitute our packet-level model, discussed below.

D. Packet-Level Model (M-3P)

The PER vs SNR results from our Matlab simulations are shown in Fig. 4(a) and 4(b). We make the following observations.

a. The PER with MIMO is better with slow fading than with fast fading, for a given SNR: This is a surprising result since the BER was worse with slow fading. It is an artifact of two effects. On the one hand, the FEC codes are less effective with a slow fading channel than with a fast fading channel. However, the errors are confined to fewer packets (than with fast fading where the errors are more spread out). The first effect hurts PER, while the second helps PER. The second factor is dominant with MIMO.

b. For a given SNR, the PER with SISO is still worse with slow fading as compared to fast fading, although to a lesser extent (than BER): The aforementioned two conflicting factors impact SISO performance as well. However, the second factor (error clustering) is unable to offset the poor performance of FEC codes when bursty errors are present.

Comparing the measured results with the packet-level models: Fig. 4(c) depicts the averages of the PER values for each SNR value from our measured results (see Fig. 2). A cursory comparison between the PER vs SNR curves with the packet-level models and the measured results indicates that there is no complete match as yet. However, the results with M-3P (SF) are much more similar to what is observed with measurements than with any of the other models. In addition, as we describe in the next section, M-3P (SF) also yields higher layer behaviors that are close to what is possible with the real system. In what follows, we discuss why there are still discrepancies between the results with the models and the measurements at the link level.

There are four differences between M-3P and the measurements.

1. Bursty errors of varying size: M-3P (FF) assumes a uniform distribution of bit errors within a packet; this is based on the premise that due to interleaving/de-interleaving bursty errors are converted into randomly placed errors. In other words, the model implicitly considers an infinite-length interleaver. However, in reality interleaving is done over finite blocks of data. Thus, bursty errors are likely to exist even after interleaving. In fact, as discussed in Section III, our measurements (and other prior work [18]) suggest that a burstiness in errors exists even in packets that are received at the MAC layer. To characterize the impact of bursty errors M-3P (SF) assumes that the fade duration is equal to the length of the packet. As discussed earlier, this impacts the performance of FEC with slow fading; in addition, errors are clustered into fewer packets (rather than being spread out among a large number of packets).

At lower rates, the clustering effect dominates; at these rates the number of bit errors are few, and there are both bursty and isolated bit errors. Although isolated errors are corrected with

FEC codes for the most part even with SISO, bursty errors remain. As discussed earlier, MIMO diversity does not help significantly in these settings. Compared to M-3P (SF), the measurements display less prevalent burstiness; consequently, more packets are affected for a given BER. As a result, the measured PER is higher than PER with M-3P (SF).

At the higher rates, there are many bits in error. In particular on SISO links, the FEC codes are unable to correct all bit errors. With MIMO, many of these errors are overcome due to diversity (in these regimes, the gain due to MIMO diversity is more significant). The burstiness of errors is more prevalent with M-3P (SF) compared to the measured results. Thus, the FEC codes are more effective in reality than what is seen with M-3P (SF). As a result, the measured results yield a lower PER for a given SNR as compared to the model.

In reality, the burst length of errors varies depending on the duration of the fades. Fade duration is an environment-specific factor and it is extremely difficult to measure. Furthermore, the time-correlation between bit-errors can vary temporally. A high-accuracy characterization of the burstiness of errors, to closely reflect what is observed in practice, remains a challenge. Markovian models have been used to characterize bursty errors [19], but in the complex real systems that we attempt to characterize, there may be long range dependence. Note that the Viterbi encoding/decoding itself induces error propagation among bits. As we see later, M-3P (SF) seems to reasonably characterize the performance at the higher layers. Thus, we believe that this model is sufficient for modeling static networks.

2. The use of OFDM: Unlike in the real 802.11n system, M-3P models do not implement an OFDM sub-system. The use of FEC in conjunction with OFDM (as in IEEE 802.11n) gives an additional level of diversity that may boost performance. Potentially we could better conform with measurement results if we could incorporate OFDM into our PHY layer simulations. The slopes of the measured results are much steeper than what is observed with M-3P (SF) (Fig. 4); this is likely to be due to additional orders of diversity which are possible with OFDM. However, measuring the frequency selectivity order L (the number of resolvable multipath components) is not an easy task; this is dependent on the environment and could also temporally vary. One could perform simulations with a number of values for L . Due to the very high complexity of simulations with OFDM, we plan to examine these issues in the future.

3. Correlated fading paths: In all of our MIMO models we assume that the path between a pair of antenna elements (one transmit and one receive) experiences a fade that is independent of the fade between any other such pair. This is the basis for most of the reported PHY layer results [4]. However, in an indoor wireless setting, this may not be the case. In other words, the fading would be Ricean [4] as opposed to Rayleigh as assumed in the simulations. Estimating the Ricean factor K [4] in a real system is difficult. Again, incorporating Ricean fading and evaluating the performance for many K values can improve fidelity but at increased simulation costs.

4. Path Loss and Shadowing Effects: In addition to the above factors, the BER vs SNR (or PER vs SNR) curves resulting from the models do not consider path loss or shadowing, whereas the measurement-based curves do. Since the network is static, one might expect these factors to cause long term “fixed” attenuations and to not affect the behavioral results (although

Simulation Parameters		
Channel Parameters	Path Loss Exponent (α)	3
	Stdev. of lognormal shadowing (β)	6
	Thermal Noise	-96 dBm
	Rayleigh fade duration	~ 0.02 s. for -5 dBm below RMS [20]
MAC Parameters	Receiver sensitivity	-80 dBm
	ECC threshold	1/12000
Traffic Load	5 CBR flows	10 pps/flow
	Packet Size	1500 bytes (12 Kbits)

TABLE III
HIGHER LAYER SIMULATION PARAMETERS.

the accuracy may be slightly affected). We seek to carefully evaluate these aspects in future work.

V. HIGHER-LAYER PERFORMANCE EVALUATIONS

In this section, we evaluate the effect of each model on the higher layer performance. We first describe our simulation setting. Subsequently, we discuss our results.

A. Simulation set-up and methodology

We implement the models in OPNET ver. 11 [21]. We provide implementation details below.

Channel model: OPNET uses a free space path loss propagation model. We implement lognormal shadowing between each pair of nodes. For each setting, the shadow fading attenuation among node pairs is held constant for the duration of the simulation (4 minutes); the same shadow fading is maintained for a given link for all of the models. Each packet experiences an instantaneous multi-path fade. The PER (or BER) vs SNR curves (in the previous section) were generated assuming Rayleigh fading. For consistency, we use a Rayleigh fading distribution in our higher layer simulations. The channel parameters are listed in Table III.

Simulating the models: Packets that are received with a power less than the receiver sensitivity are considered as noise. For the remaining packets, the SINR is computed from the reception power of the packet, the thermal noise, and the sum of powers from all interferers. If M-1 is used, the computed SINR is incremented by the diversity gain from Table I. The resulting SINR is compared to that required for a BER of 10^{-3} . If the value is higher than the requirement, the packet is deemed successful. If M-2 or M-3 (FF) is used, the reception SINR is mapped to a BER from the appropriate BER-SNR curve. Random bit errors are then placed within a received frame as per a Bernoulli distribution based on this BER. The packet is deemed successful if the total number of bit errors is zero. We do not simulate M-3 (SF)). In this case, we would need to generate bursty errors of the size of a packet (which essentially translates to generating errors at the packet level). Furthermore, directly applying the bit-level diversity gain with M-3 (SF) provides PERs that are much different with what was observed with measurements (especially on SISO links) as seen in Fig. 3 (c); this does not motivate a further study in this case. With the packet-level models, we measure the SINR and discard packets at the PER as per the corresponding PER-SNR curves. Finally, we perform higher layer simulations with the model created directly from measurements. Using this model, we discard packets as per the PER in Fig. 4(c). We call this our measurement driven model or **MDM**. We use MDM in multi-hop settings, since the IEEE 802.11n draft does not yet support the ad-hoc mode (essential for constructing routes). On single hop links, we directly compare the performance of our models with the measured performance.




Fig. 5. Link throughput performance at selected SNR values, simulated with each PHY model and measured on the testbed.

Miscellaneous issues relating to the PER computation:

We map the SINR (not the SNR) to the PER; in essence we assume that the interference also has a Gaussian distribution (like thermal noise). This leads to some inaccuracies, which however exist in all cases including the simulations driven by the measurements. It is extremely difficult (if not impossible) to accurately characterize topology dependent interference. Finally, the PER vs SNR (or the BER vs SNR) curves in Fig. 3 and 4 are plotted only for a range of SNR values. The SNR value could be high and the corresponding PER value may not be directly available from these plots. In such a case, we perform a linear extrapolation to compute the appropriate PER (or the corresponding BER values).

MAC layer implementation: We use the DCF mode of 802.11 operations. We use the SIFS, minimum backoff window and data rate values as specified in the 802.11g standard and the 802.11n draft while simulating the SISO and MIMO links, respectively. In the MIMO scenario, the Reduced Inter-frame Spacing (RIFS) equivalent of $2\mu\text{s}$ is used instead of the $10\mu\text{s}$ SIFS interval of the SISO setting. We do not use RTS/CTS messages, as this is a common design decision in practice [5]. In order to examine the impact of diversity gain with MIMO at the different rates of operation (1 bps/Hz, 2 bps/Hz and 3 bps/Hz), we fix the transmission rate to the desired value. In practice, a rate adaptation mechanism is typically used. Our results apply with such mechanisms; at each transmission rate, the rate-specific gain with MIMO will be achieved.

Routing in the multi-hop scenarios: At the routing layer we use the popular Dynamic Source Routing (DSR) protocol [22]. DSR has also been used in prior work on routing on MIMO links [1]. For the route discovery control packets we use the same fixed rate as for data packets; this is in order to ensure that the discovered route is sustained for data communications at that rate. We simulate scenarios when DSR chooses the minimum-hop routes or minimum-ETX [5] routes. The first policy refers to the DSR version that uses hop count as the routing metric. With this policy, DSR chooses the minimum hop route from among those that are discovered. When this route fails, other cached routes may be used (we examine DSR without caching later). With the second policy, we compute the ETX metric to be the inverse of the link's reliability. In particular, ETX is quantified in terms of the packet delivery rate (PDR) on the link for the corresponding rate [5]. Intermediate nodes are not allowed to report cached routes to the source as in the default ETX implementation.

Topologies: We first consider **single-hop** flows to study the link throughput performance. We vary the SNR on these links by adjusting miscellaneous channel parameters (such as the distance and shadowing). We only consider links in isolation; thus, the only noise is the thermal noise. Next, we consider




Fig. 6. Experiments in multi-hop settings: Simple models overestimate achievable E2E performance improvement with MIMO.

randomly generated **multi-hop** topologies of different densities, where nodes are distributed in a $1000\text{m} \times 1000\text{m}$ area. In our first set of simulations, five UDP flows are generated between random source-destination pairs at a rate specified in Table III (the results with different numbers of flows are similar). We change the location of the source and destination nodes to vary the end-to-end hop-distance. Each scenario is simulated ten times (hence, each scenario is evaluated with fifty different source-destination pairs). In our second set of simulations, we capture the impact of the routing metric. We modify routing policies to disallow the reporting of cached routes. Furthermore, we only consider flows in isolation to reduce route failures due to interference. These experiments accentuate the impact of the routing metric on the gains due to MIMO diversity.

B. Evaluation of the MIMO Models

We first present the results in single-hop settings. Next, we examine single and multiple flows in multi-hop scenarios. Our key findings are: (a) the bit-level models almost always overestimate the possible performance improvements with MIMO, and (b) both the performance improvement and the absolute SISO and MIMO performances using the packet-level model on the slow fading channel conform to a large extent with those observed using MDM.

Recall from the discussions in the previous section that there was not an exact match between the PER vs SNR results with M-3P (SF) and MDM. However, we observe a closer match in terms of application throughput results (to follow). This is primarily a consequence of retransmissions and the corresponding backoffs at the MAC layer.

Results from single-hop settings: We compare the link throughput obtained in our simulations with those obtained via measurements. Towards this, we first measure the SISO and MIMO throughput on different links in our testbed, at all three rates. Then, we simulate several single-hop topologies and obtain the throughput on SISO and MIMO links for a range of SNR values. We present the comparative results in Fig. 5 at two sample SNR values. We observe that the simulated throughput with MDM differs by at most 11% from what we measure on our testbed. This experiment justifies the effectiveness of MDM

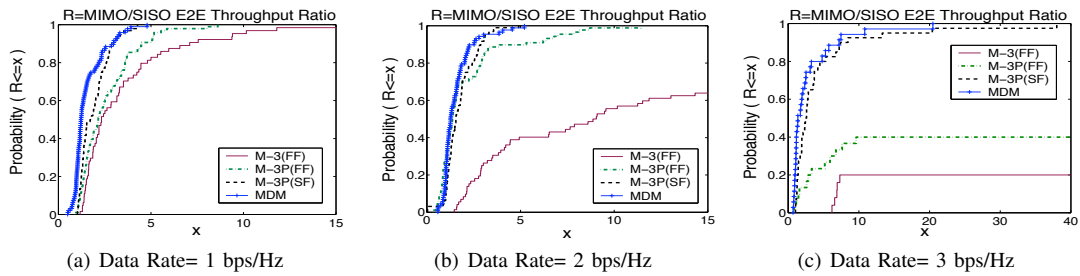


Fig. 7. Cumulative distribution of the per-flow throughput ratio (MIMO vs SISO).




Fig. 8. Distribution of the average E2E throughput and delays achieved by the flows at the low and high rates.

as an indicator of what might be possible (a benchmark) in multi-hop settings. The results with simulations are not exactly identical to those with measurements, because the wireless channel could be different in reality (as discussed earlier) from what is implemented in the simulator.

We also observe that the throughput with M-3P (SF) differs by at most 21% from the throughput with MDM, and 24% from the throughput on the testbed (at 3 bps/Hz data rate). On the other hand, the use of M-2 can lead to a drastic overestimation in the performance improvement with MIMO; the absolute throughput is much lower (77% lower with MIMO and 3,000 times lower with SISO) than what is observed on the testbed. Both of these observations can be attributed to the use of FEC codes. Since M-2 does not account for FEC codes, the packet transmissions with this model experience higher losses leading to poor throughput. This effect is more pronounced with SISO than with MIMO (in fact the throughput is close to zero) and thus, the throughput increase with MIMO diversity is extremely high. With the fast fading models, M-3(FF) and M-3P(FF), the improvements with MIMO over SISO are again much higher than what is observed on the testbed. This is because the errors are less bursty. With these models, while the FEC codes are more effective than they are in practice, the errors tend to be spread over many more packets, leading to poorer throughput. With M-3P (SF), the impact of the environment in the indoor setting (bursty errors) is better captured; thus, the behavioral results are similar to what is observed on the testbed.

Results from multi-hop settings: Next, we compare the end-to-end performance using each of the MIMO models from Section IV, with that using MDM in multi-hop settings.

We first discuss the results with the DSR implementation on OPNET (using minimum-hop routes).

Simply using the bit-level diversity gains may lead to erroneous conclusions: First, we examine the ratio of the achievable end-to-end network throughput using MIMO links to that using SISO links. In Fig. 6, we show these ratios in two networks with different node densities. The standard deviations among the samples collected for a given setting




Fig. 9. Performance improvements due to MIMO relative to SISO decrease with ETX-based routing.

(using a particular model and data rate) are shown as error bars. As with the single hop case, all models but M-3P (SF) *overestimate* the end-to-end performance improvement with MIMO compared to SISO, irrespective of the node density. In particular, the end-to-end throughput increase with M-1 and M-2 is as high as 180 times of what is achieved with MDA (for clarity of representation, we truncate these results in Fig. 6). Furthermore, both the fast fading models M-3 (FF) and M-3P (FF) overestimate the benefits due to MIMO, especially at the high rates (by 29 times and 13 times respectively, at 3 bps/Hz). M-3P (SF) shows a good degree of conformance to the results with MDM; in all simulated densities it differs by at most 19%. The reasoning behind these results is similar to that in the single-hop case.

In Fig. 7, we plot the CDF of the aforementioned ratio for all flows (*i.e.*, source-destination pairs) in the considered 50-node topologies. We denote this ratio by R ; the abscissa x represents the value of this ratio and the ordinate represents the probability that $R \leq x$.

Given that the bit-level models as well as M-3P (FF) fail to conform with MDM, for clarity of presentation we exclude the performance results with these models from our discussion on absolute performances.

M-3P (SF) provides absolute throughput results that exhibit a high degree of conformance with the results with MDM. In Fig. 8, we plot the CDFs of the end-to-end throughput using MIMO links at the lowest and highest rates. At the lower rate, the performance with M-3P (SF) is slightly superior; the




Fig. 10. Cumulative distribution of the per-flow throughput ratio, when minimum-ETX routes are used.

opposite is true at the highest rate. We observe that with both MDM and M-3P (SF), routes of similar lengths are used (see Table IV). However, the number of retransmissions incurred using these two models differ at different rates. At the lower rate, links experience more retransmissions with MDM, while the opposite is true at the higher rate. This is because the SNR requirement for a target PER is higher with MDM than with M-3P (SF) at the lower rates (as seen from a comparison of Fig. 4(b) and 4(c)). Consequently, long links that are established with MDM have poorer reliability than those with M-3P (SF). The opposite effect is seen at the highest rate of 3 bps/Hz. At this rate, the SNR requirement for a target PER is lower with MDM than with M-3P (SF).

	M-3P (SF)			MDM		
	1 bps/Hz	2 bps/Hz	3 bps/Hz	1 bps/Hz	2 bps/Hz	3 bps/Hz
<i>Hops (avg.)</i>	5.67	5.82	5.50	5.35	5.71	5.41
<i>ReTX (avg.)</i>	334.76	301.44	684.16	523.08	364.24	407.94

TABLE IV
HOP STATISTICS REGARDING THE FLOWS CONSIDERED IN FIG. 8.

Table IV shows that the average number of hops with 2 bps/Hz (or even 1 bps/Hz) is higher than with 3 bps/Hz. This is counter-intuitive; one might expect the average hop count to increase with rate. Upon closer inspection, we find that this is an artifact of using the corresponding rate for route query packets. There are two conflicting factors that occur when a route query is transmitted at a high rate. On the one hand, the queries are less reliably transported across each hop. On the other hand, each query transmission takes a shorter time thereby reducing interference on other query retransmissions. While the absolute shortest path discovered with 1 bps/Hz is shorter (or in some cases equal) than the one with 3 bps/Hz (or 2 bps/Hz) due to the former effect, there are fewer cached routes with 1 bps/Hz due to the latter effect. Some of the additional cached paths when using the higher rates bias the average hop count to a value lower than that with 1 bps/Hz. Recall that we transmit route discovery packets at the same rate as the data packets. We more closely examine the impact of route caching on a single flow later in this section.

The average end-to-end delay with M-3P (SF) exhibits a high degree of conformance with that using MDM. In Fig. 8(c) and 8(d) we plot the the CDF of the average delays with M-3P (SF) and MDM at the highest and lowest rates. The cause of the slight differences follow from our discussion earlier. To re-iterate, at 1 bps/Hz, M-3P projects a lower end-to-end delay compared to MDM due to a slightly greater diversity gain. The opposite is observed at the highest rate. The end-to-end delay performance with these two models conform best at 2bps/Hz (we omit this result due to space constraints).




Fig. 11. E2E delay and throughput performances with the most comprehensive model conform to those with our benchmark, MDM.

The most pronounced benefits with diversity are seen at the highest rate (Fig. 6). At the lower rates, the errors are few and the FEC codes mostly correct these errors. MIMO diversity does not provide significant *additional* benefits. However, at the higher rate, the errors are many and the FEC codes are not as effective with SISO. MIMO diversity reduces the extent of errors. The FEC codes are now more effective and thus, a significant gain is observed as compared to SISO.

Next we discuss the results when the route choice is based on the ETX metric [5]. We examine the ratio of the achievable end-to-end network throughput using MIMO links to that using SISO links, using ETX based routing. We observe that with ETX, this ratio decreases as compared to the minimum hop routing considered earlier. We define D to be the reduction in this ratio compared to minimum hop routing (Fig. 6). We consider the same 50-node topologies. As depicted in Fig. 9, **the gains due to MIMO decrease with the use of ETX by as much as 20%**. With ETX, the use of the more stable (albeit longer) routes with more reliable links (higher average SNR) improves the SISO performance as compared to minimum-hop routes. With the higher average SNRs on these SISO links, the FEC codes are more effective and MIMO does not provide significant *additional* benefits. Similarly, we observe that the conformance of the M-3P (SF) to MDM in terms of the delay performance increases³; the two curves (corresponding to M-3P(SF) and MDM) are closer to each other in Fig. 11(b) than in Fig. 8(d). This is because links chosen by ETX-based routing operate at relatively high SNRs, where the PER with M-3P (SF) is closer to that with measurements (see Fig. 4(c)).

A closer look at isolated flows: Our results so far provide macroscopic views of the impact of MIMO diversity on higher layer performance. Next, we consider isolated flows for further investigation; this significantly reduces the extent of route failures due to interference-induced link failures. In this controlled study, we also modify the DSR implementation to eliminate

³Due to lack of space, we only show the absolute ETX performance at the highest rate; the behavioral results are similar with the other two rates.

the impact of caching on route selection. With route caching in DSR, any node can respond to a route query based on the routes in their cache [22]. This can cause the minimum hop version to choose much longer paths than the actual shortest path. Here, we disable this feature and only the destination can respond to RREQs. We consider 30 scenarios and pick 7 scenarios where the source and the destination pairs are separated by a distance of approximately 5 hops (the results were similar for other pairs that were separated by different distances). As expected, we observe that the minimum-ETX routes are slightly longer (but more reliable) than the shortest paths on average, with both SISO and MIMO. But, the PDRs with SISO increase on such routes and thus, the gains due to MIMO are reduced. As a direct consequence, the MIMO to SISO throughput ratios are smaller using minimum-ETX routes compared to minimum-hop routes. We do not provide detailed results due to space constraints. However, in a nutshell, we find that at 3 bps/Hz, with M-3P (SF) the average route length in hops with SISO (MIMO) increases from 6.6 (5.9) to 6.8 (6.3) with the use of ETX. Correspondingly, for this rate, the gain due to MIMO reduces from 28% to 9%, relative to SISO. The observations are similar with MDM. The MIMO gain reduces from about 35% to about 7%, relative to SISO; the route length in hops increases with SISO (MIMO) from 6.4 (5.9) to 6.8 (6.3) due to ETX.

VI. SCOPE OF OUR WORK

Generality of our results: The measured results are from a static indoor wireless network. In such a network, slow fading would be a good model (as demonstrated by our studies). However, depending on the network deployment, appropriate models will have to be used. In a rapidly changing environment or with high mobility, the fast fading model is likely to be more appropriate (study of such a system will be considered in the future). If there are direct line of sight links (such as in an open field) the multi-path fading models may be completely invalid and may have to be replaced with simple path loss models. In fact, in such cases, MIMO may not offer any benefits. At even higher rates, the gains due to MIMO diversity may be more prominent. Finally, we wish to point out that the PER results that we provide are for a fixed packet length of 1500 bytes. The PHY simulations will have to be re-run if we had packets of varying sizes. One way of making the results more general is to classify packets into different types (short, medium long) and have PER rates for the average packet size in each class. Our studies provide a rich set of results that can be applied in a large set of possible scenarios (although not exhaustive).

On the choice of the transmission rates: We primarily consider the bit rates of 1-3 bps/Hz. It is evident from our work that at low rates MIMO does not offer significant benefits. As a result, with transmission rates lower than 18 Mbps (1 bps/Hz), we would not see any gain. On the other hand, 54 Mbps (3 bps/Hz) is the highest possible transmission rate with SISO. At even higher rates, one might expect the gains due to MIMO diversity to be more pronounced. For a fair comparison, we choose this to be the highest rate that we experiment with. Finally, we pick 36 Mbps (2 bps/Hz), since it is the median rate between 18 and 54 Mbps.

More complex MIMO systems: While the work here focuses on MIMO diversity and 2x2 systems, one might expect that similar issues will exist with higher orders of diversity, multi-user MIMO systems or with space division multiplexing.

Obtaining a deeper understanding and developing high fidelity models for these systems is a challenge.

VII. CONCLUSIONS

We undertake an in depth study of the impact of MIMO diversity on the higher layers. Our study takes into consideration various dependencies between the PHY, MAC and routing layers that influence the gains possible due to diversity from an application throughput perspective. The study is based on comparing the measured performance on an indoor 802.11n testbed with performance derived from a set of PHY layer models of varying complexity. Our study reveals that the diversity gain at the bit-level does not always appear at the packet level. Furthermore, when MIMO is used in conjunction with other PHY layer features such as FEC codes, it may not offer benefits at certain rates. Finally, we find that the routing metric used may impact the gains possible with MIMO. Our study can be useful in the design of higher layer protocols for use with spatial diversity; as an example, the behaviors of MIMO at different rates might influence the way in which the rate adaptation algorithm is designed. It may lead to the design of new routing metrics that jointly consider rate and the gains that are possible with diversity. We will consider these possibilities in the future.

REFERENCES

- [1] K. Sundaresan and R. Sivakumar. Routing in ad-hoc networks with MIMO links. In *IEEE ICNP*, 2005.
- [2] G. Jakllari, S. Krishnamurthy, M. Faloutsos, P. Krishnamurthy, and O. Ercetin. A framework for distributed spatio-temporal communications in mobile ad hoc networks. In *IEEE INFOCOM'06*.
- [3] E. Gelal, G. Jakllari, and S.V. Krishnamurthy. Exploiting diversity gain in MIMO equipped ad hoc networks. In *Asilomar Conference on Signals and Systems*, 2006.
- [4] H. Jafarkhani. *Space-Time Coding, Theory and Practice*. Cambridge University Press, 2005.
- [5] D. Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *ACM MobiCom'03*.
- [6] S. M. Alamouti. A simple transmit diversity technique for wireless communications. *IEEE JSAC*, 16(8), 1998.
- [7] V. Tarokh, H. Jafarkhani, and A. Calderbank. Space-time block codes from orthogonal designs. *IEEE Trans. Info. Theory*, 45(5), 1999.
- [8] IEEE P802.11 - TASK GROUP N. IEEE 802.11n: Standard for Enhancements for Higher Throughput.
- [9] V. Shrivastava, S. Rayanchu, J. Yoon, and S. Banerjee. 802.11n under the microscope. In *IMC*, 2008.
- [10] A. Toledo and X. Wang. TCP performance over wireless MIMO channels with ARQ and packet combining. *IEEE Trans. Mobile Comp.*, 5(3), 2006.
- [11] D. Kotz, C. Newport, R.S. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental evaluation of wireless simulation assumptions. In *ACM MSWIM*, 2004.
- [12] M. Takai, R. Bagrodia, K. Tang, and M. Gerla. Efficient wireless network simulations with detailed propagation models. *Wireless Networks*, 7(3), 2001.
- [13] M. Takai, J. Martin, and R. Bagrodia. Effects of wireless physical layer modeling in mobile ad hoc networks. In *ACM MobiHoc'01*.
- [14] B. Bjerke et al. Packet error probability prediction for system level simulations MIMO-OFDM based 802.11n WLANs. In *IEEE ICC*, 2005.
- [15] G. Daniels H. Suzuki, M. Hedley and J.Yuan. Performance of MIMO-OFDM-BICM on measured indoor channels. In *IEEE VTC'06*.
- [16] Ralink technology: MIMO. <http://www.ralinktech.com>.
- [17] The MadWiFi driver. <http://madwifi.org>.
- [18] A. Miu, H. Balakrishnan, and C.E. Koksall. Improving loss resilience with multi-radio diversity in wireless networks. In *ACM MobiCom*, 2005.
- [19] A. Willig A. Kopke and H. Karl. Chaotic maps as parsimonious bit error models of wireless channels. In *IEEE INFOCOM*, 2003.
- [20] R.J Punnoose, P.V. Nikitin, and D.D Stancil. Efficient simulation of Ricean fading within a packet simulator. In *IEEE VTC*, 2000.
- [21] Opet user's documentation. <http://www.opnet.com>.
- [22] D. B Johnson and D. A Maltz. Dynamic source routing in ad hoc wireless networks. *Mobile Computing*, 353, 1996.