

Modulation Rate Adaptation in Urban and Vehicular Environments: Cross-Layer Implementation and Experimental Evaluation

Joseph Camp, *Member, IEEE*, and Edward Knightly, *Fellow, IEEE*

Abstract—Accurately selecting modulation rates for time-varying channel conditions is critical for avoiding performance degradations due to rate overselection when channel conditions degrade or underselection when channel conditions improve. In this paper, we design a custom cross-layer framework that enables: 1) implementation of multiple and previously unimplemented rate adaptation mechanisms; 2) experimental evaluation and comparison of rate adaptation protocols on controlled, repeatable channels as well as residential urban and downtown vehicular and nonmobile environments in which we accurately measure channel conditions with 100- μ s granularity; and 3) comparison of performance on a per-packet basis with the ideal modulation rate obtained via exhaustive experimental search. Our evaluation reveals that SNR-triggered protocols are susceptible to overselection from the ideal rate when the coherence time is low (a scenario that we show occurs in practice even in a nonmobile topology), and that “*in situ*” training can produce large gains to overcome this sensitivity. Another key finding is that a mechanism effective in differentiating between collision and fading losses for hidden terminals has severely imbalanced throughput sharing when competing links are even slightly heterogeneous. In general, we find trained SNR-based protocols outperform loss-based protocols in terms of the ability to track vehicular clients, accuracy within outdoor environments, and balanced sharing with heterogeneous links (even with physical layer capture).

Index Terms—Cross-layer implementation, mobility, modulation rate, rate adaptation, urban, vehicular, wireless.

I. INTRODUCTION

RATE adaptation protocols adjust the modulation rate according to the quality of the channel. When there is mobility of the sender, receiver, or scatterers within the environment, the channel characteristics change, thereby inducing fluctuations of the channel quality, i.e., channel fading. Depending on the degree of such fluctuations, the previously appropriate rate could become underselected if the channel state has improved or overselected if the channel state has become worse. The inability to accurately choose the modulation rate for the

current channel condition leads to loss or unnecessarily long packet transmission times and inefficient use of the channel.

Rate adaptation protocols address channel fading in one of two ways. In *loss-triggered* rate adaptation, the transmitter interprets channel state based upon timeouts (failed delivery) or receipt of acknowledgments (successful delivery) following the transmission of data packets. Loss-triggered protocols use this delivery result of multiple packets to determine the appropriate modulation rate [1]–[6]. These protocols require limited MAC and PHY interaction and are widely implemented and evaluated in indoor and outdoor settings. In *SNR-triggered* rate adaptation, the receiver uses the signal-to-noise ratio to determine the modulation rate and informs the transmitter via the four-way handshake [7], [8]. These protocols have not been implemented previously due to the closed and inflexible MAC and PHY of legacy systems.

In this paper, we implement a custom cross-layer framework for rate adaptation that enables the evaluation of rate selection accuracy on a per-packet basis, revealing the reasons for throughput differences between protocols. We measure rate adaptation accuracy for diverse channel conditions characterized by fast-fading, multipath, and interference. We perform experiments in a laboratory setting with controlled and repeatable channels as well as in residential urban and downtown propagation environments, and use accurate channel measurements taken at 100- μ s time scales. These environments are also characterized by heterogeneous links, hidden terminals, and physical layer capture.

In particular, we make the following four contributions. First, we design a cross-layer rate adaptation framework and implement five key mechanisms used by rate adaptation protocols, out of which three are used by loss-triggered protocols and two by SNR-triggered protocols. We are the first to implement SNR-triggered protocols on hardware at MAC time scales comparable to commercial systems. In in-lab and urban outdoor environments, we evaluate rate adaptation protocols by measuring the success or failure of the protocols’ selected rate as compared to the ideal rate. We determine the ideal rate via exhaustive experimental search by replaying channel conditions through multiple rate adaptation mechanisms and experimentally identifying the rate decisions that maximize throughput. In this way, we characterize the multirate mechanisms’ inaccurate rate decisions to reveal the origins of poor throughput performance. In contrast, prior work neither compared protocols’ rate selection with optimal rate selection nor evaluated rate adaptation decisions on a packet-by-packet basis.

Second, we evaluate rate adaptation accuracy on diverse channel operating conditions including fast-fading, multi-

Manuscript received October 13, 2009; revised April 14, 2010; accepted May 20, 2010; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor K. Papagiannaki. This work was supported by NSF grants CNS-0325971 and CNS-0721894 and by a grant from the Cisco University Research Program Fund at Silicon Valley Community Foundation.

J. Camp is with the Department of Electrical Engineering, Southern Methodist University, Dallas, TX 75205 USA (e-mail: camp@smu.edu).

E. Knightly is with the Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005 USA (e-mail: knightly@rice.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2010.2051454

path, and interference. We find that as coherence time decreases (fast-fading), both loss-based and SNR-based mechanisms have low throughput. However, we show via per-packet evaluation that this poor performance is due to opposing rate selection inaccuracies: Loss-triggered mechanisms *underselect* when they require consecutive successful packets to increase their transmission rate, as this occurs with low probability in fast-fading environments. In contrast, SNR-triggered protocols *overselect* with a fast-fading channel due to sensitivity to coherence time. Yet, we show that when SNR protocols are *trained* according to the environment's coherence time, significant throughput gains can be achieved. Furthermore, we show that the need for such training increases with the presence of multipath, an effect we observe is strongly present in the downtown scenario but not in the residential environment.

Third, with controlled in-lab experiments, we investigate rate adaptation accuracy with competing, heterogeneous links (links of differing average quality), as commonly measured in outdoor environments. We show that a protocol designed to overcome the misinterpretation of collision-based losses and fading-based losses with out-of-range senders (the hidden terminal scenario) [6] is effective (i.e., high aggregate throughput and equal sharing) when the competing links are statistically equal in quality. However, we find that the protocol has a severe throughput sharing imbalance whenever even slight differences in average link quality exist between competing transmitters.¹ We show that this is due to the slight difference in channel quality driving the system to a state in which only one transmitter uses the four-way handshake significantly more often, thereby giving it increased protection from hidden terminal collisions. With higher link heterogeneity between competing transmitters, the physical layer capture effect occurs in which the stronger link is able to successfully transmit packets to the receiver even with simultaneous transmissions from a weaker transmitter. We present the first evaluation of rate adaptation performance coupled with capture and find that their joint interaction can cause significant unnecessary reductions in modulation rate.

Finally, we perform experiments in two practical outdoor environments: residential urban and downtown scenarios. Independently and jointly, we evaluate each of the in-lab factors of fast-fading (now induced by mobility of the sender, receiver, or scatterers), interference (from an operational mesh network²), multipath (due to closely set buildings), and heterogeneous links (by spatial differences and obstructions between nodes). We characterize these environments with fine-time-scale channel measurements and find that even without sender or receiver mobility, and a maximum speed of only 30 km/h for scatterers (passing vehicles), the coherence time is 300 μ s, which corresponds to a speed of 250 km/h in an idealized propagation environment. This contrasts with a common assumption within rate adaptation work that the coherence time is much greater than the packet's transmission time, e.g., [8]. Although loss-triggered protocols have been widely deployed in practice for outdoor scenarios, we find that even in a static topology, these mechanisms are highly susceptible to rate selection inaccuracies triggered by the large number of environmental factors contributing to loss,

such as mobility, interference, path loss, and multipath. Yet, we find that although SNR-based protocols are indeed sensitive to changes in coherence time in outdoor environments, their rate selection accuracy is more tolerant to the frequent losses that occur in these scenarios. Moreover, through an experiment with a static sender and a mobile receiver at vehicular speeds, we find that loss-triggered mechanisms are unable to track channel changes due to the sequential rate stepping and sensitivity to failure of the protocol, whereas SNR protocols can track such mobility. Lastly, we find that once the propagation environment of an outdoor setting is characterized by the above factors such as coherence time, our in-lab experiments are able to predict the rate adaptation mechanisms' behavior in outdoor environments.

By providing the first implementation of SNR-triggered protocols and providing a framework for direct comparison of key rate adaptation mechanisms on a single platform with repeatable channels as well as well-characterized residential urban and downtown environments, our results have key implications for design of 802.11 clients and infrastructure. In particular, our results indicate that SNR-based protocols, as compared to loss-based protocols: 1) are better able to track mobility; 2) have higher robustness to heterogeneous links (including physical layer capture); 3) have higher accuracy in outdoor environments, especially with the presence of interference-induced losses; and 4) are able to overcome the overhead penalty of the four-way handshake using opportunistic, back-to-back transmissions to ensure equal airtime [8]. We conclude that with the addition of *in situ* training to ensure robustness to varying coherence time, SNR-based protocols are a sound alternative to loss-triggered protocols.

The paper is structured as follows. We first describe our cross-layer implementation in Section II. We experimentally evaluate the accuracy of rate adaptation protocols with different channel operating conditions (Section III) and heterogeneous links (Section IV). We then study the rate adaptation mechanisms in residential urban and downtown scenarios in Section V. We discuss related work on existing rate adaptation protocols in Section VI and conclude in Section VII.

II. MULTIRATE PROTOCOL IMPLEMENTATION

In this section, we describe our custom cross-layer framework and the design steps in implementing the rate adaptation mechanisms. We are the first to implement SNR-triggered protocols at MAC time scales comparable to commercial systems.

A. Wireless Open-Access Research Platform

The WARP platform, depicted in Fig. 1, was designed at Rice University, Houston, TX, and is used by a number of academic and industrial research labs for clean-slate protocol implementation of the MAC and PHY. Three main components of the WARP platform are of interest: 1) Xilinx Virtex-II Pro FPGA: MAC protocols are written in C and targeted to embedded PowerPC cores, whereas PHY protocols are implemented within the FPGA fabric to achieve the required parallelization; 2) multiple-input-multiple-output (MIMO)-capable radios: Up to four 2.4/5-GHz radio boards can be configured and can support wide-band applications such as OFDM; and 3) 10/100 Ethernet port: source and sink traffic and report feedback about the performance of the protocols. At the time of this work, WARP uses a 64-subcarrier, OFDM physical layer supporting modulations

¹Namely, scenarios in which the channel difference is insufficient to require a modulation rate change or to yield physical layer capture of data packets.

²TFA-Rice Mesh Network Deployment (<http://tfa.rice.edu>)

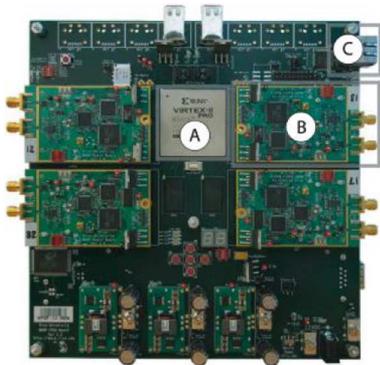


Fig. 1. WARP FPGA (A), MIMO-capable radios (B), and Ethernet (C).

of BPSK, QPSK, and 16QAM, within 10 MHz and without the use of convolutional codes. The WARP physical layer is similar to IEEE 802.11a with the same FFT size, cyclic extension, and preamble length, but the frame format and timing differs slightly. For example, the turnaround time of WARP exceeds the short interframe space (SIFS) interval of $20 \mu\text{s}$ by approximately $10 \mu\text{s}$. For further discussion of the WARP design, refer to [9].

B. CSMA Protocol Mechanisms

To implement a suite of multirate protocols, a first key step is to instrument the basic random access functions. Hence, we implemented a MAC protocol with the following five elements, analogous to mechanisms in 802.11: 1) carrier sense; 2) binary exponential backoff; 3) network allocation vector; 4) timeout; and 5) four-way handshake.

Carrier Sense. Since the MAC layer on WARP is single-threaded embedded C, much of the functionality must be interrupt-driven to efficiently transmit and receive packets. One such function is monitoring channel activity. Within the FPGA fabric, a timer accepts a specified amount of time and alerts the MAC via interrupt when the medium has been idle for that duration. This allows a transmitter to wait for a specified idle period before sending (i.e., carrier sense).

Binary Exponential Backoff. Another function of the channel-dependent timer is counting down only when the medium is idle. This is necessary for binary exponential backoff. Before transmitting a packet, the contention window is set and the MAC counts down with the channel-dependent timer. When the medium is busy, the timer must freeze and resumes counting down only after an idle period.

Network Allocation Vector. When a packet is received that is destined to another node, the duration field of the packet is used for virtual carrier sensing of the medium. This network allocation vector counts down regardless of the channel state in order to cover the packet exchange period.

Timeout. When a data packet is transmitted, a non-channel-dependent timer is set according to the timeout period for an acknowledgment. If the ACK is received before the timeout period, the timer is cleared, and for the purposes of loss-triggered rate adaptation, the data packet has succeeded. Otherwise, the data packet has been lost.

Four-Way Handshake. Finally, we implemented the RTS/CTS mechanism and added a field to the header in the

CTS message for modulation rate information to be sent back to the transmitter with SNR-based rate adaptation.

C. Cross-Layer Rate Adaptation Framework

We implemented five key rate adaptation mechanisms that existing multirate protocols utilize.

Consecutive-Packet Decision Loss-triggered Rate Adaptation. This mechanism increases the modulation rate after a number of consecutive successful transmissions and decreases after a number of consecutive failures. For this transmitter-based protocol, only counters are needed at the timeout (failures) and the reception of the ACK (successes). We use the specifics in [3] to implement the consecutive-packet decision mechanism. In particular, the mechanism increases the rate after 10 consecutive successes and lowers the rate after two consecutive failures. By default, RTS/CTS is disabled.

Historical-Decision Loss-triggered Rate Adaptation. A family of rate adaptation protocols [1], [2], [6] use a window of packets to select the modulation rate as opposed to consecutive successes or failures. Since [6] empirically outperforms [2] (and transitively [1]), we use the specifics described in [6] to implement the historical-decision mechanism, and the thresholds for rate increase and decrease are computed from the effective rates of the three modulations on WARP (5.4, 8.5, and 12 Mb/s). In particular, if a maximum tolerable loss threshold is exceeded, the rate is lowered, and if an opportunistic rate increase threshold is maintained, the rate is increased. By default, the mechanism uses basic access.

Collision/Fading Differentiation. Though still triggered by ACK/timeout loss interpretation, the rate adaptation schemes that implement the collision/fading differentiation are more immune to misinterpretation of collision-based loss. Protocols in [6] and [4], for example, dynamically enable the RTS/CTS mechanism upon loss with the assumption that a data timeout following a successful RTS/CTS exchange is likely due to poor channel conditions. While [4] toggles RTS on after a loss for a single packet and then off for the following packet, [6] uses a window of packets to enable RTS and is thus more robust to hidden terminals. We implement this mechanism and the dynamic use of the four-way handshake according to the specifics in [6]. The RTS window is halved if the last packet was lost with RTS or succeeded without RTS. If the packet is successful with RTS, the window remains unchanged.

SNR-Triggered Rate Adaptation. Rate adaptation based upon signal quality requires feedback from the receiver. In our case, we use the CTS packet for feedback. Use of this mechanism requires a mapping of channel conditions to modulation rates. In [7] and [8], this mapping was chosen according to the SNR-rate specification defined by the simulator itself. However, in our case, we measure the performance of the modulation rates according to SNR (see Section III for further discussion of this topic). The four-way handshake is used per-packet within this mechanism.

Equal Airtime Assurance. SNR-triggered rate adaptation with equal airtime assurance adds opportunistic transmission to the above SNR-based scheme mechanism: When the receiver sends a CTS back to the transmitter with a modulation rate that is above the base rate, the transmitter sends a burst of data packets in proportion to that modulation rate over the base rate. No backoff is performed between data packets to hold the channel.

Also, a queue has to be implemented to buffer packets for the burst. We use the specifics described in [8] to implement equal airtime assurance in which multiple data packets may follow an RTS/CTS exchange.

III. IN-LAB EVALUATION OF DIVERSE PHY OPERATING CONDITIONS

In this section, we explore the effects of physical channel conditions such as channel fading, multipath, and interference on rate adaptation protocols. To evaluate the accuracy of the rate choice by the protocols, we use per-packet evaluation by measuring the success or failure of the actual rate versus the ideal rate. We then measure rate adaptation accuracy according to these different channel conditions.

A. PHY Operating Conditions

We identify four channel conditions that have an effect on rate adaptation: 1) coherence time; 2) delay spread; 3) interference; and 4) physical layer capture. We define *coherence time* as the interval over which the channel is sufficiently constant (or coherent) to decode the received symbols with a particular modulation rate. We define fast and slow channel fading based on whether the coherence time of the channel is greater or less than the packet period, respectively. Multipath-induced fading occurs when two or more paths exist from sender to receiver, thereby inducing a delay between the same symbol from two different paths, called *delay spread*. Two or more paths can combine constructively or destructively at the receiver and, thus, also depend on the relative power level of the symbol versions. We define *interference* as channel activity that is undecodable by the sender and receiver. Finally, *physical layer capture* occurs when overlapping transmissions have sufficient signal power differences for one to be received correctly (we address physical layer capture jointly with its relevant hidden terminal scenario in Section IV).

Per-Packet Rate Evaluation. As stated in Section VI, unlike prior work, the observability between MAC and PHY of the custom cross-layer design allows per-packet evaluation of the rate adaptation mechanisms for a broad set of operating conditions as opposed to single-condition scenarios (such as only long coherence times). We say that a protocol selects the *ideal rate* when the modulation rate that is chosen has the highest throughput for the given channel condition. Specifically, for a given coherence time (repeatable with the channel emulator), there is an ideal modulation rate with the highest throughput for a given, mean SNR (recorded at the receiver).³ To evaluate this, the channel conditions must be repeatable, and the receiver must be able record statistics of each packet according to the actual modulation, its performance (correctly or incorrectly received), and the ideal modulation rate.

In evaluating rate adaptation, we test rate choice accuracy where at least one of the modulations is able to transmit data packets successfully. Since the header is sent at the base rate and is much shorter than the payload, it is almost always received correctly in this scenario. Thus, we find the selected modulation rate information for the data payload within the header. We then compare the selected rate with the ideal rate for the

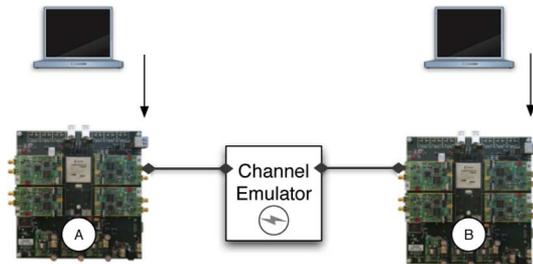


Fig. 2. Scenario to measure rate adaptation accuracy with different PHY operating conditions.

channel condition found via exhaustive search through all possible SNR and coherence time combinations to evaluate each rate adaptation mechanism. We infer from each packet the actual modulation rate of the payload, signal-to-noise ratio, ideal modulation rate, and if the packet payload is correctly or incorrectly decoded. Therefore, we classify packets according to three categories: 1) underselected (decoded payload, selected rate less than ideal rate); 2) accurate (decoded payload, selected rate same or greater than ideal rate); or 3) overselected (undecodable payload, selected rate greater than ideal rate).

Scenario. For repeatable channel conditions, we use a channel emulator for the experiments within this section. Fig. 2 depicts the experiment setup where the antenna ports of the two WARP boards are connected via wire to a Spirent Channel Emulator (SR5500). We use the emulator to produce Rayleigh distributed channels containing a wide range of coherence times and multipath delay spreads. The traffic is 1500-byte UDP at 20 Mb/s from sender A to receiver B. For the injection of interference, we use an Agilent Signal Generator (ESG-D Series) on the channel from the sender to the receiver.

B. Impact of Coherence Time

Ability to Track Changing Channels. We now evaluate the accuracy of the rate adaptation mechanisms presented in Section II-C with respect to the coherence time of the fading channel to test each mechanism's ability to track changes in channel conditions as a function of the time scale of the change. To achieve this, we measure the achievable throughput and rate selection accuracy for each multirate mechanism while varying the coherence time on a single Rayleigh fading channel with high average quality (-40 dBm). Specifically, we vary the coherence time from $100 \mu\text{s}$ to 100 ms, and we measure the accuracy of the four rate adaptation mechanisms triggered by consecutive-packet decision, historical-decision, SNR, and SNR with equal airtime assurance.

Fig. 3 shows the achievable throughput (Mb/s) as a function of the coherence time for each of the four mechanisms. For long coherence times (right portion of the graph), all protocols except the SNR-triggered protocol converge to similar performance as they are able to track the slowly fading channel. Unfortunately, we show in Section V that in practical outdoor environments, such scenarios are not encountered due to various factors. For long coherence times, the SNR-triggered protocol has significantly lower throughput than the other three protocols due to per-data-packet RTS overhead (including the equal airtime assurance mechanism that overcomes this overhead penalty). This result contrasts with simulation-based findings of performance

³In WARP, SNR is computed from the physical layer gain control, referenced to 1 mW (dBm), whereas SNR comparison is relative (dB).

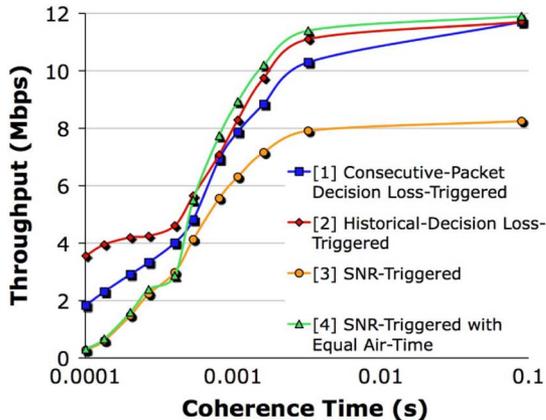


Fig. 3. Throughput versus coherence time for a high-quality Rayleigh fading channel.

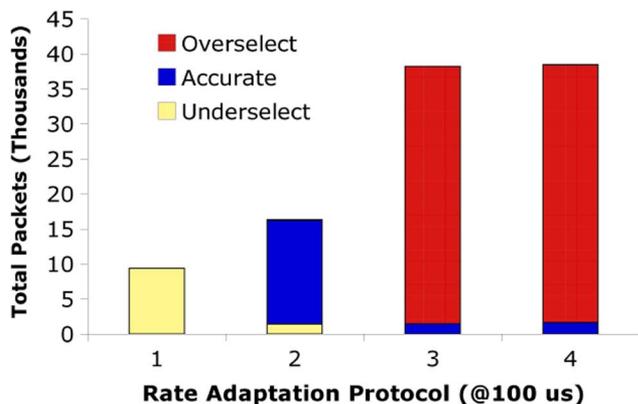


Fig. 4. Per-packet accuracy statistics for experiment depicted in Fig. 3 at $100 \mu\text{s}$ coherence time.

improvements over the consecutive-packet loss-triggered mechanism [7]. For the left portion of the graph (short coherence of the channel), the SNR-triggered protocols are now the worst performing. In contrast, the historical-decision loss-triggered protocol becomes the protocol with highest throughput, and the consecutive-packet decision loss-triggered protocol is between the two extremes.

Fig. 4 further describes the fast-fading case for a coherence time of $100 \mu\text{s}$. Each of the four protocols is on the x-axis and total classified packets (for a 60-s test) are on the y-axis according to three categories: underselected, accurate, and overselected. We see that the low throughput in the consecutive-packet decision protocol is due to underselection, and in the SNR-triggered protocol is due to overselection. Fig. 4 does not include packets that are at or below the ideal rate choices and are unavoidably lost due to severe channel fluctuations mid-packet (14.5 k and 17 k packets for consecutive-packet decision and historical-decision protocols, respectively). Further experiments below investigate these inaccurate rate decisions. *Finding: Multirate mechanisms triggered by both consecutive-packet decision and SNR have low throughput in fast-fading scenarios, but the low performance is due to opposite rate choice inaccuracies. Thus, evaluation of rate adaptation accuracy on a packet-by-packet basis is necessary to identify the reasons for the poor performance.*

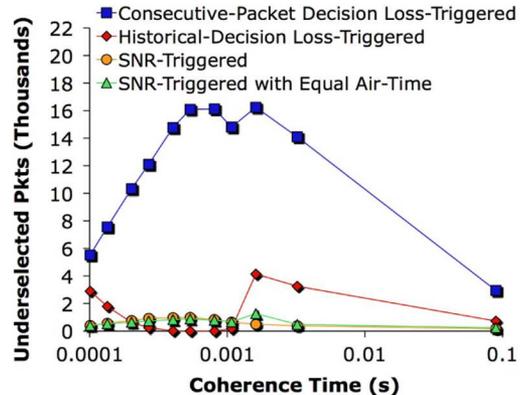


Fig. 5. Total number of underselected packets per protocol as a function of the coherence time on a Rayleigh fading channel.

Cause of Loss-Based Underselection. The first cause of low throughput is rate underselection by the consecutive-packet decision in the class of loss-triggered protocols. To show this effect from the prior experiment, we focus on the packets that are received correctly, yet are below the ideal rate of the channel from each of the four mechanisms now on a channel with medium average quality (-55 dBm).

Fig. 5 shows underselected packets in thousands (for 60-s test) as a function of the coherence time for each of the four mechanisms. For each protocol other than the consecutive-packet decision loss-triggered protocol, the number of packets received correctly below the ideal modulation is low (less than 4 k). However, for the loss-triggered protocol using the consecutive packet decisions, the total number of underselected packets increases as the fading increases on the channel up to 16 k packets at the packet transmission time of BPSK (2 ms) and then steadily decreases. The increase is due to the inability of the protocol to successfully transmit 10 consecutive packets (required to achieve a rate increment) and the increased likelihood of two consecutive failures. The unexpected decrease for coherence times greater than the packet transmission time is due to both the ideal modulation rate decreasing (the underselected rate now becoming the appropriate choice) and the number of undecodable headers increasing (thereby reducing the total decodable headers considered). *Finding: Loss-triggered protocols underselect from the ideal rate in fast-fading environments due to the consecutive-packet decision mechanism. We later show both loss-triggered mechanisms additionally suffer from packet failures due to various environmental factors.*

Cause of SNR-Based Overselection. The second cause of low throughput in a fast-fading scenario is the increased overselection of the SNR-triggered protocols. We revisit the prior experiment (Fig. 3) to consider the number of corrupt payloads received that are transmitted with a modulation rate that is above the ideal rate for the channel.

Fig. 6 shows the total overselected packets (in thousands) as a function of the coherence time. We observe that none of the protocols overselect as the coherence time increases to 1 ms (twice the packet period for 16QAM) since each of the protocols is able to transmit at the highest modulation. However, as coherence time decreases, SNR-triggered protocols transmit at a modulation rate greater than the ideal for up to 36 k packets,

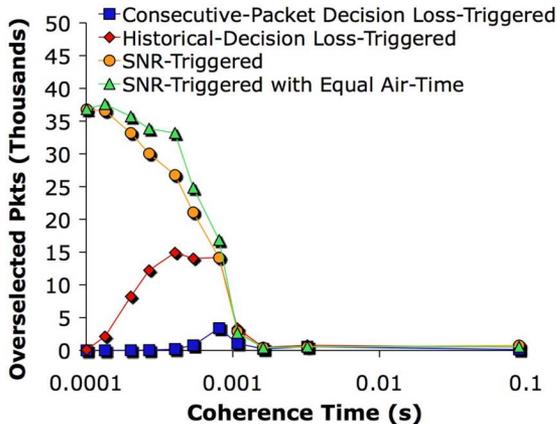


Fig. 6. Total number of overselected packets per protocol as a function of the coherence time.

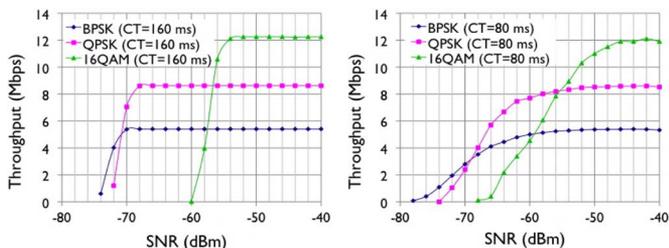


Fig. 7. Performance of modulation schemes with a coherence time of (left) 160 and (right) 80 ms.

much greater than for loss-triggered protocols. This overselection is due to the SNR-triggered protocols deriving their rate decisions from SNR-rate relationships for a channel with long coherence times. *Finding: SNR-triggered protocols overselect from the ideal rate due to coherence time sensitivity.*

C. Coherence Time Training

In situ Training for SNR Protocols. Because we found that SNR protocols are not robust to a broad range of coherence times, here we explore training SNR-based protocols. Training refers to obtaining the SNR-rate profile for the mobile node's actual operating environment, thereby incorporating the environment's coherence time. To explore training in a controlled environment, we begin by holding the coherence time constant on a single Rayleigh fading channel (no multipath present) and vary the SNR across the full range of allowable received power for the WARP radio board (-80 to -40 dBm). We repeat the experiment for many different coherence times (induced by speeds of 0.9, 25, 50, 75, 100, 150, 200, 400, and 800 km/h). For each coherence time, SNR, and modulation rate, we measure the achieved throughput.

Fig. 7 shows the achieved UDP throughput (Mb/s) as a function of the SNR (dBm) for each of the WARP modulations with a coherence time of 160 (left) and 80 ms (right). In Fig. 7 (left), we observe that the modulations have the highest throughput in the following regions: -57 dBm and higher (16QAM), from -71 to -57 dBm (QPSK), and less than -71 dBm (BPSK). These SNR thresholds correspond to the ideal rate. In contrast, Fig. 7 (right) indicates that the SNR thresholds shift to the right, requiring two additional decibels for the same modulation rate. Therefore, if the rate adaptation protocol was making a decision based upon

TABLE I
COHERENCE TIME AND SNR NECESSARY FOR RATE INCREASE
TO IDEAL MODULATION RATE

Speed	Coh.Time	QPSK \rightarrow 16-QAM	BPSK \rightarrow QPSK
Static	160 ms	-57 dBm	-72 dBm
0.9 kph	80 ms	-54 dBm	-72 dBm
25 kph	3.2 ms	-52 dBm	-72 dBm
50 kph	1.6 ms	-51 dBm	-72 dBm
75 kph	1.1 ms	-46 dBm	-72 dBm
100 kph	0.8 ms	0 dBm	-72 dBm

a longer coherence time than reality, the protocol would tend to overselect, resulting in loss. In the right-hand side figure, observe that BPSK has a larger tail at poor channel conditions due to the increased channel fluctuations (and occasional acceptable conditions) compared to the infinitely coherent channel in the left-hand side figure. Table I shows these thresholds separating the SNR regions for ideal rate on a given coherence time. All of the rows in Table I represent channels with a Rayleigh fading distribution according to the velocity in k/h. The static case (160 ms) and the case with a coherence time of 0.8 ms are the two extremes, i.e., there is no lower coherence time for which 16QAM should be used for coherence time values of less than 0.8 ms. While we did not directly measure the effect of convolutional codes, a similar SNR to throughput relationship is seen in [10, Fig. 9].

Per our findings in Figs. 4 and 6, we now quantify the gains of retraining an SNR-triggered protocol according to different coherence times. To achieve this, we first change the implementation of the two aforementioned SNR-triggered protocols (with and without equal airtime assurance) to make a rate decision based on Table I and measure the achievable throughput over the range of coherence time from our previous experiments (100 μ s to 100 ms).

Fig. 8 depicts the achievable throughput from each of the SNR-triggered protocols as a function of coherence time. For short coherence times, the SNR-triggered protocols trained at a coherence time of 0.8 ms achieve approximately 3 Mb/s more than the SNR protocols that are trained with static channels. Conversely, for long coherence times, the protocols trained at 0.8 ms underselect and have an achievable throughput of 3 and 1 Mb/s less than the statically trained, SNR-triggered protocols with and without equal airtime assurance, respectively. *Finding: SNR-based protocols can obtain significant throughput gains by incorporating coherence time training into the modulation rate decision thresholds.*

Coherence Time Detection and Implementation. If additional bits are required to be sent to measure the channel coherence, the gains of jointly considering SNR and coherence time will quickly be negated. In fact, such overhead would greatly exceed these gains. Thus, we now describe how to dynamically use both instantaneous and long-term (coherence time) SNR information at the receiver to achieve the aforementioned gains in modulation rate adaptation without injecting additional overhead. In the WARP physical layer, RSSI can be known throughout the packet as opposed to just at the beginning of the packet from the Automatic Gain Control. Furthermore, there is a noise level of the radio from which the SNR can be calculated. Recall with the carrier sense mechanism, there is a counter of how many samples have passed since the channel has been below a certain threshold or idle. Thus, in a similar fashion, there could be a running counter according to how many 10-MHz

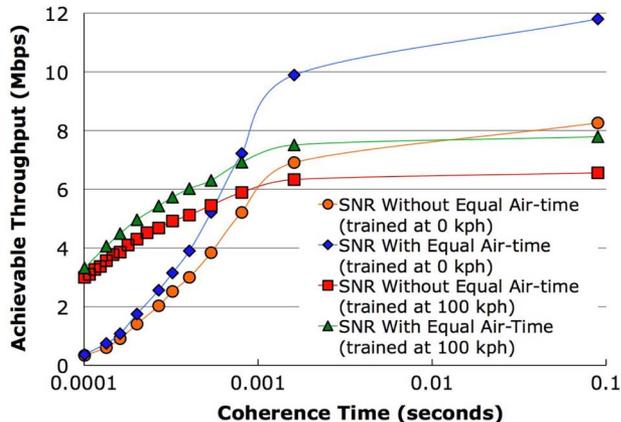


Fig. 8. Achievable throughput of SNR-triggered protocols before and after coherence time training.

samples (100-ns granularity) since the link quality crossed the SNR thresholds for each modulation rate. By averaging over these durations, the coherence time can be measured with the data that is already being sent over the link (i.e., a coherence time measurement could be performed without injecting additional bits, which induces overhead on the link). The hardware resources necessary to implement the mechanism is minimal.⁴ Furthermore, since the logic is only used while receiving a packet, the power consumption of the logic is infinitesimal compared to the act of receiving a packet. The receiver then uses this long-term characterization of the channel (coherence time) with the instantaneous SNR measured from the RTS packet to select the modulation rate. The rate choice would then be sent in the CTS packet according to the typical operation of the protocol.

However, though we have shown that there are gains to be had by considering instantaneous and long-term SNR, there are factors that prevent SNR from directly mapping to whether a packet is received correctly. One notable factor is the carrier frequency offset, which is the difference in carrier frequency between the transmitter and receiver. Mobility causes modulation constellations to spin, making the offset difficult to track. Thus, there are certain (high) vehicular speeds that the carrier frequency offset recovery in WARP is expected to fail. In this case, both measurements of SNR (instantaneous and long-term) would not result in modulation rate adaptation that maximizes the throughput, and it would be best to take a loss-triggered approach where the modulation is dropped to the lowest rate to maintain communication. To do so, the existing cyclic redundancy check (CRC) could be used since it is a good indicator of packet success or failure.

For devices that can make such an SNR-decision but are not able to implement the averaging at the physical layer, a context-aware approach could be taken. Namely, increasingly devices are multifunctional with cellular interfaces and accelerometers in addition to 802.11. Thus, context information could be used about a cellular tower location and movement detected by an accelerometer that could map to typical coherence time values. Such an approach would require training in many different environments such as indoor, downtown, and residential

⁴For example, the number of FPGA slices used for the carrier sense mechanism is 36 out of the 21 223 slices used for the entire WARP design (0.17%).

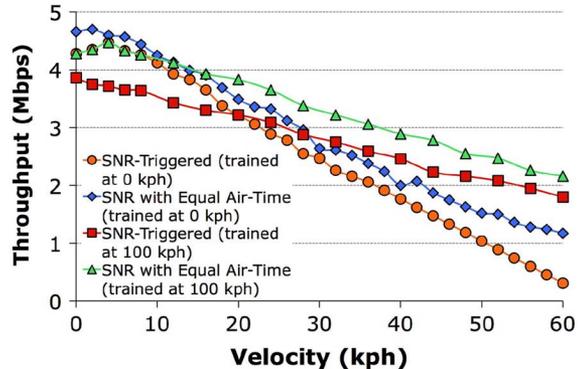


Fig. 9. Throughput with multipath-induced fading for SNR-based protocols.

to understand typical coherence time values. These measurements would be similar to what we have performed, but in more scenarios and with additional sample velocities to form the mapping. The aforementioned schemes require training that contrasts an online threshold adjustment proposed in [11]. While the online schemes reduce the amount of training, they must ensure robustness to collision/fading differentiation and interference from devices that do not obey the IEEE 802.11 protocol.

D. Multipath and Interference Effects

Our last PHY operating condition experiments evaluate rate adaptation accuracy with multipath and interference.

Multipath-Induced Fading. As discussed in Section VI, prior studies have shown multipath to be a dominant effect in the packet delivery ratio of a particular modulation rate and has been shown to lead to increased small-scale fluctuations [12], [13]. Thus, we evaluate rate adaptation accuracy within this context. To achieve this, we use the prior experimental setup (Fig. 2) with multiple Rayleigh channels where multipath delay is present. We use the case where five Rayleigh channels from the Commercial A setting set forth by JTC [14] with an RMS delay spread of 55 ns.

Fig. 9 shows the achievable throughput as a function of speed for the SNR-triggered protocols with and without training. At speeds of less than 10 km/h, the SNR-triggered protocol with equal airtime assurance that was trained at 0 km/h has the highest throughput. However, at only 10 km/h, the protocol that is trained at 100 km/h becomes the highest performing, showing that multipath has shifted the speed of the most appropriate coherence time training. *Finding: When multipath is present in the environment, the coherence-time sensitivity of SNR-triggered mechanisms is increased, and in situ training becomes more critical, even at lower vehicular speeds.*

Interference from External Devices. We now investigate rate adaptation accuracy with interference from undecodable noise sources since the spectrum is populated by numerous devices such as cordless telephones, microwaves, and other networks. In our experiment, we use a slow-fading channel with packet-sized noise (2 ms) with varying idle times.

Fig. 10 depicts the total number of packets underselected (for 60-s test) as a function of the idle period between bursts of noise. We find that as the idle period shortens, the consecutive-packet decision protocol increases in underselection up to 2 ms. At that point, the number of packets that are successful greatly decreases, and the transmitter sends fewer packets due

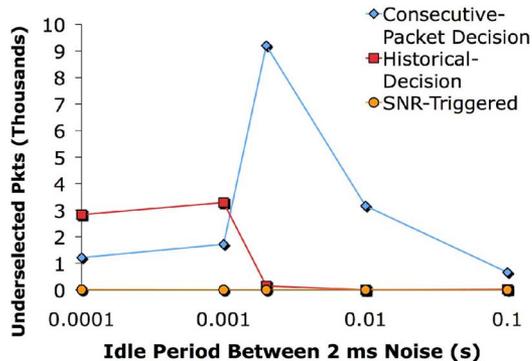


Fig. 10. Underslected packets by the loss-triggered protocols as interference is injected.

to long backoff intervals. The historical-decision mechanism is less susceptible, but where the number of losses are at the threshold values of the protocol to reduce rate, the protocol underselects. SNR-triggered protocols (only one shown since both are nearly identical) have lower overall throughput due to interference, but the rate decisions remain accurate. Thus, like the fast-fading scenario, interference causes loss-triggered protocols to underselect. *Finding: Interference forces both loss-triggered and SNR-triggered protocols to have lower throughput, but additionally causes the loss-based mechanisms to alter future rate decisions and underselect.*

IV. IN-LAB EVALUATION UNDER HETEROGENEOUS LINKS

In this section, we evaluate the accuracy of rate adaptation when transmitters are competing for bandwidth and have differing channel qualities among nodes. We study this heterogeneity in link quality at a mutual receiver for the case where competing transmitters are out of range or within range.

A. Heterogeneous Links

Differences in link qualities can exist among forwarding links (those selected by the routing algorithm) and nonforwarding links (not selected by the routing algorithm). For heterogeneous *nonforwarding* links, the behavior and coordination of competing transmitters depend on whether neighbors are in range (as depicted in Fig. 11, left) and can decode header packets from each other or are out of range and can neither decode headers nor sense channel activity from each other. (The latter case is a hidden terminal scenario and is depicted in Fig. 11, right.) For heterogeneous *forwarding* links, a wide range of link qualities can exist, e.g., links 1 and 2 in Fig. 11 can vastly differ in quality. In the most extreme case, physical layer capture has been shown to occur when the stronger transmission of two simultaneously transmitting terminals can be correctly received at a mutual receiver (e.g., the quality of link 1 \gg link 2 resulting in capture of A 's packets over C 's packets at node B).

Scenarios. As discussed in Section VI, although heterogeneous forwarding links with hidden terminals and the physical layer capture effect are common scenarios in practice [15], neither scenario has been explored with respect to rate adaptation. To explore these effects, we use the topologies depicted in Fig. 11, in which both scenarios have A and C sending data to B via links 1 and 2, respectively, and the transmitters are either in range (left) or out of range (right). The experiments are

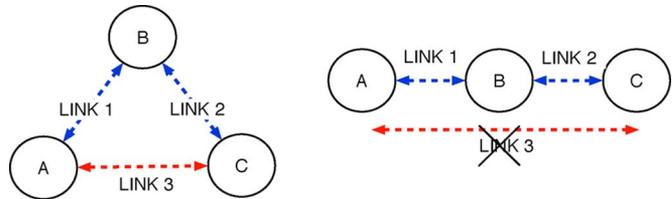


Fig. 11. Two scenarios with and without links from node A to node C , varying links 1 and 2.

performed with WARP boards using 3-dBi external antenna. We vary the qualities of forwarding links (1 and 2) by adjusting the transmit power and force the nonforwarding link (3) to be out of range by placing an obstruction along the direct path from A to C . The traffic pattern is 1500-byte UDP traffic with constant bit rate and 20-Mb/s offered load.

B. Hidden Transmitters With Heterogeneous Forwarding Links

Collision- and Fading-Based Loss with Heterogeneous Competing Links. Recall that reference [6] considered hidden terminals with *homogeneous* links, whereas we now explore the joint effect of hidden terminals (Fig. 11, right) and heterogeneous forwarding links on rate adaptation, a scenario not previously studied. We evaluate the accuracy of rate adaptation protocols for hidden terminals with heterogeneous forwarding links by varying their relative link quality (i.e., links 1 and 2 in Fig. 11, right).

To achieve this, we first establish a hidden terminal scenario where the channel from each of the transmitters to the mutual receiver is of high quality (-45 dBm). We next measure the achievable throughput of each of the simultaneously active flows ($A \rightarrow B$ and $C \rightarrow B$). We then repeat the experiment and hold the link quality of Node A constant (-45 dBm) and lower the link quality of Node B transmitter in steps of 5 dB up to 20 dB. Note from Fig. 7 (left) that in the region of -45 to -55 dBm, the highest modulation rate should be chosen.

Table II contains the per node throughput (kb/s) for each of the link differences (dB) per protocol, for each of the four mechanisms: historical-decision loss-triggered with and without collision/fading differentiation and SNR-triggered with and without equal airtime assurance (see Section II-C). Where equivalent links exist (0 dB), all protocols nearly perfectly achieve equal throughput sharing for the two flows. However, as the links increasingly differ in quality, the protocols obtain vastly different throughput sharing profiles. Consider the loss-triggered protocols with and without collision/fading differentiation: The protocol with the collision/fading differentiation mechanism has the highest aggregate throughput (throughput of Node A and Node C) across all differences in forwarding links, whereas the loss-triggered protocol without the mechanism has low aggregate throughput due to lack of protection from the four-way handshake to collision-based losses. While the aggregate throughput is highest with the mechanism, at a link difference of only 5 dB (which both transmitters should still be able to transmit at 16QAM), there is a 69%–31% throughput sharing as opposed to the 49%–51% sharing without the mechanism. In contrast, the SNR-based protocols have lower aggregate throughput but are able to more equally share the throughput despite the link quality differences. *Finding: Slight differences in link quality, even*

TABLE II
PERFORMANCE OF RATE ADAPTATION UNDER HETEROGENEOUS LINKS IN HIDDEN TERMINAL SCENARIO

Relative SNR	Loss-Triggered Rate Adaptation				SNR-Triggered Rate Adaptation			
	Without Collision Diff. Throughput (kbps - %)		With Collision Differentiation Throughput (kbps - %)		Without Equal Air-time Throughput (kbps - %)		With Equal Air-time Throughput (kbps - %)	
	Node A	Node C	Node A	Node C	Node A	Node C	Node A	Node C
0 dB	655 (50.0)	655 (50.0)	5211 (50.3)	5139 (49.7)	2888 (49.1)	2998 (50.9)	4039 (50.6)	3945 (49.4)
5 dB	665 (51.3)	631 (48.7)	3129 (30.7)	7050 (69.3)	2850 (49.4)	2919 (50.6)	3940 (48.7)	4155 (51.3)
10 dB	619 (48.1)	668 (51.9)	1480 (15.7)	7924 (84.3)	1741 (35.0)	3236 (65.0)	2038 (29.8)	4809 (70.2)
15 dB	214 (9.3)	2084 (90.7)	233 (2.6)	8651 (97.4)	1854 (43.3)	2429 (56.7)	2300 (40.1)	3440 (59.9)
20 dB	45 (0.7)	6295 (99.3)	917 (10.2)	8084 (89.8)	1272 (35.3)	2330 (64.7)	1890 (34.5)	3592 (65.5)

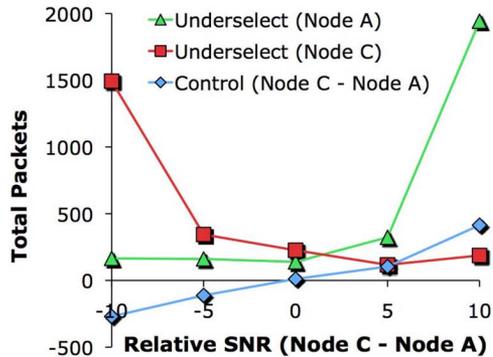


Fig. 12. Sensitivity to small differences in SNR of competing links with collision/fading differentiation.

in the same modulation rate region, cause collision/fading differentiation to have large differences in throughput sharing between competing hidden transmitters.

Origins of Throughput Sharing Imbalance for Collision/Fading Differentiation. Next, we evaluate the reason for the large differences in throughput sharing of the collision/fading differentiation mechanism. To do so, we use the rate adaptation accuracy statistics of the above result within the region where the two competing, hidden transmitters are able to transmit at the highest modulation rate.

In Fig. 12, Node *C* has a relative SNR with respect to Node *A* from -10 to 10 dB at the mutual receiver (Node *B*). The total packets underselected by *A* and *C* are shown (y-axis) as well as the difference of transmitted control messages (Node *C* minus Node *A*). Even at a difference of 5 dB (*C* is the stronger transmitter), the number of control messages sent by *C* is much greater than *A* (104 packets), thereby removing some of the protection for data packets from *A*. This can be explained by the increase in probability of short, low-rate control packets of *C* being captured over the control packets of *A*. Thus, *C* is less likely to back off when RTS packets overlap in time, thereby decreasing the probability that *A* can complete an RTS/CTS exchange (when the links are equivalent both back off with equal probability and this problem does not manifest). Considering the operation of collision/fading differentiation, after a successful RTS/CTS exchange, either the protocol halves or decrements the RTS window based on whether the data is lost or successful, respectively. In either case, both nodes would have periods where the RTS/CTS is not used. However, the weaker transmitter using collision/fading differentiation would consider a smaller relative number of these lost data packets sent when RTS/CTS is not used, thereby selecting a rate below the ideal rate (i.e., underselection occurs). *Finding: The exaggeration of*

slight differences in link quality of the collision/fading differentiation mechanism is due to unequal channel accesses of the four-way handshake, favoring the flow that uses the four-way handshake.

Interaction of Capture and Rate Adaptation. Extreme heterogeneity in the forwarding links results in physical layer capture (even with data packets) and drives rate adaptation to previously unexplored behavior. In Table II, the loss-triggered protocol has an increase in throughput for the stronger transmitter when the relative SNR is greater than or equal to 15 dB. This increase in throughput is due to the stronger transmitter no longer experiencing consecutive losses from collision, revealing that physical layer capture occurs even for data packets at these SNR differences. In our measurements, we find that although capture occurs almost completely with 1500 -byte data packets with a 20 -dB difference in SNR, the delivery ratio of RTS messages is reduced by only 10% . This is due to exponential backoff within the MAC allowing sufficient spacing for a small-sized control message.

In addition to the aforementioned imbalance of the loss-triggered mechanism, note that even with collision/fading differentiation, there is a 90% share taken by the stronger transmitter due to the imbalanced use of the four-way handshake (965 more RTS packets for the stronger transmitter). The SNR-triggered protocols, however, have the expected distribution according to the appropriate rate choice (approximately $2:1$ from the modulation rate performance depicted in Fig. 7 with SNR of -45 and -60 dBm). One might expect that turning on RTS (thereby ensuring equal use per transmitter) would allow rate adaptation to be robust to physical layer capture. However, in the same scenario, if RTS is enabled for loss-triggered rate adaptation, the throughput distribution is 6.1 Mb/s (86%) and 945 kb/s (14%) since the RTS losses trigger a lowering of the modulation rate. Although there is sufficient spacing for the RTS packet of the capture-losing transmitter to fit within the exponential backoff window of the capture-winning transmitter, the four-way handshake only provides protection for the rate adaptation algorithm to physical layer capture when the RTS messages do not trigger the channel-state interpretation of the protocol. *Finding: The joint interaction of rate adaptation and the physical layer capture effect causes significant reductions in throughput for the capture-losing node, which can be avoided if RTS losses are independent of rate selection triggering.*

C. In-Range Transmitters With Heterogeneous Forwarding Links

Competing Multirate Links with Ability to Carrier Sense. As [16] showed in the case of in-range heterogeneous forwarding links, low-quality links can cause even high-quality

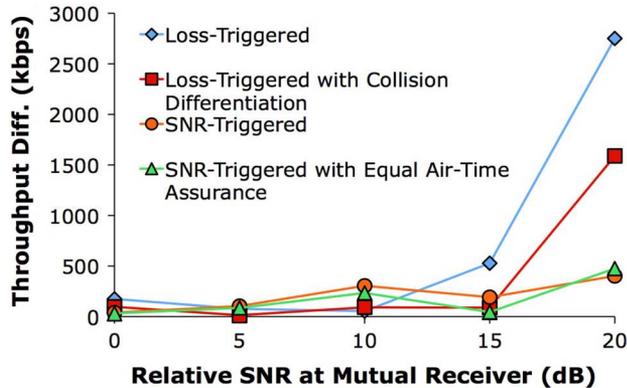


Fig. 13. Throughput difference of in-range flows varying relative SNR between senders (A minus C).

links to yield low throughput. Here, we investigate the performance of the aforementioned protocols in such a scenario. To evaluate this issue, we repeat the same heterogeneous forwarding link experiment as before, but with in-range transmitters.

Fig. 13 shows the difference in achieved throughput (kb/s) for the two flows as a function of the relative SNR (dB). The difference in achieved throughput is negligible for SNR differences of less than 15 dB. At link quality differences greater than 15 dB, the throughput that the flows achieve diverge by approximately 2 Mb/s for the loss-triggered protocols, which is according to expectation caused by the modulation rate change (difference divided by 2). The SNR-triggered protocols differ by approximately 500 kb/s since the RTS overhead reduces the relative difference of the two flows with different rates since part of the time is used for base-rate transmissions by both senders. *Finding: The transmitters-in-range scenario: 1) does not suffer from the sensitivity of the collision/fading differentiation under heterogeneous forwarding links; and 2) the severe throughput imbalance caused by the physical layer capture effect within the hidden terminal scenario does not occur due to the lack of simultaneous transmissions from increased coordination (virtual and physical sensing) of in-range competing transmitters.*

V. RESIDENTIAL URBAN AND DOWNTOWN SCENARIOS

In this section, we perform experiments in a residential urban environment consisting of dense foliage and homes and a commercial downtown environment having strong multipath due to closely set buildings. These scenarios enable evaluation of rate adaptation protocols in outdoor environments similar to those encountered in large-scale wireless deployments—scenarios that can have increased fading, delay spread, and interference over indoor networks. We characterize these environments with measurements from the cross-layer implementation and study rate adaptation accuracy for both mobile and nonmobile sender–receiver pairs.

A. Residential Urban and Downtown Experiment Design

In these two outdoor environments, we evaluate the combination of physical-layer operating conditions and heterogeneous link factors tested independently in Sections III and IV. Thus, we first characterize these environments for perspective on the

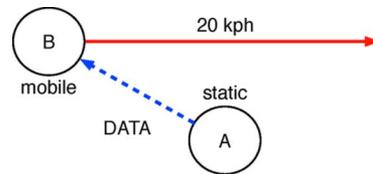


Fig. 14. Two-node scenario for mobility in both environments, static sender to mobile receiver.

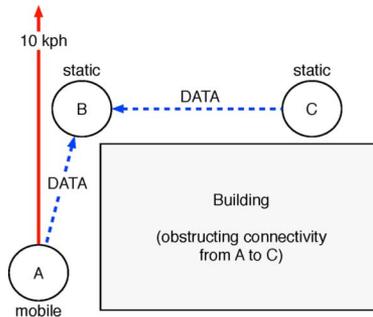


Fig. 15. Heterogeneous links in residential urban environment.

experiments in prior sections, and then evaluate the rate adaptation accuracy within these scenarios.

Scenarios. The following three scenarios are used to explore the environments. First, a pair of static nodes is used to measure channel conditions and to test rate adaptation accuracy in such conditions. Then, a mobile topology with two nodes depicted in Fig. 14 is used where a static node (A) sends data to a mobile node (B). The mobile node starts moving with a nominal link (i.e., the link has greater than 90% packet loss at a distance of approximately 200 m), passes the static node, and continues until the link becomes nominal again, evaluating the ability of the protocols to ramp up and down the rate. The last scenario is similar to the heterogeneous link topology (both in-range and hidden) in the previous section where two outside nodes (A and C) are contending to a middle node (B), except here, one of the outside nodes (A) is mobile as depicted in Fig. 15. Note that collision/fading differentiation is not used until the third scenario (hidden terminals).

Our residential urban measurements are performed within a densely populated, single-family residential neighborhood with dense foliage as measured in [17]. The downtown measurements are performed on streets in Houston, TX, where buildings tens of stories in height line each side. Measurements were performed for scenarios in which one or both sides of the street have no such structures, but are not presented here due to space constraints. In all experiments, nodes are placed inside the vehicle, and a 3-dBi antenna is mounted on the vehicle roof at a height of 2 m. Before each experiment, we first measure the coherence time of the scenario. To do so, we generate UDP traffic of varying packet sizes and record the SNR variance between different windows of packets to determine the coherence time of the channel.

B. Impact of Environment (Static Pair)

Measurements of the outdoor propagation environment allow comparison with in-lab experiments and understanding of the

performance of the hardware (i.e., delay spread tolerance, modulation performance, etc.). Multipath and delay spread are potential factors that effect performance as in [18], for example, the packet delivery performance was shown to be somewhat uncorrelated with SNR, and the authors concluded that a strong multipath effect is the cause.

Raw Characterization of Outdoor Environments. We first measure the channel characteristics of the environments with a static pair of nodes. In both the residential and downtown environments, vehicles pass at approximately 30 km/h within 5 m from the location of one of the nodes in the experiment. On average, we find the coherence time to be 100 and 80 ms within the residential urban and downtown scenarios, respectively (the duration of tens of packets). However, we find that passing cars can drive the coherence time of the residential urban scenario to as low as 15 ms and the downtown scenario to as low as 300 μ s (less than one packet transmission time). *Finding: Despite the static topology, the coherence time can be as low as 300 μ s, which, in an idealized propagation environment, corresponds to a velocity of 250 km/h.*

Rate Adaptation Accuracy in Outdoor Environments. Next, we measure the performance of multirate mechanisms within both residential urban and downtown scenarios. In our experiment, we generate UDP traffic from sender to the receiver for each of the rate adaptation mechanisms at a distance of approximately 100 m in both environments. We tested differing ranges for the maximum reach of the transmitters while still being able to transmit packets successfully at the highest modulation rate and the highest transmit power. We record the per-packet variance of SNR to measure the fading of the channel during the experiment. We note that the average SNR between sender and receiver in the downtown case is 10 dB stronger than the residential urban scenario.

Fig. 16 shows the results from the residential urban scenario (left) and the downtown scenario (right). The total number of packets at the receiver (y-axis) are depicted according to whether they are underselected, accurate, and overselected during the test (60 s). Each of the four mechanisms are on the x-axis in the following order, triggered by: 1) consecutive-packet decision; 2) historical decision; 3) SNR; and 4) SNR with equal airtime assurance. In the residential urban scenario (Fig. 16, left), the consecutive-decision mechanism largely underselects, while the historical-decision mechanism largely overselects. However, the SNR-triggered protocols are highly accurate. Although the coherence time is multiple packets in duration, the consecutive-decision mechanism underselects since the mobility of scatters (when present) disallows the required 10 consecutive successful packets to raise the rate and is further reduced from sources of loss not present within the indoor setting. However, the historical-decision mechanism overselects from the ideal rate since the outdoor modulation rates achieve different delivery ratios than in the indoor setting (where the WARP modulations can achieve the throughput with which the loss thresholds of the historical-decision mechanism were established). Both loss-triggered inaccuracies reveal that these protocols, although widely used in practice in outdoor settings, are tuned for indoor channel types. *Finding: Consecutive-packet decision and historical-decision loss-based protocols are largely inaccurate at adapting the rate within a practical outdoor setting.*

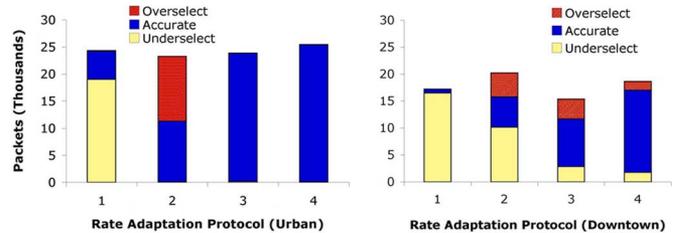


Fig. 16. Rate adaptation accuracy with static pair in (left) residential urban and (right) downtown.

In the downtown scenario (Fig. 16, right), recall that the average coherence time in this scenario is 80 ms, but driven as low as 300 μ s when cars pass. Along the experienced range of coherence times, when the coherence time is approximately equal to the packet transmission time (2 ms), we observe the effect shown in Section III in which the modulation rate is unable to perform given the short coherence time. The result is overselection for SNR-triggered protocols and underselection for loss-based protocols due to excessive loss triggers. For the lowest coherence times, the duration of 300 μ s is even smaller than the turnaround time of the RTS/CTS exchange. Consequently, an SNR-based decision at the time of the RTS reception is stale by the time of the data packet reception, resulting in underselection by SNR-triggered protocols. *Finding: Even in the static scenario, the short coherence time caused by the mobility of scatterers forces SNR-triggered protocols to both under- and overselect and forces loss-triggered protocols to underselect due to effects analyzed in Section III.*

We also observe that in the downtown scenario of Fig. 16 (right), all of the mechanisms have a much lower number of successfully received packets than for the residential urban scenario. Despite the better link quality (10 dB higher), the performance of the protocols is driven lower due to a strong multipath component in the downtown scenario that is not present within the residential urban environment. *Finding: Multipath is a dominant effect in rate adaptation that drives throughput lower within the downtown scenario, but not in the residential urban scenario.*

C. Impact of Mobility

Tracking Channels Under Vehicular Mobility. We now evaluate rate adaptation accuracy within the same two environments with the increased fading and channel quality changes that occur with mobility. To achieve this, we position a node statically on the side of the street and measure the achievable throughput to a mobile node that approaches and passes on the same street at 20 km/h, as depicted in Fig. 14. We track the per-packet variance of the SNR to measure the channel fading during the experiment.

For the residential urban scenario, Fig. 17 depicts normalized throughput as a function of time as the receiver is moving toward then away from the sender. All rate adaptation mechanisms increase rate as the receiver approaches the sender and decrease after passing the sender. However, the SNR-triggered protocols have much longer periods (4 s) of normalized throughput close to 1 as compared to the short-duration spikes of lower peak value for the loss-triggered protocols. Thus, the loss-triggered

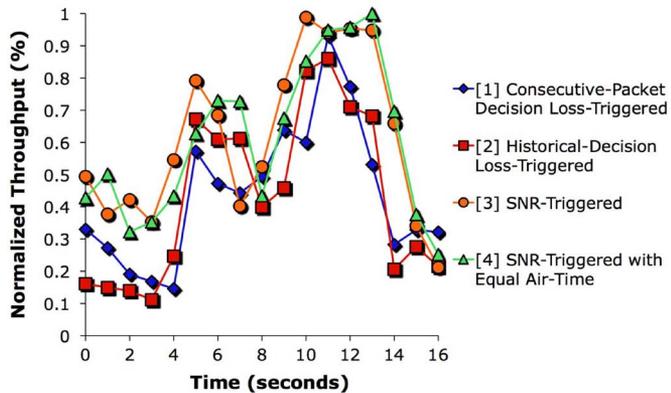


Fig. 17. Normalized throughput (from max value in environment) for each of the multirate protocols within a residential urban setting.

protocols are not able to track the mobile client, even at relatively low speed. (Similar results hold for downtown but are not shown here.) *Finding: Sequential rate stepping and the sensitivity to failure of the loss-triggered protocols cannot track mobile environments, but SNR-based protocols are able to accurately adapt.*

Mobility with High Levels of Interference. Next, we evaluate the combined effect of interference and fast-fading on rate adaptation accuracy within the aforementioned mobile scenario. To do so, we compare the performance of the rate adaptation accuracy when the two nodes are isolated on a channel (i.e., no other devices cause interference) from results depicted in Fig. 17, and then when the two nodes are on the same channel as the TFA network containing 20 access points (i.e., interference induced by beacons and traffic on the same channel but undecodable to our sender–receiver pair). We confirm the activity of user and backhaul traffic on the TFA network by *tcpdump* traces taken at the gateway mesh nodes during these experiments.

Fig. 18 contains the underselected, accurate, and overselected packets for each of the four mechanisms first without (left) and then with interference (right). We observe that interference causes poor rate decisions by the loss-triggered protocols (1 and 2, in the figures) as each of the first two protocols increase the number of underselected packets by at least 3 k packets from the left figure to the right figure. The rate decisions of the SNR-triggered protocols (3 and 4, in the figures) remain accurate, however the throughput is brought lower (less total packets) by the presence of interference since the available idle time is reduced. *Finding: Interference from external devices shortens the interval of available airtime, causing an increased number of packet failures, triggering both types of loss-based protocols to underselect. Furthermore, the shorter interval reduces the likelihood of a successful four-way handshake and drives the throughput of SNR-triggered mechanisms lower (although the rate decisions are accurate).*

D. Impact of Heterogeneous Links

Heterogeneous Links with Mobility in Residential Urban Scenario. Here, we evaluate the combined factors of heterogeneous links with mobility within an urban environment. To achieve this, we use the topology pictured in Fig. 15, which is a dynamic version of each of the scenarios in the previous section:

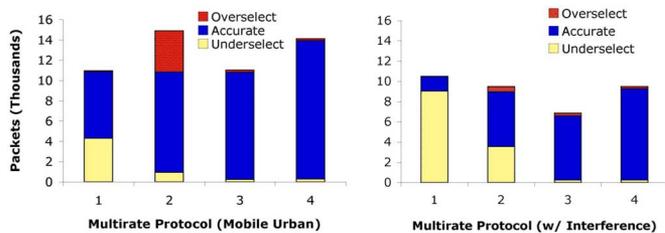


Fig. 18. Rate adaptation accuracy with mobile scenario in residential urban setting (left) without and (right) with interference.

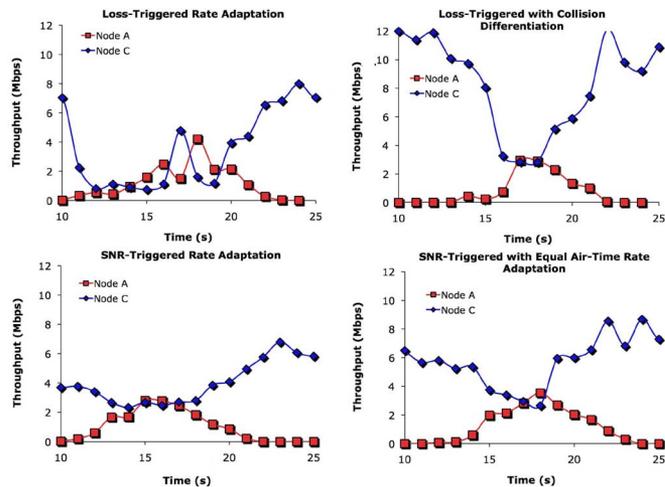


Fig. 19. Throughput of each of the four protocols in the heterogeneous links topology (Fig. 15).

hidden terminals and in-range terminals with heterogeneous forwarding links. We expect that the static node would have an advantage over the mobile node since the link’s channel conditions do not suffer from both the degree of fading of a mobile link and the longer-term changes of link quality from nominal to good. Nonetheless, the mobile node reaches a physical location within the experiment that has better channel conditions to the receiver than the static node.

Each of the four graphs in Fig. 19 depicts throughput for both of the contending transmitters (*A* and *C*) over time. In the left half of each of the graphs, the two transmitters are hidden from one another, and approximately halfway through the experiment, they become in-range. For the consecutive-packet decision loss-triggered protocol without collision/fading differentiation (top left), the aggregate throughput is low in the hidden terminal scenario, then highly fluctuates once the transmitters are in-range. The historical-decision loss-triggered protocol with collision/fading differentiation (top right) has difficulty with heterogeneous links, as previously observed in Section IV, and the mobile node (*A*) has only a small share of throughput until the links become equivalent. The SNR-triggered protocol (bottom left) has the longest period of equivalent throughput of the four protocols. Lastly, for the SNR-triggered protocol with equal airtime assurance (bottom right), as the mobile link becomes higher quality than the static link, the mobile link sends back-to-back packets in proportion to the selected modulation rate over the base rate. *Finding: In-lab experiments can predict the outdoor behavior of the rate adaptation mechanisms*

with heterogeneous links from competing in-range and hidden transmitters.

VI. RELATED WORK

Rate Adaptation Protocol Design. There are two classes of rate adaptation mechanisms that have been developed that differ in how they determine the appropriate physical layer rate according to the perceived channel state. These schemes can be classified into loss-triggered [1]–[6] and SNR-triggered [7], [8] multirate protocols. *Loss-triggered* protocols are the most commonly implemented due to their transmitter-based simplicity. These protocols use sequential rate stepping based upon either consecutive successes and failures [3] or delivery statistics over a window of time based upon historical performance of the modulation rates [1], [2], [6]. *SNR-triggered* protocols infer channel state at the receiver based upon signal strength of control messages. The transmitter then either sends data packets per RTS/CTS exchange [7] or sends a burst of data packets in proportion to the modulation rate over the base rate for time-share fairness [8]. [19] implements the protocol proposed in [7] with a software-defined MAC and PHY resulting in a turnaround time of multiple *ms*. However, our custom cross-layer implementation (operating within the fabric of the FPGA and embedded processor) operates at MAC time scales comparable to commercial systems (300 μ s turnaround time), allowing accurate SNR-based rate adaptation even in fast-fading channels and overhead according to prior expectations from the RTS/CTS exchange. Hence, in contrast to multirate protocol design and prior software-based implementation, we focus on the cross-layer implementation and evaluation of the key rate adaptation mechanisms in a large class of scenarios and topologies [20]. As discussed in Section III-C, context-awareness [21] (and channel reciprocity discussed in [11]) can further increase the gains we showed by SNR-based protocols.

Evaluation of Rate Adaptation. Prior work has investigated the effectiveness of rate adaptation protocols via throughput comparison. The issues investigated have been a fast-fading channel performance comparison of two protocols [7], [22], collision/fading differentiation with hidden terminals indoors with off-the-shelf hardware [6], equal airtime assurance for SNR-triggered rate adaptation simulations where the coherence time is assumed to be much greater than the packet period [8], or single-active, one-hop flow performance of loss-triggered protocols compared on a mesh deployment [2]. However, in our study, we are able to evaluate a broad set of rate adaptation mechanisms on a per-packet basis via the observability between the MAC and PHY layer, revealing the rate choices which lead to relative differences in throughput per protocol (i.e., our work is the first that is able to assess rate selection on a packet-by-packet basis).

PHY Operating Conditions. In prior work, the channel conditions considered when testing multirate protocols have been immobile sender and receiver in a predominantly line-of-sight outdoor environment [2], nonvehicular mobility in a simulator [7], [8], or indoor environment [6], [11]. In such scenarios, the channel fading is almost entirely isolated to the case where the coherence time of the channel is much greater than the packet period. However, the increase of citywide wireless networks and

other large-scale mesh networks such as [17] bring into question how these protocols perform when the coherence time approaches the packet period, whether with vehicular speeds or a high mobility of scatterers within the environment. Here, we evaluate the protocols on a broad set of emulated channel conditions including fast-fading, interference, and multipath and then test the protocols in urban and downtown settings with these conditions.

Heterogeneous Links. Heterogeneous links have been shown to cause problems in rate adaptation in the following contexts: 1) a weaker (i.e., more distant) transmitter consumes a majority of the airtime and causes the stronger transmitter to have reduced rate [16]; and 2) a hidden terminal scenario where loss-triggered protocols misinterpret collision-based losses as channel-state-based losses, erroneously reducing the selected rate. Within the latter context a dynamic enabling of the RTS mechanism has been shown to mitigate the misinterpretation of the cause of loss [6]. However, within such a hidden terminal scenario, only homogeneous competing links have been explored which, within the context of a deployed wireless network, is not the norm [15]. Therefore, we evaluate how accurate the rate adaptation mechanisms are with heterogeneous forwarding links within a hidden terminal scenario. Recent promising work on distinction of loss types at the physical layer [23], [24] has been leveraged to aide the performance of loss-based rate adaptation [25], yet the USRP platform currently used for the experiments precludes real-time evaluation at commercial time scales.

Moreover, with extreme heterogeneity in forwarding links, physical layer capture occurs, causing the MAC of the weaker transmitter to be subject to the performance of the physical layer. Since capture has been shown to occur with negligible differences in link quality [26], the effect is common within deployed networks [15], and it is necessary to consider the capture effect on rate adaptation accuracy, which we are the first to explore. For further discussion of the rate adaptation issues with heterogeneous links, refer to Section IV.

Residential Urban and Downtown. [18] concluded that while there was some correlation with link performance and SNR, multipath was the dominant effect in the MIT Roofnet network. Other mesh network studies have shown the correlation between SNR and link performance to be high [17]. We show that while the effect of multipath is severe in the downtown scenario, it is far less severe in the residential urban scenario. Furthermore, we find that the assumption of coherence time being much greater than the packet period [8] does not hold even in static topologies in downtown scenarios due to effects such as the mobility of scatterers.

VII. CONCLUSION

In this work, we developed a custom cross-layer rate adaptation framework that has high levels of interaction and observability between MAC and PHY layers. We are the first to implement SNR-based rate adaptation at MAC time scales comparable to commercial systems and evaluate protocol accuracy compared to optimal rate selection on a packet-by-packet basis. Using this cross-layer implementation, we found that loss-triggered mechanisms underselect in the presence of fast-fading and interference and are unable to track channel changes in mobile environments. Furthermore, we found that

in situ training of SNR-triggered protocols to overcome their coherence time sensitivity allows significant throughput gains. We show that even in static topologies, in practical outdoor environments, coherence time training is necessary. Finally, we show that a mechanism designed to equally share throughput in the hidden terminal scenario has a severe imbalance in throughput sharing with only slight heterogeneity in average link quality of competing transmitters. In future work, we will consider a MIMO dimension to the robustness versus throughput tradeoff as emerging systems such as IEEE 802.11n have tens of data rates to consider.

ACKNOWLEDGMENT

The authors would like to thank N. Warnes, J. Gulati, and C. Sequeira for their hard work in assisting with the experiments.

REFERENCES

- [1] "MadWifi project," Sep. 2005 [Online]. Available: <http://madwifi.org>
- [2] J. C. Bicket, "Bit-rate selection in wireless networks," M.S. Thesis, MIT, Cambridge, MA, Feb. 2005.
- [3] A. Kamerman and L. Monteban, "WaveLAN II: A high-performance wireless LAN for the unlicensed band," *Bell Labs Tech. J.*, pp. 118–133, Summer, 1997.
- [4] J. Kim, S. Kim, S. Choi, and D. Qiao, "CARA: Collision-aware rate adaptation for IEEE 802.11 WLANs," in *Proc. IEEE INFOCOM*, 2006.
- [5] M. Lacage, M. Hossein, and T. Turletti, "IEEE 802.11 rate adaptation: A practical approach," in *Proc. MSWiM*, Oct. 2004, pp. 126–134.
- [6] S. Wong, S. Lu, H. Yang, and V. Bharghavan, "Robust rate adaptation for 802.11 wireless networks," in *Proc. ACM MobiCom*, 2006, pp. 146–157.
- [7] G. Holland, N. Vaidya, and P. Bahl, "A rate-adaptive MAC protocol for multi-hop wireless networks," in *Proc. ACM MobiCom*, Rome, Italy, Jul. 2001, pp. 236–251.
- [8] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic media access for multirate ad hoc networks," in *Proc. ACM MobiCom*, Atlanta, GA, Sep. 2002, pp. 24–35.
- [9] A. Khattab, J. Camp, C. Hunter, P. Murphy, A. Sabharwal, and E. Knightly, "WARP: A flexible platform for clean-slate wireless medium access protocol design," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 12, no. 1, pp. 56–58, 2008.
- [10] H. Rahul, F. Edalat, D. Katabi, and C. G. Sodini, "Frequency-aware rate adaptation and mac protocols," in *Proc. ACM MobiCom*, Beijing, China, Sep. 2009, pp. 193–203.
- [11] G. Judd, X. Wang, and P. Steenkiste, "Efficient channel-aware rate adaptation in dynamic environments," in *Proc. ACM MobiSys*, Boulder, CO, Jun. 2008, pp. 118–131.
- [12] T. Rappaport, T. Rappaport, Ed., *Wireless Communications, Principles & Practice*, ser. Emerging Technologies. Upper Saddle River, NJ: Prentice-Hall, 1996.
- [13] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge, U.K.: Cambridge Univ. Press, 2005.
- [14] "Technical report on RF channel characterization and system deployment modeling," JTC (Air) Standards Contribution, Tech. Rep. JTC(AIR)/94.09.23-065R6, Sep. 1994.
- [15] J. Camp, V. Mancuso, O. Gurewitz, and E. Knightly, "A measurement study of multiplicative overhead effects in wireless networks," in *Proc. IEEE INFOCOM*, 2008, pp. 76–80.
- [16] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proc. IEEE INFOCOM*, San Francisco, CA, Apr. 2003, vol. 2, pp. 836–843.
- [17] J. Camp, J. Robinson, C. Steger, and E. Knightly, "Measurement driven deployment of a two-tier urban mesh access network," in *Proc. ACM MobiSys*, Uppsala, Sweden, Jun. 2006, pp. 96–109.
- [18] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11 mesh network," in *Proc. ACM SIGCOMM*, Portland, OR, 2004, pp. 121–132.
- [19] K. Mandke, S.-H. Choi, G. Kim, R. Grant, R. Daniels, W. Kim, R. Heath, and S. Nettles, "Early results on Hydra: A flexible MAC/PHY multihop testbed," in *Proc. IEEE Veh. Technol. Conf.*, Dublin, Ireland, Apr. 2007, pp. 1896–1900.
- [20] J. Camp and E. Knightly, "Modulation rate adaptation in urban and vehicular environments: Cross-layer implementation and experimental evaluation," in *Proc. ACM MobiCom*, San Francisco, CA, Sep. 2008, pp. 315–326.
- [21] P. Shankar, T. Nadeem, J. Rosca, and L. Iftode, "CARS: Context-aware rate selection for vehicular networks," in *Proc. IEEE ICNP*, Orlando, FL, Oct. 2008, pp. 1–12.
- [22] P. Chevillat, J. Jelitto, A. N. Barreto, and H. L. Truong, "A dynamic link adaptation algorithm for IEEE 802.11a wireless LANs," in *Proc. IEEE ICC*, Anchorage, AK, May 2003, vol. 2, pp. 1141–1145.
- [23] K. Jamieson and H. Balakrishnan, "PPR: Partial packet recovery for wireless networks," in *Proc. ACM SIGCOMM*, Kyoto, Japan, Aug. 2007, pp. 409–420.
- [24] S. Rayanchu, A. Mishra, D. Agrawal, S. Saha, and S. Banerjee, "Diagnosing wireless packet losses in 802.11: Separating collision from weak signal," in *Proc. IEEE INFOCOM*, Phoenix, AZ, Apr. 2008, pp. 735–743.
- [25] M. Vutukuru, H. Balakrishnan, and K. Jamieson, "Cross-layer wireless bit rate adaptation," in *Proc. ACM SIGCOMM*, Barcelona, Spain, Aug. 2009, pp. 3–14.
- [26] A. Kochut, A. Vasani, A. U. Shankar, and A. Agrawala, "Sniffing out the correct physical layer capture model in 802.11b," in *Proc. IEEE ICNP*, Berlin, Germany, Oct. 2004, pp. 252–261.



Joseph Camp (S'03–M'09) received the B.S. degree with honors in electrical and computer engineering (ECE) from the University of Texas at Austin in 2003, and the M.S. and Ph.D. degrees in ECE from Rice University, Houston, TX, in 2005 and 2009, respectively.

He is an Assistant Professor of electrical engineering with Southern Methodist University, Dallas, TX, which he joined in 2009 after received the Ph.D. degree. He was the Chief Network Architect from 2004 to 2009 for the Technology For All (TFA) Network, a mesh deployment in Houston, TX, which serves thousands of users in an underresourced community. His research interests are in the areas of wireless networking and embedded systems, specifically focused on deployment, measurement, and analysis of large-scale wireless networks and development of embedded protocols for network hardware.

Dr. Camp received the Ralph Budd Award in 2010 for the best engineering thesis at Rice University.



Edward Knightly (S'91–M'96–SM'04–F'09) received the B.S. degree from Auburn University, Auburn, AL, in 1991, and the M.S. and Ph.D. degrees from the University of California at Berkeley in 1992 and 1996, respectively.

He is a Professor of electrical and computer engineering with Rice University, Houston, TX. His research interests are in the areas of mobile and wireless networks and high-performance and denial-of-service resilient protocol design.

Prof. Knightly is a Sloan Fellow and a recipient of National Science Foundation CAREER Award. He received the Best Paper Award at ACM MobiCom 2008. He served as Associate Editor of numerous journals and special issues including the IEEE/ACM TRANSACTIONS ON NETWORKING and the IEEE JOURNAL ON SELECTED AREAS OF COMMUNICATIONS Special Issue on Multi-Hop Wireless Mesh Networks. He served as Technical Co-Chair of IEEE INFOCOM 2005 and General Chair of ACM MobiHoc 2009 and ACM MobiSys 2007. He has served on the program committee for numerous networking conferences including ICNP, INFOCOM, MobiCom, and SIGMETRICS.