

RICE UNIVERSITY

**Experimental and Analytical Evaluation of Embedded Link
Performance with Small-Scale Channel Fluctuations**

by

Joseph D. Camp

A THESIS SUBMITTED
IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE
DOCTOR OF PHILOSOPHY

APPROVED, THESIS COMMITTEE:

Edward W. Knightly, Chair
Professor of Electrical and Computer
Engineering

David B. Johnson
Professor of Computer Science and
Electrical and Computer Engineering

T. S. Eugene Ng
Assistant Professor of Computer Science
and Electrical and Computer Engineering

Ashutosh Sabharwal
Assistant Professor of Electrical and
Computer Engineering

Houston, Texas

May, 2009

ABSTRACT

Experimental and Analytical Evaluation of Embedded Link Performance with Small-Scale Channel Fluctuations

by

Joseph D. Camp

We have deployed a first-of-its-kind, urban-scale wireless mesh network which provides Internet access to 1000's of users spanning multiple square kilometers in an underserved area in Houston, TX. However, in this and other urban environments, IEEE 802.11 link performance is both misunderstood and poor-performing due to complex node interactions which are affected by a vast array of factors including topology, channel conditions, modulation rate, packet sizes, and physical layer capture. In this thesis, I draw from 100's of thousands of urban measurements and develop an analytical model to understand the performance of links embedded in the aforementioned complex scenarios. My focus is on two fundamental concepts involving embedded links. First, choosing the modulation rate which maximizes the throughput is imperative since each bit of the (overly-)shared medium is critical. Yet, all existing rate adaptation mechanisms fail to track the ideal rate even in a

simple, non-mobile urban scenario. Using a custom cross-layer framework, I implement multiple and previously un-implemented rate adaptation mechanisms to reveal the reasons for the failure and design rate adaptation mechanisms which are able to track urban and downtown vehicular and non-mobile environments. Second, I pose a basic, yet unsolved problem: given a time-varying channel and traffic matrix in the aforementioned complex scenario, predict the throughput of an embedded link and understand the complex interactions of factors that lead to its performance. By performing thousands of measurements of embedded links on an urban mesh network and developing an analytical model, this work is the first to show that even a 1 dB change in channel state can yield a bi-modal shift in throughput that emulates a change in node connectivity. Finally, I apply our model and experimentation to modulation rate selection and the interaction of control and data traffic to show that understanding these complex interdependencies leads to operation in improved performance regimes. My work has implications for this and other urban communities which have unequal access to Internet resources, enabling a high-speed access infrastructure at extremely low cost.

Acknowledgments

First and foremost, I would like to thank my advisor, Ed Knightly, who has been a model example of an excellent researcher and provided seemingly endless resources to allow success in my work. Also, I would like to thank Ashu Sabharwal for challenging me over the years. Additionally, I would like to thank Behnaam Aazhang, Dave Johnson, and Eugene Ng for serving on my 599, M.S., and/or Ph.D. committees, providing novel perspectives to my work.

I am extremely grateful for Technology For All, especially Jim Forrest and Will Reed. They have not only been mentors for me on community engagement but have allowed me to perform research for a real need. Moreover, I thank James Carlson and Nell Warnes for managing the TFA Network, freeing my time for academic work.

I am thankful for the friendships and collaborations within our research group, including Bahar, Aleksandar, Soups, Violeta, Joshua, Jingpu, Misko, Ahmed, Ehsan, Tasos, and Eugenio. For my work on the Wireless Overhead Multiplier, I am thankful for collaboration with Vincenzo Mancuso and Omer Gurewitz [21], and for the the embedded flow modeling, I am thankful for collaboration with with Ehsan Aryafar [27]. Further, I would like to thank the entire WARP Team, especially Chris Hunter and Patrick Murphy for their guidance on the platform. Further, Jaimeet Gulati and Chirag Sequeira tremendously helped me in performing rate adaptation experiments.

My wife, Molly, has been a pillar of strength, counsel, love, and encouragement for me through it all—including a number of paper deadlines that prevented things she would have enjoyed. To my parents who have always believed in me and older sister, Juliana, who has always looked out for me: thank you. I am especially indebted to my Mom for how sacrificial she has been throughout my life. Lastly, I appreciate my father-in-law, Powell, helping me to enjoy research.

I dedicate this thesis to two dear family members who have passed. The first is my late grandfather, Lawrence David (LD) Camp, who, though he did not say much, managed to always tell me how proud he was of me. He was probably the hardest worker I know and I consistently reflect upon how he would approach life. The second, is my late mother-in-law, Martina Robinson. Though I knew Tina a relatively short amount of time, I knew her well and she was one of my biggest fans. She could understand people well before others could and for this reason, I consider it an honor that she thought well of me. I am extremely thankful for the blessing that they both have been to me and the richness that they brought to my life.

Finally, a vital piece of my Ph.D. story began just over eleven years ago in high school with no money for college. Since then, the LORD has generously provided for eleven years of full-time higher education to pursue the giftings and passions that he has placed within me without incurring debt. In the future, I hope to use this overwhelmingly-generous gift to serve others and participate in His provision for them.

Contents

Abstract	ii
Acknowledgments	iv
List of Figures	x
List of Tables	xv
1 Introduction	1
1.1 Summary of Thesis Contributions	4
1.2 Thesis Overview	6
2 Background	8
2.1 Technology For All Network Deployment	8
2.2 Wireless Open-Access Research Platform	10
3 Modulation Rate Adaptation	12
3.1 Introduction	13
3.2 Multirate Protocol Implementation	20
3.2.1 CSMA Protocol Mechanisms	20
3.2.2 Cross-Layer Rate Adaptation Framework	22
3.3 In-Lab Evaluation of Diverse PHY Operating Conditions	25
3.3.1 PHY Operating Conditions	25

3.3.2	Impact of Coherence Time	28
3.3.3	Coherence Time Training	33
3.3.4	Multipath and Interference Effects	38
3.4	In-Lab Evaluation Under Heterogeneous Links	42
3.4.1	Heterogeneous Links	42
3.4.2	Hidden Transmitters with Heterogeneous Forwarding Links . .	45
3.4.3	In-range Transmitters with Heterogeneous Forwarding Links .	49
3.5	Residential Urban and Downtown Scenarios	52
3.5.1	Residential Urban and Downtown Experiment Design	52
3.5.2	Impact of Environment (Static Pair)	54
3.5.3	Impact of Mobility	58
3.5.4	Impact of Heterogeneous Links	61
3.6	Summary	64
4	Embedded Link Performance	65
4.1	Introduction	65
4.2	Link Variation and Capture in a Large-Scale Urban Mesh	68
4.2.1	TFA Network Link Variation	68
4.2.2	Background: Timing Impact on Capture	70
4.2.3	Capture Experiment Set-up of Prism Chipset	71

4.2.4	Capture Threshold as a Function of Modulation Rate and Packet Size	72
4.2.5	Capture Prevalence in TFA	74
4.3	Embedded Link Model	76
4.3.1	Background: Embedded Link Scenario	76
4.3.2	Link Throughput Model	79
4.4	Topological Profile Inversion and the Bi-Modal Effects of 1 dB	95
4.4.1	Experimental Set-up and Measured Model Inputs	95
4.4.2	Baseline Scenario: Symmetric Cross-flow Connectivity	96
4.4.3	Inverted Traffic Profile for Asymmetric Cross-flow Connectivity	99
4.4.4	Asymmetric Profile Inversion: Effects of Forward and Reverse Capture	101
4.4.5	Reverse Capture Shifts the Traffic Profile for Symmetric Cross-flow Connectivity	104
4.4.6	Link Variation of 1 dB Driving Bi-Modal Topological Profile Inversion	106
4.5	Applications of Embedded Link Model and Experimentation	108
4.5.1	Altered Rate Selection Problem Due to Capture	109
4.5.2	Experimentally Discovering the Wireless Overhead Multiplier	111
4.5.3	Isolating Link Effects for the Wireless Overhead Multiplier	118

4.5.4	Predicting the Throughput Reduction Factor for all Traffic Types and Rates	130
4.6	Summary	133
5	Related Work	135
5.1	Modulation Rate Adaptation	135
5.2	Embedded Link Performance	139
6	Conclusion	141
	References	144

List of Figures

1.1	Heterogeneous links inherently exist in mesh networks even in an idealized topology and propagation environment.	3
2.1	Connectivity graph of the TFA backhaul with appropriate scaling for distance between nodes. There are 4,000 residential users (not shown).	9
2.2	WARP FPGA and MIMO-capable radios.	10
3.1	Scenario to measure rate adaptation accuracy with different PHY operating conditions.	27
3.2	Throughput versus coherence time for a high-quality Rayleigh channel.	29
3.3	Per-packet accuracy statistics for experiment depicted in Fig. 3.2 at 100 μ s coherence time.	30
3.4	Total number of undersampled packets per protocol as a function of the coherence time.	31
3.5	Total number of oversampled packets per protocol as a function of the coherence time.	32
3.6	Performance of modulation schemes at 0 kph (left) and at 100 kph (right).	34

3.7	Achievable throughput of SNR-triggered protocols before and after coherence time training.	36
3.8	Throughput with multipath-induced fading for SNR-based protocols.	39
3.9	Undersampled packets by the loss-triggered protocols as interference is injected.	40
3.10	Two scenarios with and without links from node A to node C, varying links 1 and 2.	43
3.11	Sensitivity to small differences in SNR of competing links with collision/fading differentiation.	47
3.12	Throughput difference of in-range flows varying relative SNR between senders (A minus C).	50
3.13	Two-node scenario for mobility in both environments, static sender to mobile receiver.	53
3.14	Heterogeneous links in residential urban environment.	54
3.15	Rate adaptation accuracy with static pair in residential urban (left) and downtown (right).	56
3.16	Normalized throughput (from max value in environment) for each of the multirate protocols within a residential urban setting.	59
3.17	Rate adaptation accuracy with mobile scenario in residential urban setting without (left) and with interference (right).	60

3.18	Throughput of each of the four protocols in the heterogeneous links topology (Fig. 3.14).	62
4.1	Link Variation in the urban mesh network.	69
4.2	Node <i>A</i> leads <i>B</i> by the synchronization bits in the preamble which allows the receiver to lock on to <i>A</i> 's packet.	71
4.3	Capture experiment set-up for Prism 2.5 Chipset.	72
4.4	Capture probability of the TFA hardware in-lab on two controlled channels of a channel emulator.	73
4.5	Histogram of all competing backhaul link pairs within TFA according to relative SNR.	75
4.6	A snapshot of an embedded flow in a complex urban scenarios including various interfering sources.	77
4.7	Discrete time renewal process for the joint channel state evolution. . .	80
4.8	Snapshot of embedded flow coupled with symmetric and asymmetric cross-flow connectivity.	83
4.9	Four cases for symmetric cross-flow connectivity based on timing and capture behaviors.	85
4.10	Four cases for asymmetric cross-flow connectivity based on timing and capture behaviors.	92
4.11	Symmetric cross-flow connectivity (from Table 4.6) with basic access.	97

4.12 Symmetric cross-flow connectivity (from Table 4.6) with the four-way handshake.	98
4.13 Asymmetric cross-flow connectivity (from Table 4.7) with basic access.	100
4.14 Asymmetric cross-flow connectivity (from Table 4.7) with the four-way handshake.	101
4.15 Asymmetric cross-flow connectivity with basic access.	102
4.16 Asymmetric cross-flow connectivity with 2-way handshake (a) and 4-way handshake (b-d) where the disadvantaged flow has increased chance to win the contention due to favorable capture relationships.	103
4.17 Symmetric cross-flow connectivity with (a) 2-way access and (b) 4-way access with both reverse captures of Bb to 1.	105
4.18 Asymmetric topology (from Table 4.7) over a month of measurements and two days' measurements where the maximum and minimum difference between sharing is achieved.	107
4.19 Embedded link throughput when hidden source with uncoupled backoff saturates the channel.	110
4.20 Achieved throughput with and without overhead (isolated) injected from the TFA network.	113
4.21 Node s sends data and overhead (OH) to node r while all other nodes within set O transmit only OH.	116

4.22 Node s sends data and overhead (OH) to node r while node o transmits only OH.	119
4.23 802.11-behavior experiment set-up for the off-the-shelf wireless cards.	120
4.24 802.11 card behavior when noise is injected at the transmitter only.	121
4.25 WOM considering the link class (transmission range or out of range) of the link between the WOM-inducing node o and the data-sender s	123
4.26 WOM (left) and aggregate TCP (right) considering use of RTS mechanism in an out of range scenario.	125
4.27 WOM of out of range links considering relative RSSI at data-receiver r from data-sender s and WOM-inducing node o	127
4.28 Asymmetry of WOM of two nodes with respect to one another.	128
4.29 WOM considering the different effects of physical layer capture effect within the out of range link class.	129
4.30 Range of throughput reduction multipliers based upon the data rate of the embedded link.	131
4.31 Packet size and physical layer rate effect on throughput reduction factor.	133

List of Tables

3.1	Coherence time and SNR necessary for rate increase to ideal modulation rate.	35
3.2	Performance of rate adaptation under heterogeneous links in hidden terminal scenario.	44
3.3	Performance of rate adaptation under heterogeneous links in hidden terminal scenario.	44
4.1	Parameters for model from TFA hardware.	81
4.2	Two-way access with symmetric connectivity.	82
4.3	Four-way access with symmetric connectivity.	84
4.4	Two-way access with asymmetric connectivity.	88
4.5	Four-way access with asymmetric connectivity.	89
4.6	Symmetric cross-flow connectivity sub-topology where a positive value favors Aa and negative, Bb	97
4.7	Asymmetric cross-flow connectivity sub-topology (positive value favors Aa and negative, Bb).	99

Chapter 1

Introduction

Urban areas are densely populated with IEEE 802.11 nodes that span many different network architectures and domains, many of which interact with one another. One such network is the Technology for All Network we have deployed in Houston, TX which provides free Internet to 1000's of users over multiple square kilometers. The TFA Network is a first-of-its-kind, multi-tier wireless mesh network where research can be performed in a densely populated urban environment. Namely, the largest research mesh network prior to our deployment was 40 users [1]. Further, not only are thousands of wireless devices interacting in the network but numerous access points provide localized Internet access to homes, businesses, schools, and libraries which belong to different networks but are on the same frequency. As a result, it is clear why the conventional notion of "scale" has to do with network size and why researchers consider simulation or modeling scenarios which have thousands of wireless devices (since prior testbeds have not existed with thousands of wireless nodes). However, in doing so, there are inevitable assumptions to allow the tools (e.g., models and simulators) to be tractable and thereby, reducing the "scale" of the factors considered. Thus, as a predecessor to understanding the network-wide performance, we must first understand the complex factors affecting link performance in the TFA

Network which produces an entirely new problem space. We term such a link within a topology an *embedded link*.

In particular, time-varying factors affect performance such as topology, channel conditions, modulation rates, packet sizes, and physical layer capture, which we now define. In multi-hop wireless networks, primary consideration is given to the performance of the forwarding links, i.e., links selected by the routing protocol to forward traffic to and from wired gateways. For example, prior work has studied the tradeoff between node spacing and the performance of the resulting links and multi-hop paths [2]. However, since the set of nodes within the main forwarding path use a shared medium, the addition of mesh nodes along the forwarding path also creates a large number of non-forwarding links, or links that are not selected or cannot be selected by the routing protocol to forward data. In the strictest sense, every node forms a link with every other node, even if the resulting link yields near negligible interference. In any case, the resulting connectivity matrix of forwarding and non-forwarding links within a *topology* is vastly heterogeneous in quality due to relative differences in spacing and wireless propagation characteristics among nodes.

The aforementioned heterogeneity is unavoidable, whether these non-forwarding links are foreseen or not during the design process. For example, consider a hexagonal topology in which all neighbors are one-hop and have identical distance to the gateway, as shown in Fig. 1.1. Even if the propagation environment is homogeneous (e.g.,

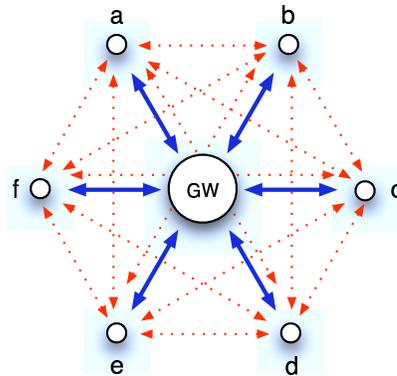


Figure 1.1 Heterogeneous links inherently exist in mesh networks even in an idealized topology and propagation environment.

with a uniform path loss), the links formed within the topology are not identical. Although the forwarding links to the gateway could, in principle, be homogeneous, the non-forwarding links are inherently heterogeneous due to the geometry. Thus, an effective multi-hop topology (non-triangle) necessarily yields heterogeneous links. In a real-world deployment, time-varying *channel conditions* exist with even greater link heterogeneity due to topological irregularities, non-uniform path loss, and the presence of channel fluctuations due to movement of the sender, receiver, or scatterers in the environment.

Since each of these links are of vastly different quality different *modulation rates* are used to allow robustness for a given link. For example, within IEEE 802.11b, there are four modulation rates with the lowest rate (1 Mbps) being least susceptible to losses on a poor quality link and the highest rate (11 Mbps) being most susceptible to

losses [3]. When links compete and their packets overlap in time, differing relative link quality can cause one packet to be successfully received while the other one is dropped, i.e., resulting in *physical layer capture* [4, 5, 6]. Further, there are diverse *packet sizes* within a given network due to different applications having large- or small-sized data packets or even, small-sized control packets for network management.

In this thesis, we perform 100's of thousands of urban measurements and develop an analytical model which is able to predict embedded link throughput in the aforementioned complex scenarios. There are two main areas of focus: *(i)* choosing the modulation rate which maximizes the throughput of embedded links, and *(ii)* predicting and understanding the performance of embedded links with the aforementioned complex factors. We show that the *joint* consideration of these factors yields bi-modal performance with only small changes to channel conditions. Moreover, by understanding embedded link performance, we show that network-wide gains can be achieved.

1.1 Summary of Thesis Contributions

Accurately choosing the modulation rate which maximizes the throughput on embedded links in the network is imperative due to the high usage and shared aspect of the wireless medium. However, we find that all existing modulation rate selection protocols are unable to track a simple urban scenario consisting of a non-mobile sender and non-mobile receiver. Thus, we develop a custom, cross-layer framework

for rate adaptation where multiple and previously un-implemented rate adaptation mechanisms can be directly compared. Unlike prior work evaluating modulation rate selection, we evaluate rate choices on a packet-by-packet basis against the ideal for the channel condition to reveal the reasons behind the rate selection inaccuracies where the ideal rate is found via exhaustive search across modulation rates and channel conditions. The experiments are performed in both in-lab, controlled channels to outdoor, urban channels where links are both in isolation and in competition. To find the ideal rate in urban channels, the channel condition must be measured which is done with 100- μ s granularity. By revealing the reasons for these failures, we then design rate adaptation mechanisms which are able to track the modulation rate across many different urban scenarios including urban and downtown vehicular and non-mobile environments.

Once the modulation rate is accurately selected, there exists a simple, yet unsolved problem: given a time-varying channel and traffic matrix in the aforementioned complex scenario, predict the throughput of an embedded link and understand the complex interactions of factors that lead to its performance. To achieve this, we perform thousands of measurements on embedded links within the TFA Network and develop an embedded link model which incorporates the interdependencies among the complex factors of topology, channel conditions, modulation rates, packet sizes, and physical layer capture. While there has been prediction of CSMA flow throughput

since [7] and even with consideration of capture since [8], no existing prior model or experimentation that can predict the following result: that even 1 dB change in channel state can yield a bi-modal shift in throughput that emulates a change in node connectivity. We term this phenomenon *topological profile inversion* and identify the reasons for these bi-modal shifts to occur via analysis of all possible capture relationships within different sub-topologies. We show that capture of the *reverse* traffic can allow a previously starving flow to compete fairly. Finally, we apply our model and experimentation to two different domains: *(i)* modulation rate selection and *(ii)* interaction of control and data traffic. In each, we show that understanding aforementioned complex interactions allows embedded links to operate in an improved performance region.

1.2 Thesis Overview

The thesis proceeds as follows. In Chapter 2, we present background information on both hardware platforms used for the urban experiments. The first platform is the Technology For All Network which was deployed according to the strategy presented in my prior thesis [9]. The second platform is the Wireless Open-Access Research Platform (WARP), a wireless hardware platform which allows clean-slate design of the physical and medium access layers [10]. In Chapter 3, we develop a custom, cross-layer modulation rate adaptation framework on WARP to perform channel measurements as well as implement, evaluate, and design modulation rate

adaptation protocols. The evaluation is performed a controlled, in-lab environment and in residential and downtown urban vehicular and non-mobile environments. In Chapter 4, we perform embedded link measurements on the TFA topology as well as develop an embedded link model to predict its throughput. On the TFA topology, we experimentally find that the control traffic has a disproportionately large effect on the data traffic which the model is able to reveal the reasons for. In Chapter 5, related work is presented on both topics. Finally, in Chapter 6 we conclude by discussing the implications and future directions that result from the thesis.

Chapter 2

Background

In this chapter, I describe the hardware platforms that are used in both urban experiment data sets: the Technology For All Mesh Network and the Wireless Open-Access Research Platform.

2.1 Technology For All Network Deployment

The TFA network is a multi-tier mesh access network deployed in a densely populated, single-family residential neighborhood. At the time of the study, 17 backhaul nodes are predominantly deployed on single-story residences with the exception of three schools, two businesses, and a public library. The spatial distribution of the backhaul nodes are shown in Fig. 2.1 and are graphically connected if a wireless link can be established between two nodes. All of the wireless links are omni-directional in nature with the exception of three long-haul directional links, pictured darker. The backhaul nodes share Internet bandwidth from a single 100 Mbps fiber and currently serve 4,000 users. The coverage area is 3 km² and has a population density of 4,760 residents per km². For additional details of the deployment and community, refer to [2] and <http://tfa.rice.edu>.

The TFA platform is programmable and observable. Each of the TFA nodes runs an open-source operating system. We perform extensive, non-intrusive, and

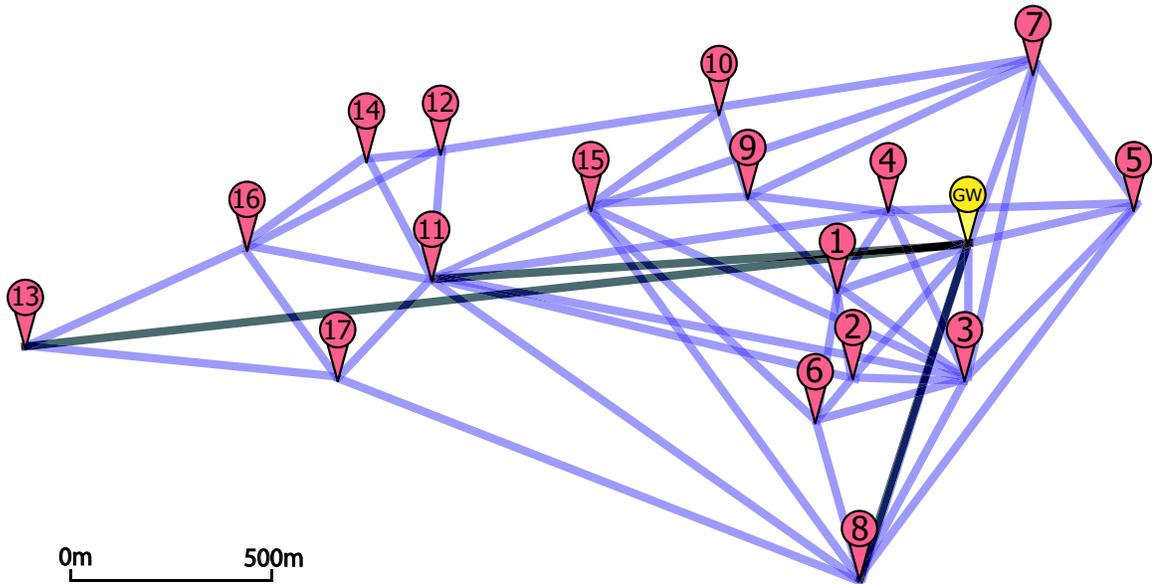


Figure 2.1 Connectivity graph of the TFA backhaul with appropriate scaling for distance between nodes. There are 4,000 residential users (not shown).

privacy-respecting measurements consisting of detailed packet and signal measurements for network operations, modeling, and protocol design. The TFA nodes have much greater processing power (1GHz) and storage (4 GB) than most commercial mesh nodes to handle protocol design and data logging. The Linux operating system is derived originally from the open-source LocustWorld mesh networking software which uses AODV routing and HostAP drivers.

Each mesh node has a single, SMC 2532-B 802.11b wireless adapter with 200 mW transmission power to serve both backhaul and access traffic. The cards connect to a 15 dBi omni-directional antenna with a vertical beamwidth of 8 degrees. The gateway node has multiple radios for added capacity via directional long-haul links on

a different frequency. A TFA backhaul node is a mini-ITX motherboard encased in a waterproof enclosure installed on the outside of building structures of deployment locations. The backhaul antennas are attached to the sides of homes at 10m height, and at slightly greater height (maximum of 20m) at the library, schools, and businesses. The client access node hardware is in many cases unknown to us. Yet, it is clear that a wide variety exists, from PCs employing an external USB WiFi antenna placed near a window to laptops.

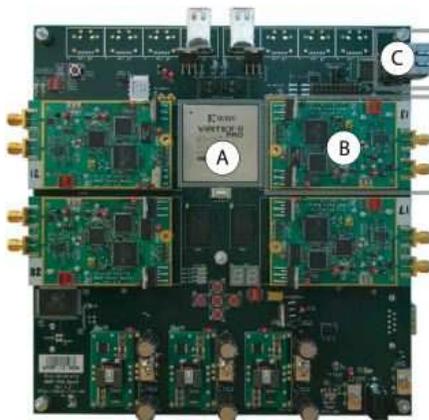


Figure 2.2 WARP FPGA and MIMO-capable radios.

2.2 Wireless Open-Access Research Platform

The WARP platform,* depicted in Fig. 2.2, was designed at Rice University 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: (a) Xilinx Virtex-II Pro FPGA: MAC protocols are written in

*<http://warp.rice.edu>

C and targeted to embedded PowerPC cores whereas PHY protocols are implemented within the FPGA fabric to achieve the required parallelization, (b) MIMO-capable radios: up to four 2.4/5GHz radio boards can be configured and can support wideband applications such as OFDM, and (c) 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 of BPSK, QPSK, and 16QAM, within 10 MHz. For further discussion of the design of WARP, refer to [11].

Chapter 3

Modulation Rate Adaptation

3.1 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 undersampled if the channel state has improved or oversampled 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 hence, 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, see for example, [12, 13, 14, 15, 16, 17]. All current 802.11 rate adaptation protocols are based on the loss-triggered mechanism, and correspondingly, there has been evaluation 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 [18, 19]. These protocols have not been implemented previously

due to the closed and inflexible MAC and PHY of legacy systems.

In this chapter, 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 channel measurements taken on 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 mul-

tirate 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, multipath, 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 opposite 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. Further, we show that the need for such training increases with the presence of multipath, an effect we observe to be strongly present within the downtown scenario but not within the residential urban environment.

Third, with controlled in-lab experiments, we investigate rate adaptation accuracy with heterogeneous links (links of differing average quality), as commonly measured in outdoor environments. We show that a protocol designed to overcome the misinter-

pretation of collision-based losses and fading-based losses with out-of-range senders (the hidden terminal scenario) [17] 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.* 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 and downtown urban 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 network*), multipath (due to

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

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

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 kph for scatterers (passing vehicles), the coherence time is 300 μ s which corresponds to a speed of 250 kph 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., [19]. 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 of the protocol, whereas SNR protocols can track such mobility. Lastly, we find that once the propagation environment of an outdoor setting is characterized in terms of instantaneous and long-term channel conditions (via 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 residential urban and downtown environments which we measure the channel conditions, 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: (i) are better able to track mobility, (ii) have higher robustness to heterogeneous links (including physical layer capture), (iii) have higher accuracy in outdoor environments, especially with the presence of interference-induced losses, and (iv) are able to overcome the overhead penalty of the four-way-handshake using equal air-time assurance. We conclude that when SNR-based protocols consider instantaneous and long-term (coherence time) channel qualities to ensure robustness to varying coherence time, they are a sound alternative to loss-triggered protocols.

The rest of the chapter proceeds as follows. We first describe our custom cross-layer implementation in Section 3.2. We then experimentally evaluate the accuracy of rate adaptation protocols with different channel operating conditions (Section 3.3) and heterogeneous links (Section 3.4). In Section 3.5, we then study the rate adaptation mechanisms in residential urban and downtown scenarios. Lastly, we discuss related work on existing rate adaptation protocols in Section 5.1 and summarize in

Section 3.6.

3.2 Multirate Protocol Implementation

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

3.2.1 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 5 elements, analogous to mechanisms in 802.11: (i) carrier sense, (ii) binary exponential backoff, (iii) network allocation vector, (iv) timeout, and (v) 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) where an idle period corresponds to the amount of time since the received signal strength indicator (RSSI) is below a threshold.

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 to virtual carrier sense 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, the (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, if the ACK is not received before the timeout, 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.

3.2.2 Cross-Layer Rate Adaptation Framework

We implemented five key rate adaptation mechanisms which 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 [14] to implement the consecutive-packet decision mechanism (10 successes, 2 failures). The mechanism uses the two-way handshake (no RTS/CTS exchange) unless otherwise specified.

Historical-Decision Loss-triggered Rate Adaptation

A family of rate adaptation protocols [12, 13, 17] use a window of packets to select the modulation rate as opposed to consecutive successes or failures. Since [17] empirically outperforms [13] (and transitively [12]), we use the specifics described in [17] 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 Mbps). The mechanism uses the two-way handshake unless otherwise specified.

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 [17] and [15], 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 to be due to channel-based loss. While [15] toggles RTS on after a loss for a single packet and then off for the following packet, [17] 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 [17].

SNR-triggered Rate Adaptation

Rate adaptation based upon signal quality requires feedback from the receiver — in our case, we use the CTS packet for that purpose. Use of this mechanism requires a mapping of channel conditions to modulation rates. In [18, 19], 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 3.3 for further discussion of this topic). The four-way handshake is used per-packet within this mechanism.

Equal Air-time Assurance

SNR-triggered rate adaptation with equal air-time 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. To send back-to-back packets, no backoff is performed between the packets so that the transmitter can hold the channel. Also, a queue has to be implemented within the MAC so that bursts of packets can occur. We use the specifics described in [19] to implement equal air-time assurance in which multiple data packets may follow an RTS/CTS exchange.

3.3 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 the rate adaptation accuracy according to these different channel conditions.

3.3.1 PHY Operating Conditions

We identify four channel conditions that have an effect on rate adaptation: (i) coherence time, (ii) delay spread, (iii) interference, and (iv) 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 a sender-receiver pair, thereby inducing a delay between the same symbol from two different paths, called *delay spread*. The 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 simultaneous transmissions from two different transmitters 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 3.4).

Per-Packet Rate Evaluation

As stated in Section 3.4, unlike prior work, the observability between MAC and PHY of the custom cross-layer design allows per-packet evaluation at the receiver 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).^{*} 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 we at least one of the modulation rates 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

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

rate with the ideal rate for the 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: (i) underselected (decoded payload, selected rate less than ideal rate), (ii) accurate (decoded payload, selected rate same or greater than ideal rate), or (iii) overselected (undecodable payload, selected rate greater than ideal rate).

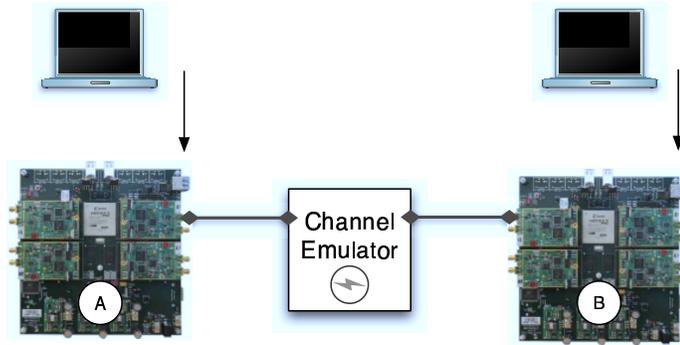


Figure 3.1 Scenario to measure rate adaptation accuracy with different PHY operating conditions.

Scenario

For repeatable channel conditions, we use a channel emulator for the experiments within this section. Fig. 3.1 depicts the experiment set-up where the antenna ports of the two WARP boards are connected via wire to a Spirent Communications 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 Mbps from sender, A , to receiver, B . For the injection of interference, we additionally use an Agilent Signal Generator (ESG-D Series) on the channel from the sender to the receiver.

3.3.2 Impact of Coherence Time Ability to Track Changing Channels

We now evaluate the accuracy of the rate adaptation mechanisms presented in Section 3.2.2 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 μ s to 100 ms. For each coherence time, we measure the accuracy of the four rate adaptation mechanisms triggered by: consecutive-packet decision, historical-decision, SNR, and SNR with equal air-time assurance.

Fig. 3.2 shows the achievable throughput (Mbps) 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,

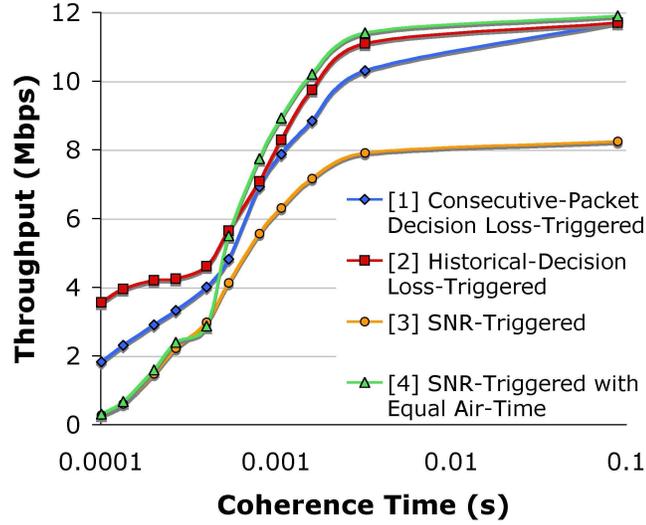


Figure 3.2 Throughput versus coherence time for a high-quality Rayleigh channel.

we show in Section 3.5, that in practical outdoor environments, such scenarios are not encountered. For high coherence time, the SNR-triggered protocol has significantly lower throughput than the other three protocols due to per-data-packet RTS overhead (including the equal air-time assurance mechanism which overcomes this overhead penalty). This result contrasts with simulation-based findings of performance improvements over the consecutive-packet loss-triggered mechanism [18]. For the left portion of the graph (short coherence of the channel), the highest performing protocols at long coherence times (consecutive-packet decision loss-triggered and SNR-triggered with equal air-time assurance) are now the worst performing at short coherence times. The historical-decision loss-triggered protocol becomes the protocol with highest throughput.

Fig. 3.3 further describes the fast-fading case for a coherence time of $100 \mu\text{s}$. Each

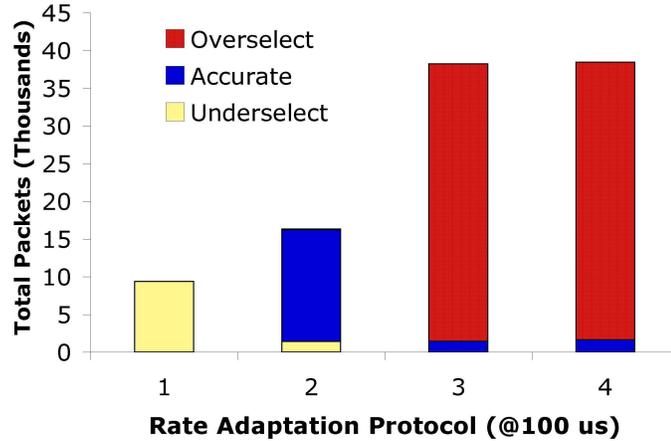


Figure 3.3 Per-packet accuracy statistics for experiment depicted in Fig. 3.2 at $100 \mu\text{s}$ coherence time.

of the four protocols is on the x-axis and total packets are on the y-axis, and are classified according to underselected, accurate, and overselected (for 60 s test). We see that the low throughput in the consecutive-packet decision protocols is due to underselection and in the SNR-triggered protocol is due to overselection. 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.

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 which are received correctly, yet are below the ideal rate of the channel from each of the four mechanisms on a channel with medium average quality (-55 dBm).

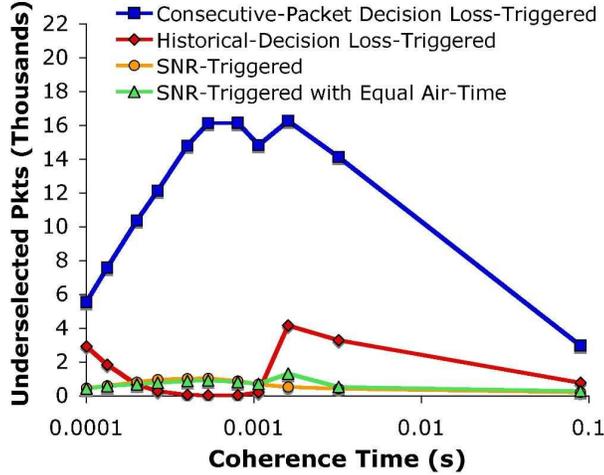


Figure 3.4 Total number of underselected packets per protocol as a function of the coherence time.

Fig. 3.4 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 4k). 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 16k 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 ten 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 undersampled rate now becoming the appropriate choice) and the number of undecodable headers increasing (thereby reducing the amount of decodable headers considered). *Finding:* Loss-triggered protocols undersample from the ideal rate in fast-fading environments due to the consecutive-packet decision mechanism.

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. 3.2) to consider the number of corrupt payloads received which are transmitted with a modulation rate that is above the ideal rate for the channel.

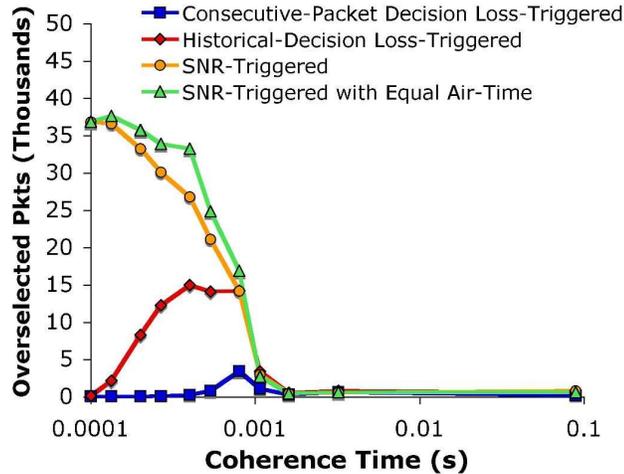


Figure 3.5 Total number of overselected packets per protocol as a function of the coherence time.

Fig. 3.5 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 decreases to 1 ms (approximately twice the packet period interval for a 16QAM packet) 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 36k packets (approaching 100 percent of the packets), 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.

3.3.3 Coherence Time Training

Considering Long-Term Channel Characteristics 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 dBm 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 kph). For each coherence time and SNR, we

measure the achieved throughput.

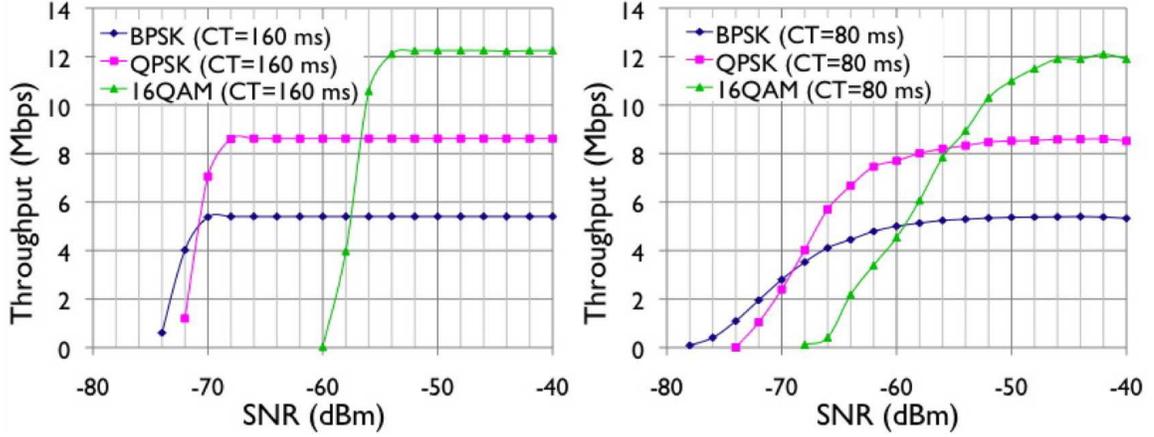


Figure 3.6 Performance of modulation schemes at 0 kph (left) and at 100 kph (right).

Fig. 3.6 shows the achieved UDP throughput (Mbps) as a function of the SNR (dBm) for each of the WARP modulations with a coherence time of 160 ms (left) and 80 ms (right). In Fig. 3.6 (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. 3.6 (right) indicates that the SNR thresholds shift to the right, requiring 2 additional dB for the same modulation rate. Therefore, if the rate adaptation protocol was making a decision based upon a longer coherence time than reality, the protocol would tend to overselect, resulting in loss. Table I shows these thresholds separating the SNR regions for ideal rate on a given coherence time.

Speed	Coh.Time	QPSK→16-QAM	BPSK→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

Table 3.1 Coherence time and SNR necessary for rate increase to ideal modulation rate. Note that 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.

Per our findings in Fig. 3.3 and Fig. 3.5, 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 air-time 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. 3.7 depicts the measured achievable throughput (Mbps) 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 Mbps 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 Mbps and 1 Mbps less than the statically-trained, SNR-triggered protocols with and without equal air-time assurance, respec-

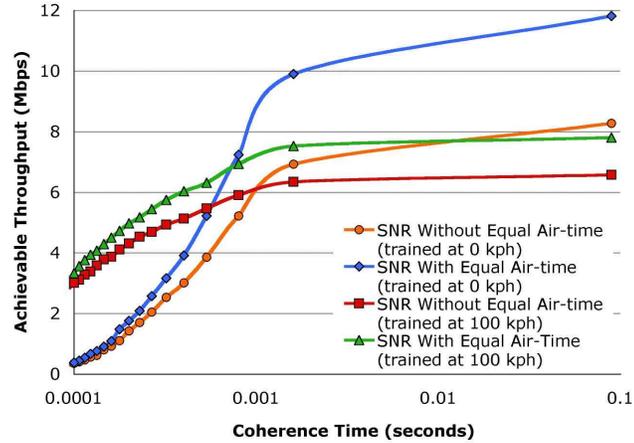


Figure 3.7 Achievable throughput of SNR-triggered protocols before and after coherence time training.

tively. *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. Further, 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 of how many 10 MHz 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.* Further, 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 the constellation to spin, making the offset difficult to track. Thus, there are certain vehicular speeds that the carrier frequency offset

*For 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 percent).

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) pin at the MAC layer (PowerPC) 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 multi-functional 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 which could map to typical coherence time values. Such an approach would require training in many different environments such as indoor, downtown, and residential 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.

3.3.4 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 5.1, prior studies have shown multipath to be a dominant effect in the packet delivery ratio of a particular modulation rate. Thus, we evaluate rate adaptation accuracy within this context. To achieve this, we use the prior experimental set-up (Fig. 3.1) 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 [20] with an RMS delay spread of 55 ns.

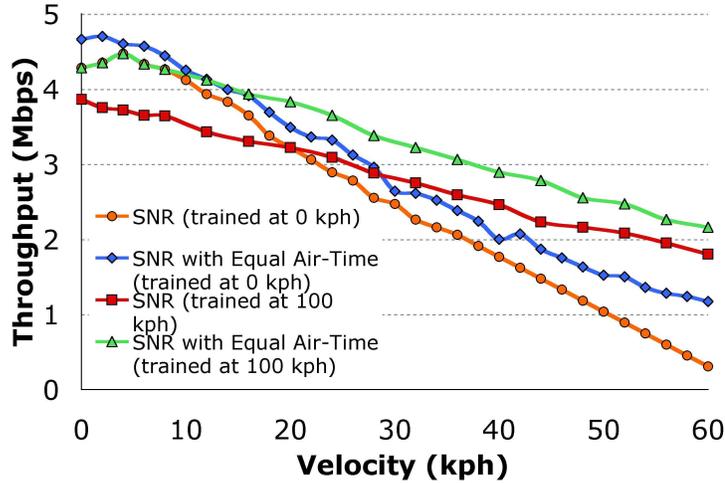


Figure 3.8 Throughput with multipath-induced fading for SNR-based protocols.

Fig. 3.8 shows the achievable throughput as a function of speed for the SNR-triggered protocols with and without training. At speeds of less than 10 kph, the SNR-triggered protocol with equal air-time assurance that was trained at 0 kph has the highest throughput. However, at only 10 kph, the protocol which is trained at 100 kph becomes the highest performing, thereby, showing that multipath has shifted the

speed of the most appropriate coherence time training. *Finding:* When multipath is present in the environment, the sensitivity of SNR-triggered mechanisms to coherence time is increased, and coherence time 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 open spectrum is populated by numerous devices including cordless telephones, microwaves, and other networks. In our experiment, we use a slow-fading channel with packet-sized noise (2 ms) and vary the idle time between noise.

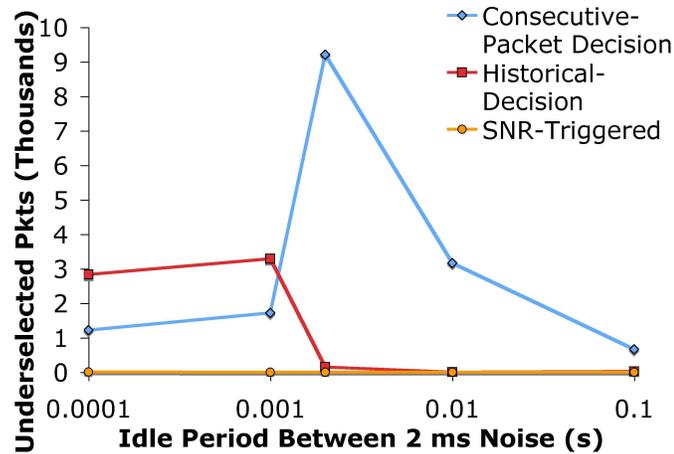


Figure 3.9 Underslected packets by the loss-triggered protocols as interference is injected.

Fig. 3.9 depicts the total number of packets underslected (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 which are successful greatly decreases and the transmitter sends fewer packets due 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 underselect.

3.4 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.

3.4.1 Heterogeneous Links

Differences in link qualities can exist among forwarding links (those selected by the routing algorithm) and non-forwarding links (not selected by the routing algorithm). For heterogeneous *non-forwarding* links, the behavior and coordination of competing transmitters depend on whether neighbors are in-range (as depicted in Fig. 3.10 (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. 3.10 (right).) For heterogeneous *forwarding* links, a wide range of link qualities can exist, e.g., links 1 and 2 in Fig. 3.10 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).

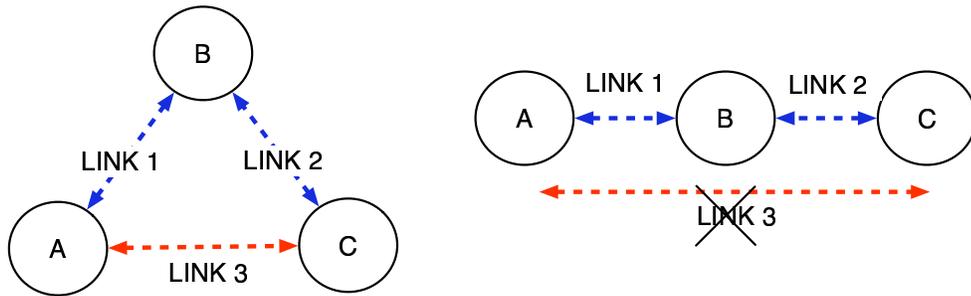


Figure 3.10 Two scenarios with and without links from node A to node C, varying links 1 and 2.

Scenario

As discussed in Section 5.1, although heterogeneous forwarding links with hidden terminals and the physical layer capture effect are common scenarios in practice [21], neither scenario has been explored with respect to rate adaptation. To explore these effects, we use the topologies depicted in Fig. 3.10, 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 or out-of-range (left and right figures, respectively). The experiments are 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 non-forwarding link (3) to be of 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 Mbps offered load.

Relative SNR	Loss-Triggered Rate Adaptation			
	Without Collision Diff.		With Collision Differentiation	
	Throughput (kbps - %)		Throughput (kbps - %)	
	Node A	Node C	Node A	Node C
-20 dB	5464 (99.6)	24 (0.4)	8258 (90.9)	822 (9.1)
-15 dB	1387 (79.1)	367 (20.9)	9618 (98.4)	160 (1.6)
-10 dB	682 (53.7)	589 (46.3)	8308 (84.8)	1490 (15.2)
-5 dB	679 (51.7)	635 (48.3)	6496 (63.7)	3699 (36.3)
0 dB	655 (50.0)	655 (50.0)	5211 (50.3)	5139 (49.7)
5 dB	665 (51.3)	631 (48.7)	3129 (30.7)	7050 (69.3)
10 dB	619 (48.1)	668 (51.9)	1480 (15.7)	7924 (84.3)
15 dB	214 (9.3)	2084 (90.7)	233 (2.6)	8651 (97.4)
20 dB	45 (0.7)	6295 (99.3)	917 (10.2)	8084 (89.8)

Table 3.2 Performance of rate adaptation under heterogeneous links in hidden terminal scenario.

Relative SNR	SNR-Triggered Rate Adaptation			
	Without Equal Air-time		With Equal Air-time	
	Throughput (kbps - %)		Throughput (kbps - %)	
	Node A	Node C	Node A	Node C
-20 dB	2502 (57.9)	1820 (42.1)	3324 (57.1)	2495 (42.9)
-15 dB	2387 (55.9)	1886 (44.1)	3319 (57.1)	2495 (42.9)
-10 dB	3186 (62.7)	1897 (37.3)	4700 (65.2)	2511 (34.8)
-5 dB	2953 (51.1)	2822 (48.9)	3994 (49.9)	4017 (50.1)
0 dB	2888 (49.1)	2998 (50.9)	4039 (50.6)	3945 (49.4)
5 dB	2850 (49.4)	2919 (50.6)	3940 (48.7)	4155 (51.3)
10 dB	1741 (35.0)	3236 (65.0)	2038 (29.8)	4809 (70.2)
15 dB	1854 (43.3)	2429 (56.7)	2300 (40.1)	3440 (59.9)
20 dB	1272 (35.3)	2330 (64.7)	1890 (34.5)	3592 (65.5)

Table 3.3 Performance of rate adaptation under heterogeneous links in hidden terminal scenario.

3.4.2 Hidden Transmitters with Heterogeneous Forwarding Links Collision- and Fading-Based Loss with Heterogeneous Competing Links

Recall that reference [17] considered hidden terminals with *homogeneous* links, whereas we now explore the joint effect of hidden terminals (Fig. 3.10 (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. 3.10 (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 one transmitter constant (-45 dBm) and lower the link quality of the other transmitter in steps of 5 dB up to 20 dB. Note from Fig. 3.6 (left) that in the region of -45 to -55 dBm, the highest modulation rate (16QAM) should be chosen as nearly all packets are still able to be received correctly at this rate.

Table 3.2 and 3.3 contain the per node throughput (kbps) 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 air-time assurance (see Section 3.2.2). Observe that in the

middle of the table where equivalent links exist (in bold), 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 high 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 there is high aggregate throughput 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. *Finding:* Slight differences in link quality, even within 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.

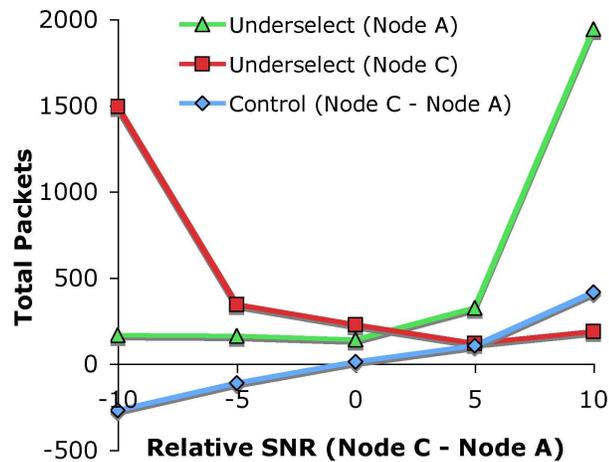


Figure 3.11 Sensitivity to small differences in SNR of competing links with collision/fading differentiation.

In Fig. 3.11, 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. Conversely, observe that when the difference is -5 dB (A is the stronger transmitter), A has greater protection from the 112 additional control packets compared to C. Therefore, the weaker transmitter using collision/fading differentiation, like the behavior of the protocol without the mechanism, has increased losses due to the lack of RTS protection and begins to lower the rate yielding increased underselection. The reason for this RTS usage difference is due to the dynamic window that is used for the number of packets using RTS

[17]. The window is halved either when RTS is lost or RTS, data, and ACK have all succeeded. Thus, since the weaker link has a slightly less RTS success rate and the event of losing RTS is much shorter than a complete successful exchange, the weaker transmitter halves its window much more often. This leads to many more data packets that are not preceded by RTS. *Finding:* The exaggeration of slight differences in link quality of the collision/fading differentiation mechanism is due to unequal use 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 and drives rate adaptation to previously unexplored behavior. In Table 3.2, 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 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). 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 Mbps (86%) and 945 kbps (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.

3.4.3 In-range Transmitters with Heterogeneous Forwarding Links Competing Multirate Links with Ability to Carrier Sense

As reference [22] showed in the case of in-range heterogeneous forwarding links, low-quality links can cause even high-quality 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.

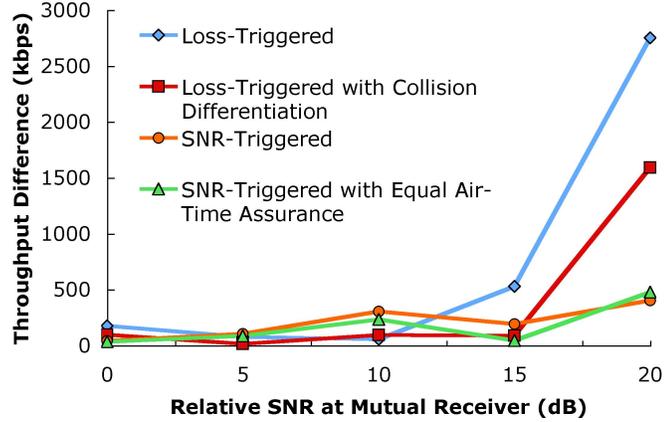


Figure 3.12 Throughput difference of in-range flows varying relative SNR between senders (A minus C).

Fig. 3.12 shows the difference in achieved throughput (kbps) 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 Mbps 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 kbps 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 (i) does not suffer from the sensitivity of the collision/fading differentiation under heterogeneous forwarding links;

and (ii) 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.

3.5 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 non-mobile sender-receiver pairs.

3.5.1 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 3.3 and 3.4. Thus, we first characterize these environments for perspective on the experiments in prior sections, and then evaluate the rate adaptation accuracy within these scenarios.

Scenario

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

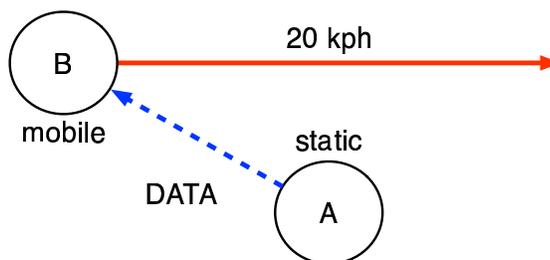


Figure 3.13 Two-node scenario for mobility in both environments, static sender to mobile receiver.

Fig. 3.13 is used where a static node (A) sends data to a mobile node (B). The mobile node starts moving with a nominal link, 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. 3.14. 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 [2]. The downtown measurements are performed on streets in Houston, Texas 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. 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.

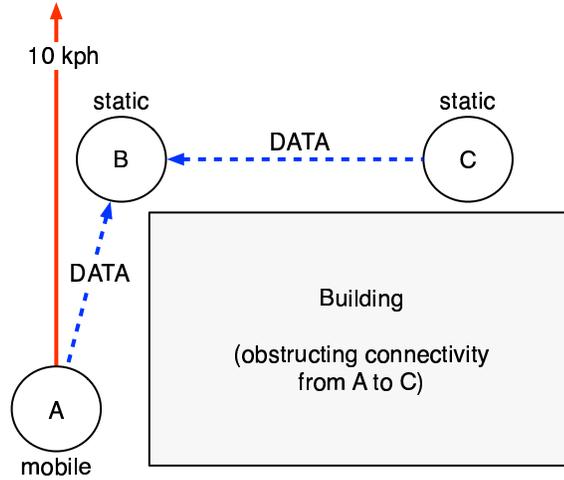


Figure 3.14 Heterogeneous links in residential urban environment.

3.5.2 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 reference [23], 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 characterize the environments with a pair of static nodes. 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. In both environments, vehicles pass at approximately 30 kph within 5 m from the location of

one of the nodes in the experiment. On average, we find the coherence time to be 100 ms and 80 ms within the residential urban and downtown scenarios, respectively. 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.

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 kph.

Rate Adaptation Accuracy in Outdoor Environments

Next, we measure the performance of the multirate mechanism within both the residential urban and downtown scenarios. In our experiment, we generate UDP traffic from the 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 transmitting with 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. 3.15 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 follow-

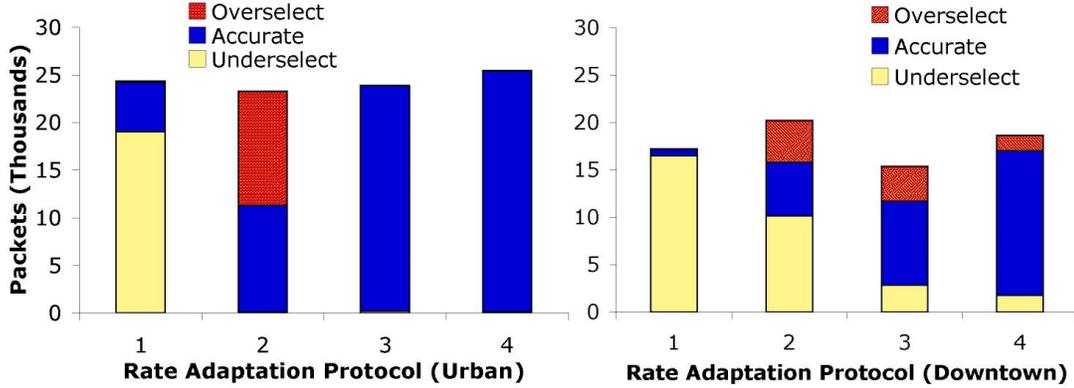


Figure 3.15 Rate adaptation accuracy with static pair in residential urban (left) and downtown (right).

ing order, triggered by: (1) consecutive-packet decision, (2) historical-decision, (3) SNR, and (4) SNR with equal air-time assurance. In the residential urban scenario (Fig. 3.15 (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 ten 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 solely for channels more representative of indoor settings. *Finding:* Consecutive-packet decision and historical-decision loss-based protocols are largely inaccurate at adapting the rate within a practical outdoor setting.

In the downtown scenario (Fig. 3.15 (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 3.3 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 turn-around 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 over-select and forces loss-triggered protocols to underselect due to effects analyzed in Section 3.3.

We also observe that in the downtown scenario of Fig. 3.15 (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.

3.5.3 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 kph, as depicted in Fig. 3.13. We track the per-packet variance of the SNR to measure the channel fading during the experiment.

For the residential urban scenario, Fig. 3.16 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 seconds) 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 protocols are not able to track the mobile client, even at relatively low

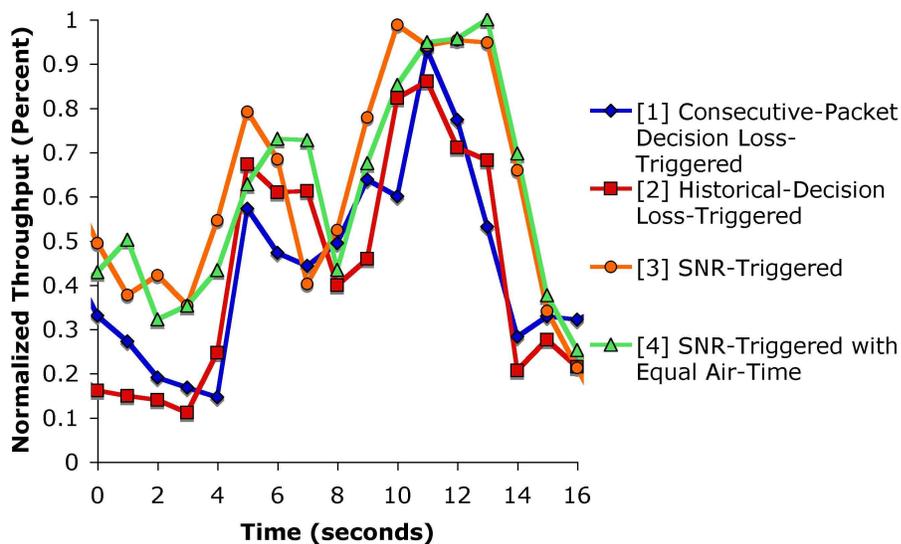


Figure 3.16 Normalized throughput (from max value in environment) for each of the multirate protocols within a residential urban setting.

speed. (Similar results hold for downtown but are not shown here.) *Finding:* Sequential rate stepping 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. 3.16, and then when the two nodes are on the same channel as the TFA network containing 17 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.

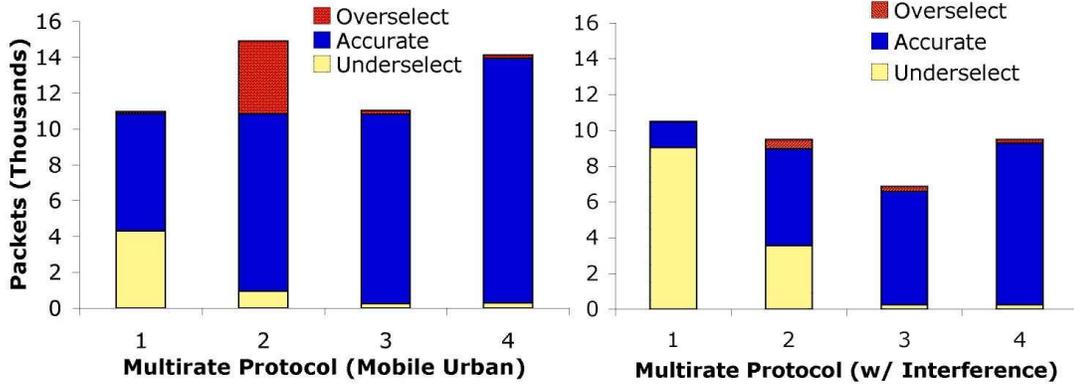


Figure 3.17 Rate adaptation accuracy with mobile scenario in residential urban setting without (left) and with interference (right).

Fig. 3.17 contains the underslected, accurate and overselected packets (y-axis) for each of the four mechanisms first without (left figure) and then with interference (right figure). We observe that some rate decisions are effected by the interference within the loss-triggered protocols (1 and 2, in the figures) as each of the first two protocols increase the number of underslected packets by at least 3k 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 accurate packets) by the presence of interference since the available idle time is reduced, causing a smaller percentage of successful four-way handshakes than if the

two-way handshake were used. *Finding:* Interference from external devices shortens the interval of available air-time, causing an increased number of losses, triggering both types of loss-based protocols to underselect. Further, 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).

3.5.4 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. 3.14, which is a dynamic version of each of the scenarios in the previous section: 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 channel conditions of such a link 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. 3.18 depicts throughput for both of the contending transmitters (A and C) as a function of time. In the left half of each of the graphs, the two transmitters are hidden from one another and approximately half-way through the experiment, they become in-range. For the consecutive-packet decision loss-triggered

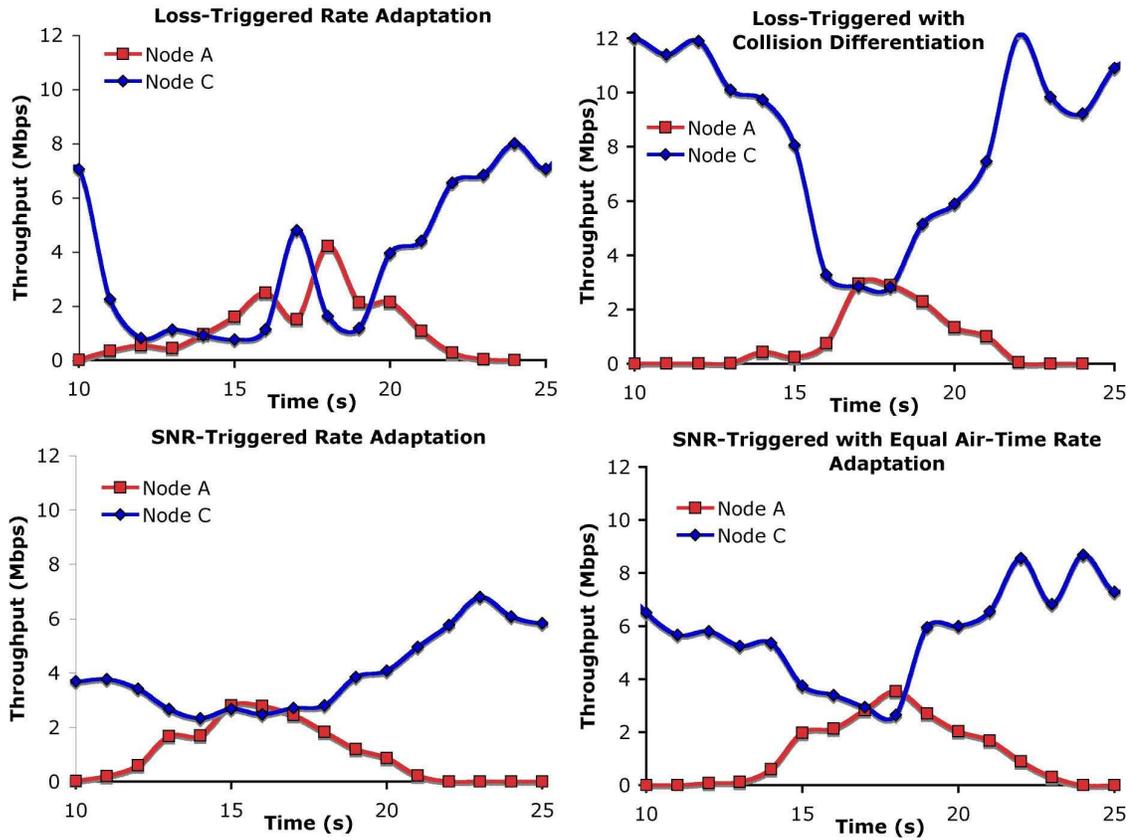


Figure 3.18 Throughput of each of the four protocols in the heterogeneous links topology (Fig. 3.14).

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 3.4 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 air-time 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.

3.6 Summary

In summary, we developed a custom cross-layer rate adaptation framework which 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 ideal 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. Further, we found that coherence time 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.

Chapter 4

Embedded Link Performance

4.1 Introduction

In urban environments, IEEE 802.11 nodes interact in many ways, e.g., within and among paths in a multi-hop network and among deployments from different domains. Moreover, node interactions are affected by a vast array of factors including topology, modulation rate, packet size, channel conditions, and physical layer capture. In this chapter, we pose the following unsolved problem: given a time-varying channel and traffic matrix in the aforementioned scenario, predict the throughput of an embedded link and understand the complex interactions of factors that lead to its performance. To do so, we perform several thousand measurements in a dense urban mesh deployment and experimentally show that a small-scale change in channel conditions can mimic a change in physical topology. We term this *topological profile inversion*, as despite physical connectivity remaining unchanged, channel dynamics of even 1 dB can yield a bi-modal performance shift typically associated with a different topology. We devise a simple model that can both predict an embedded link's throughput in this environment and provide a fundamental understanding of the origins of this behavior.

In particular, we present the following three contributions. First, we develop an analytical model that can jointly incorporate topology, modulation rate, packet size, channel conditions, and physical layer capture. The model predicts throughput

of embedded links by incorporating these system factors via use of a broad set of link interaction states in an embedded Markov chain. This model is the first to characterize our experimental finding that even a 1 dB change in channel state can yield a bi-modal shift in throughput that emulates a change in node connectivity.*

Second, we design a set of urban experiments consisting of 1000's of measurements. We first validate the analytical model and show that it is accurate in predicting embedded link throughput for diverse channel conditions and topologies. Next, we study embedded link sub-topologies that exhibit bi-modal behavior and experimentally identify the conditions for switching modes. One such sub-topology occurs when embedded links compete asymmetrically due to topological connectivity factors. Namely, a “disadvantaged flow” can starve due to lack of knowledge at the sender about when to begin contention whereas the source of the “advantaged flow” has full information about when to begin contention. Consequently, the advantaged flow “wins” the contention nearly all the time [24]. Throughout, we show that *reverse traffic* (acknowledgment and clear-to-send packets traveling in the reverse direction of data) has a critical impact. In contrast to data, this reverse channel is *not* carrier sensed before transmitting. When coupled with capture relationships, this yields new link interdependencies, interactions with forward traffic, and vulnerable sub-topologies, all characterized by the model. We show that in these sub-topologies,

*For a thorough discussion of related work refer to Section 5.2.

even a small-scale channel fluctuation can yield a switch from a starvation mode to a fair-contention mode, or vice-versa, as if the connectivity among the contending flows was changed.

Finally, we apply this understanding to two domains. *(i) Modulation rate selection.* The conventional wisdom is that the modulation rate should be as high as possible for the sender-receiver channel conditions. However, we show that that the joint effects of topology, packet size, channel conditions, and physical layer capture must be considered. *(ii) Interaction of Control and Data Traffic.* Control traffic such as routing announcements with low average rate (10's of kbps) has a disproportionate impact on the throughput of data traffic, potentially reducing data throughput by 100's of kbps [21]. We show how the coupling of small-sized low-modulation rate control packets with large-sized high-modulation rate data packets, together with topology and capture relationships yields this behavior.

The rest of the chapter is organized as follows. First, we introduce our large-scale urban mesh network and measure its link variation and capture behavior in Section 4.2. We then present our embedded-link model in Section 4.3. In Section 4.4, we study topological profile inversion and the effect of small-scale channel dynamics. We apply our experimental analysis and model in Section 4.5. Lastly, we compare to related work in Section 5.2 and conclude in Section 4.6.

4.2 Link Variation and Capture in a Large-Scale Urban Mesh

In this section, we explore link variations and physical layer capture in an urban mesh deployment. Moreover, in a controlled, in-lab environment, we perform experiments with the same wireless card used in the deployment to understand relative link quality differences that lead to physical layer capture as a function of modulation rate and packet size. We create an SNR matrix of all links within the topology and show how common the competing link pair’s relative quality exceeds these capture thresholds. We use these measurements in the following sections to seed our model with accurate physical layer behavior over time and to understand the poor performance of high and low rates interacting in the application section. Further, the prevalence of physical layer capture within the urban mesh deployment underscores the need for understanding the complex interaction of capture, small-scale channel dynamics, and MAC behavior.

4.2.1 TFA Network Link Variation

Small-scale channel fluctuations play a key role in our experiments. Thus, we first measure the variations on all the TFA links to better understand the behavior of embedded links. To quantify the link variations across the network, we measure the per-second RSSI at each mesh node in a synchronized way for ten minutes. The measurements are taken in the winter with minimum foliage and therefore these link variations are expected to be lower than if they were taken in the summer months

(e.g., due to wind moving the trees).

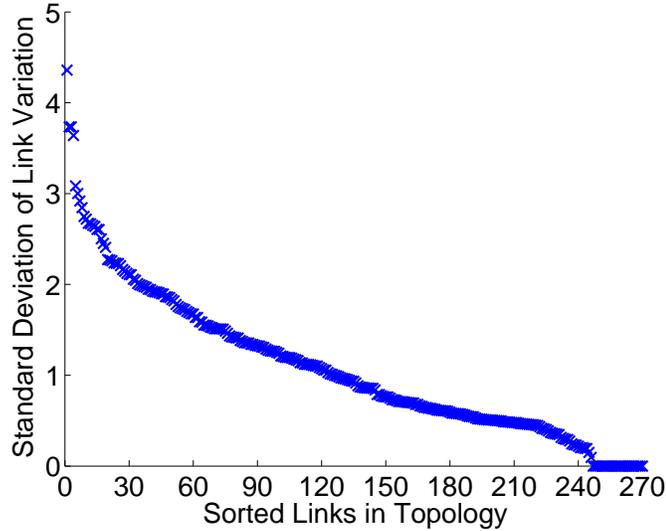


Figure 4.1 Link Variation in the urban mesh network.

Fig. 4.1 depicts the standard deviation of link variations for all 270 links within the TFA Network sorted from highest to lowest variation. Here, a link is defined as any two mesh nodes in which the receiver can hear packets from a sender and record an RSSI value. Thus, a sender-receiver pair could be accounted for twice and symmetry between directions of the pair is not assumed. We observe that 125 links have a standard deviation of above 1 dB and 145 links below 1 dB. However, only 32 low-variation links have high enough link quality to be selected by the routing protocol (greater than -80 dBm), meaning that the remainder of the low-variation links are barely able to receive packets and have RSSI at the lowest possible value for reception. Thus, even in a static urban topology with minimal (seasonal) foliage,

there are link variations of 1 dB in over 80 percent of the usable links.

4.2.2 Background: Timing Impact on Capture

Prior work has shown that the timing of the competing packets plays a critical role in physical layer capture due to the Message-in-Message (MIM) function required by the 802.11a standard and implemented in the Atheros chipset [5]. Namely, if a packet's preamble is received correctly (or, more precisely, enough synchronization bits within the preamble are received correctly), the receiver 'locks on' to that packet (as depicted in Fig. 4.2) and only switches to attempt to decode a later overlapping packet if that packet is greater than 10 dB stronger than the first transmission. As a result, there are two capture thresholds based upon timing. However, in our experiments, we use the Prism 2.5 chipset which has been shown to not implement MIM and forces overlapping transmissions to result in loss if the receiver is locked on to the weaker packet [4]. The ability of the sender to capture depends completely upon the correct Frame Check Sequence versus an interfering source. Thus, in the case where the capture occurs for the Prism chipset (i.e., the stronger packet is first or the stronger packet trails the weaker packet by less than the synchronization bits of the preamble, lasting 6 slots in 802.11b), there is a single capture threshold for the physical layer rate and packet size combination.

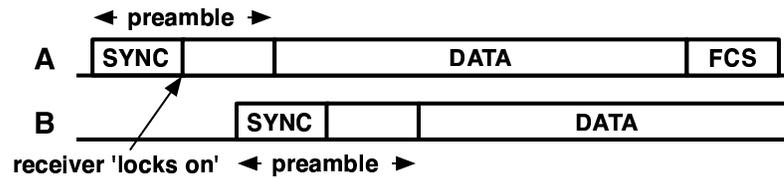


Figure 4.2 Node *A* leads *B* by the synchronization bits in the preamble which allows the receiver to lock on to *A*'s packet.

4.2.3 Capture Experiment Set-up of Prism Chipset

To inform our model about the physical layer capture behavior in TFA, we measure the the delivery ratio of the wireless card when competing against another transmitter at certain relative SNR values. The wireless card is the SMC EliteConnect 2532-B which is an 802.11b card using the Prism 2.5 chipset. Fig. 4.3 depicts our use of both channels of a channel emulator for the transmissions of the node which is trying to capture (sender on channel 1) and the competing transmitter (interferer on channel 2). The sender and interferer are unable to carrier sense one another. For the interferer, we eliminate the effects of binary exponential backoff by sending infinitely long packets. For the sender, we use broadcast traffic to not need ACKs (since we are unable to create a third channel in the reverse direction). We combine the two outputs of the two channels and connect it to the input of the mutual receiver.

For each relative value of SNR, we hold the power constant for channel 1 (-72 dBm) and vary the power of channel 2 (-70 to -84 dBm), testing a range from -2 to 12 dB. During the measurement, the relative power levels are held constant (i.e., no fading is experienced on the channels). We send equally spaced broadcast packets

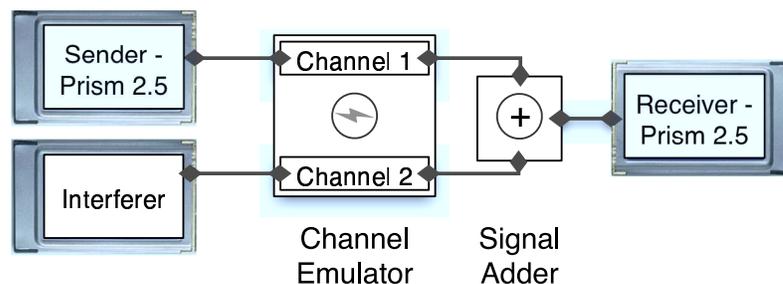


Figure 4.3 Capture experiment set-up for Prism 2.5 Chipset.

(nearly flooding the channel) over a 60-second duration, running synchronized scripts which record the card statistics for both sender and receiver before and after the experiment. Thus, we are able to calculate the packet delivery ratio as the total amount of received packets over the total transmitted packets.

4.2.4 Capture Threshold as a Function of Modulation Rate and Packet Size

Prior measurement studies on physical layer capture have shown that higher modulation rates require greater relative SNR to achieve capture [5, 6]. Since TFA has a diverse traffic profile with large-size, high-rate packets interacting with small-size, low-rate packets it is important to experimentally understand the role of both modulation rate *and* packet size, which has not been fully explored. We expect that larger packets require greater relative SNR to capture since the bit error rate is related to the channel condition, and increasing the number of bits increases the probability of error. However, it is not clear the degree to which packet size will effect these thresholds. To chose relevant packet sizes for our measurements, we refer to recent

studies on the Internet which have classified the traffic according to three different groups of approximately 100, 500, and 1500 bytes [25]. We additionally consider the case of 1000 bytes. We measure the capture thresholds for all four physical layer rate and packet size combinations (16 total configurations). These measurements inform our model about the physical layer behavior based upon both factors.

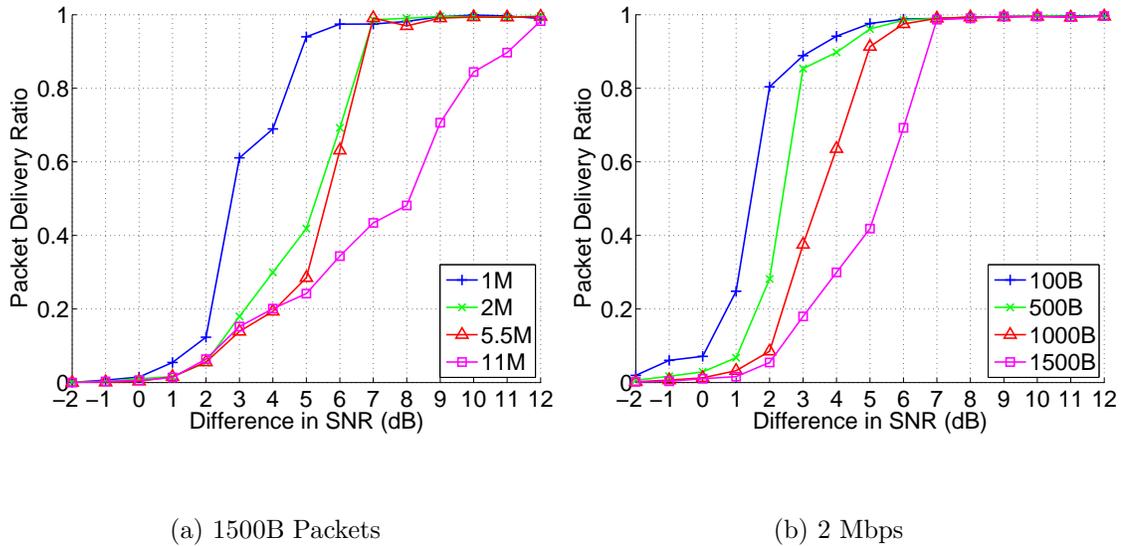


Figure 4.4 Capture probability of the TFA hardware in-lab on two controlled channels of a channel emulator.

Fig. 4.4 depicts the delivery ratio for the sender at each modulation rate for 1500 byte packets (left) and each packet size for the control rate (right) according to each relative SNR value from -2 to +12. For 1500 bytes (a common data packet size within the network), observe that nearly zero packets are delivered for 0 dB and below (when the interferer is as loud or louder than the sender). As, the difference in SNR increases, the lowest modulation rate quickly converges to nearly 1 for 5 dB yet the

highest modulation rate is only able to obtain close to 1 for a relative SNR of 12 dB. In contrast, for the control rate in the right figure (2 Mbps), 2 dB is enough to capture 80 percent of the packets for the smallest size (100 bytes). A difference of over 6 dB is required for the same performance of the largest-sized packets (1500 bytes). Referring back to the 1500 byte result, the relative SNR required to achieve the same packet delivery ratio between a small-sized, low-rate frame and a large-sized, high-rate frame is up to 8 dB different. Thus, the ability to capture is highly dependent upon both packet size and modulation rate rate of the strongest overlapping packet.

4.2.5 Capture Prevalence in TFA

Based on the in-lab measurements for when capture occurs, we now consider capture relationships across the network based on a distribution of relative signal quality of competing link pairs along the TFA backhaul tier. We consider a single point in time, though we have verified that similar distributions exist for all of our measurements which span a week's time with per-second measurements over 10 minute durations. At the beginning of each test interval, the mesh nodes are synchronized according to a global clock using *ntpdate*. From the signal measurements, we sequentially search each mesh node for any two possible links that would compete at a mutual receiver. Since reverse traffic does not carrier sense and since these cards have been shown in [21] to lack the function of physical carrier sensing, we do not exclude competing in-range link pairs from consideration.

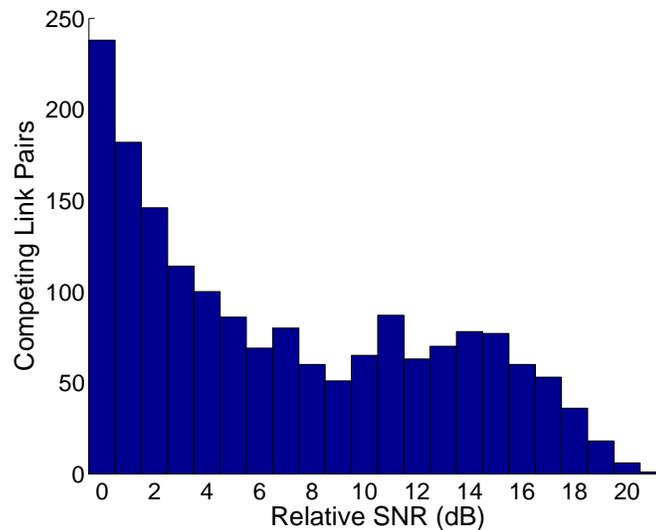


Figure 4.5 Histogram of all competing backhaul link pairs within TFA according to relative SNR.

Fig. 4.5 is a histogram of the number of competing link pairs in the TFA backhaul tier according to the relative SNR between the links. At the time of the measurement, there is a total of 1621 link pairs that compete with one another from the perspective of the 17 omni-directional access points (where each serves as a receiver for competing backhaul link pairs) within the network. Observe that there appear to be two groupings, 0 and approximately 14 dB. The latter grouping corresponds to a number of links which are barely able to receive packets and have the lowest value of RSSI in which the card can receive packets (-85 dBm to -88 dBm) versus well-engineered links (-68 to -75 dBm), creating the range of 20 dB. Out of these link pairs, over 72 percent of the link pairs have relative SNR differences of 2 dB or higher which corresponds to an 97 percent capture ratio for control traffic. Thus, there is a high degree of physical

layer capture that occurs in the TFA topology.

4.3 Embedded Link Model

In this section, we develop an analytical model that predicts the throughput of embedded links in complex scenarios which include diverse topologies, modulation rates, packet sizes, channel conditions, and physical layer capture relationships. These complex system factors are incorporated using a broad set of link interaction states embedded in a bi-dimensional Markov chain.

4.3.1 Background: Embedded Link Scenario

We study the performance of a link embedded in a static multi-hop wireless network. Fig. 4.6 depicts a snapshot in time of one such embedded link in a complex topology with sources interfering with the embedded flow from A to a .

Coupled and Uncoupled Flows and Hidden Terminals

In most cases, a flow such as Bb that is interfering with flow Aa has backoff behavior that is *coupled* to that of flow Aa . In other cases, such as with broadcast traffic, an interfering flow's backoff process is *uncoupled* with the embedded flow. When two transmitters such as A and B use 802.11, if *inter-sender* interference exists in which packets can be sensed or decoded between transmitters, one transmitter defers while the other transmits. The resulting behavior can be predicted using existing models such as [26] and extensions. Thus, we consider the case where the two

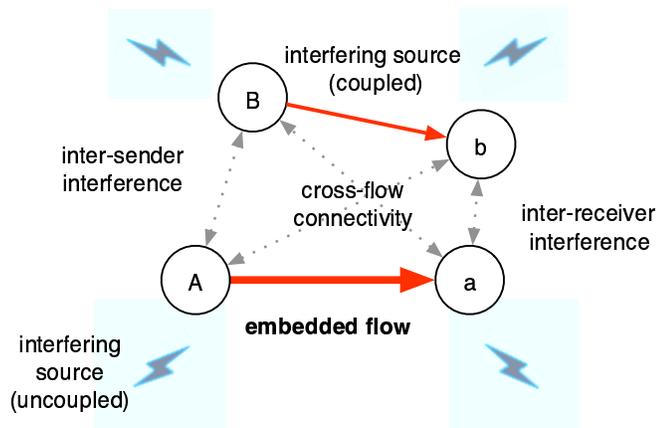


Figure 4.6 A snapshot of an embedded flow in a complex urban scenarios including various interfering sources.

transmitters have no inter-sender interference (hidden terminals). Similarly, *inter-receiver* interference exists where receivers can decode packets from one another based on the delivery ratio per modulation rate and packet size.

Cross-Flow Connectivity

Flows with coupled backoff behavior can have cross-flow connectivity where the sender of one flow is able to receive packets from the receiver of the competing flow. There can be *symmetric or asymmetric cross-flow connectivity* if the senders of both flows are able to decode packets from the receiver of the competing flow or if only one is able to do so, respectively. The symmetry or asymmetry of this relationship has been shown to cause balanced or imbalanced throughput sharing due to the MAC layer behavior [24, 27]. The cross-flow connectivity is probabilistically defined based on the delivery ratio of the links.

Though the infrastructure is fixed, the mobility within the environment can cause fluctuations in channel quality. Further, there are spatial differences between the sender and receiver of a particular data flow. Therefore, there are link quality differences between any two sender-receiver pairs which are addressed by the use of different modulation rates [28]. These differences in link qualities can also cause competing links which have packets simultaneously received at a given receiver to experience physical layer capture [5, 4].

Complexity of Capture Relationships

Physical layer capture can occur for traffic in the forward direction (e.g., data or RTS packets from B and A overlapping at a in Fig.4.6) or for traffic in the reverse direction (e.g., CTS or ACK packets from a and b overlapping at A in Fig. 4.6). For a given flow, there can be *forward traffic capture* over the forward or reverse competing traffic and *reverse traffic capture* over the forward or reverse competing traffic. There are a total of four possible capture scenarios for a given flow with respect to a competing flow and three possible results: winning capture, losing capture, or collision with loss.* Thus, when considering all possible capture states for two hidden flows, there are 729 different scenarios (3^6). Furthermore, additional complexity exists considering that the capture result is probabilistic and depends on timing (see Section 4.2). We next develop an analytical model that takes into account these complexities.

*We consider 3 as opposed to 4 capture relationships per flow since the fourth relationship does exist without inter-sender interference.

4.3.2 Link Throughput Model

Modeling Coupled Sources

We develop a bi-dimensional discrete time Markov chain embedded over continuous time to study the throughput sharing behavior of two coupled sources. In our model, we explicitly account for different capture relationships that exist among different competing nodes at a receiver, where the system state is the joint backoff stage of the coupled sources. The transition probability is determined by capture relationships and other system parameters. This allows the relation of different capture relationships to the steady state distribution of the system state. Using our analytical model, we are able to accurately predict the throughput as well as investigate the impact of capture relationships and other parameters on the system performance.

Joint Channel State Evolution

In order to correctly analyze the behavior of coupled sources, we consider the joint backoff evolution of the two flows. Fig. 4.7 shows an abstract representation of the joint channel state evolution where the arrows correspond to time instants in which both senders can potentially start transmitting the first packet of a new data exchange. We identify three main states: (i) idle state, (ii) single access state where either one flow transmits or both flows transmit but the first packet of the earlier flow does not overlap with the later flow (e.g., the first flow's RTS finishes and is now receiving a CTS while the other flow transmits an RTS), and (iii) overlapping state

where both flows transmit and their first packets overlap.

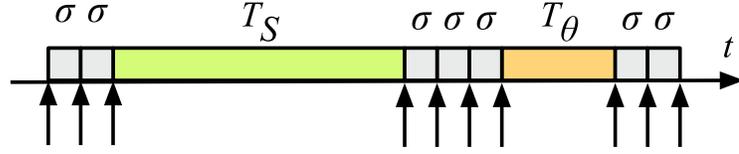


Figure 4.7 Discrete time renewal process for the joint channel state evolution.

The time intervals in which the channel remains in the above states are denoted by σ , T_S , and T_θ , respectively. While σ is a constant equal to one mini-slot duration in 802.11, the duration of the other intervals (T_S and T_θ) depend on the modulation rates of transmitting nodes, the access mechanism, and the overlapping duration when it happens. We denote the packet size and modulation rate of node n as R_n and l_n , respectively. With RTS/CTS, the duration of a successful single access state with no overlapping is equal to:

$$T_{S_n} = \frac{\text{RTS} + \text{CTS} + \text{ACK}}{R_{\text{basic}}} + 3 \cdot \text{SIFS} + \text{DIFS} + \frac{l_n}{R_n} \quad (4.1)$$

In the case of overlapping transmissions, either one or both packets are captured or both packets are dropped. Thus, the duration is variable for each access mechanism, modulation rate, and flow winning capture. These values are calculated for each case once their corresponding probabilities are calculated.

Basic Rate	2 Mbps
Data Rates	1, 2, 5.5, and 11 Mbps
SIFS/DIFS/EIFS	10/50/364 μ s
Mini-slot (σ)	20 μ s
Maximum Retry Limit	4
(CW_{\min}, CW_{\max})	(32, 512)
PHY Preamble	384 bits @2 Mbps (192 μ s)
MAC Header	30 bytes @2 Mbps
RTS/CTS/ACK Size	30/30/14 bytes @2 Mbps
Data Payload	1500 bytes

Table 4.1 Parameters for model from TFA hardware.

System State

We represent the system state as the pair (i, j) , where i and j represent the current backoff stages of transmitters A and B . Note that $0 \leq i, j \leq m$, where $m + 1$ is equal to the maximum retransmission limit. The key approximation in our model is that at each switching time the next state does not depend on the current state. This allows us to model the evolution of our bi-dimensional state process with a discrete time Markov chain embedded over continuous time at the time instants in which *both senders* can potentially start transmitting the first packet of a new data exchange (either RTS or data) if either backoff counter is zero.

We further assume that a station's backoff counter is geometrically distributed over the contention window. This allows us to exploit the memoryless property of the geometric distribution without accounting for the remaining number of backoff slots. The parameter, γ_i , of the geometric distribution that characterizes the backoff

To State	Transition Probability
i, j	$(1 - \gamma_i)(1 - \gamma_j)$
$0, j$	$\gamma_i(1 - \gamma_j)^{J_A}(1 - l_A)$
$0, j + 1$	$P_\theta S_{AB} C_{ABa,D}(1 - l_A)$
$i + 1, j + 1$	$P_\theta S_{AB}(1 - C_{ABa,D}(1 - l_A))$

Table 4.2 Two-way access with symmetric connectivity.

counter at stage i is given by $\gamma_i = \frac{2}{W_i - 1}$ where W_i is the window size of backoff stage i . Consequently, a station in stage i attempts a new transmission with probability γ_i .

Transition Probability Calculation

Nodes A and B have transmission probabilities of γ_i and γ_j , corresponding to backoff stages i and j , respectively. The transition probabilities stem from the generic state (i, j) and are summarized in Tables 4.2 to 4.5. We next calculate these probabilities with symmetric and asymmetric cross flow connectivity (Fig. 4.8) for both access mechanisms.*

Two-way Access with Symmetric Connectivity

We have summarized the transition probabilities for this group in Table 4.2. The first row is the idle state with transition probability equal to the probability that neither of the nodes is transmitting. The second row refers to a single access state leading to a successful transmission by A , where f_A denotes the duration of a data

*In our presentation, we assume that the packet loss for control messages is zero, while data packets can potentially be lost due to channel conditions and noise at receiver. Considering the packet loss of control messages, doubles the number of states. Yet, we show that this assumption still allows accurate prediction.

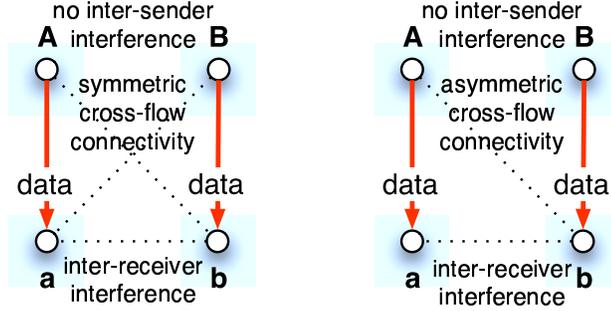


Figure 4.8 Snapshot of embedded flow coupled with symmetric and asymmetric cross-flow connectivity.

packet size in number of mini-slots. Note that the corresponding unsuccessful single access state has the probability $\gamma_i(1 - \gamma_j)^{f_A}l_A$ where l_A is the data packet loss probability. The same two probabilities are also present for successful or failed transmission by B , which we have omitted but can be easily obtained due to symmetry. On other other hand, if the state is neither an idle state nor a single access state, then it would be an overlapping state. We denote the probability that the system enters such a state by P_θ , which is equal to:

$$P(\theta) = 1 - (1 - \gamma_i)(1 - \gamma_j) - (\gamma_i)(1 - \gamma_j)^{f_A} - (\gamma_j)(1 - \gamma_i)^{f_B} \quad (4.2)$$

When packets overlap, the one that arrives first can be captured if its relative SNR is high enough to capture. On the other hand, the late overlapping packet can only be captured if the relative SNR is higher *and* it arrives during the synchronization bits of the header of the first packet (see Section 4.2). Let us denote the probability that X 's data packet is captured over Y 's transmission at x by $C_{XYx,D}$. The third

row of our table denotes a successful transmission by A when it overlaps with B . This probability is equal to the probability that packets overlap, A 's packet arrives earlier and is captured, and that A 's packet is not lost. The probability that B 's packet arrives later than A conditioned that they overlap, is denoted by S_{AB} and is equal to:

$$\frac{(1 - \gamma_j) + \dots + (1 - \gamma_j)^{f_A - 1}}{(1 - \gamma_j) + \dots + (1 - \gamma_j)^{f_A - 1} + (1 - \gamma_i) + \dots + (1 - \gamma_i)^{f_B - 1}} \quad (4.3)$$

To State	Probability
$0, j + 1$	$P_\theta S_{BA} O_{BA} C_{ABa} (1 - C_{BAb}) (1 - l_A)$
$0, j + k$	$P_\theta S_{BA} O_{BA} C_{ABa} C_{BAb} C_{Aba} C_{abA} (1 - C_{baB}) P_k(rts) C_{ABa,D}^{U(k-1)} (1 - l_A)$
$i + k, 0$	$P_\theta S_{BA} C_{BAb} (1 - O_{BA} C_{ABa} C_{Aba} C_{baB}) P_{k-1}(rts) C_{BAb,D}^{U(k-1)} (1 - l_B)$
$i + k, 0$	$P_\theta S_{BA} C_{BAb} O_{BA} C_{ABa} C_{Aba} C_{baB} (1 - C_{abA}) P_{k-1}(rts) C_{BAb,D}^{U(k-1)} (1 - l_B)$
$0, 0$	$P_\theta S_{BA} C_{BAb} O_{BA} C_{ABa} C_{Aba} C_{baB} C_{abA} C_{BAb,D} C_{ABa,D} (1 - l_B) (1 - l_A)$

Table 4.3 Four-way access with symmetric connectivity.

A 's packet would be lost if the packet is not captured or due to channel conditions. This results in a backoff stage of $(i + 1, j + 1)$ and is calculated in the fourth row of Table 4.2. The same probabilities can be easily calculated for B 's transmission due to symmetry. Note that since the late packet can arrive anywhere during the transmission of the first packet, the probability that it arrives during the synchronization bits of the first packet is very low, and hence for basic access we assume it is never captured. Finally, we assume that with overlapping packets the late packet arrives in

the middle of the first packet, and we select the overall length as the state duration.

Four-way Access with Symmetric Connectivity

We now calculate the transition probabilities when the RTS/CTS mechanism is used. The idle and single access states can be calculated the same as the basic access by replacing f with f' , where f' is the duration of RTS packet size in mini-slots. As a result, P_θ in Eq. 4.2 would correspond to the overlapping RTS packets probability. Fig. 4.9 depicts possible combinations in which B 's RTS arrives earlier than A . This probability is denoted by S_{BA} and can be calculated by Eq. 4.3 with f' instead of f .

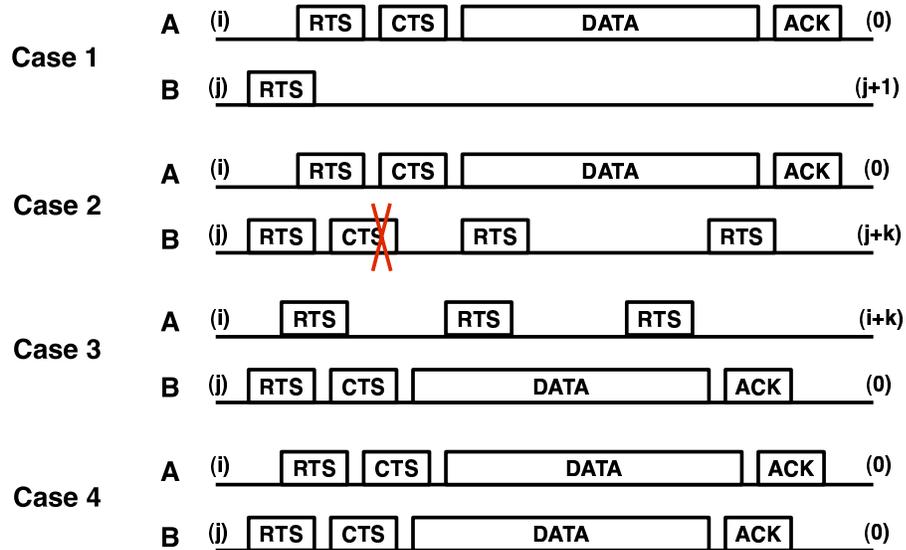


Figure 4.9 Four cases for symmetric cross-flow connectivity based on timing and capture behaviors.

For each case in Fig. 4.9, a corresponding row in Table 4.3 calculates the transition probability. In the first case, A 's RTS arrives during the synchronization bits of B 's RTS. O_{BA} calculates the probability of this event, conditioned that the two packets

overlap:

$$\frac{1 + (1 - \gamma_i) + \dots + (1 - \gamma_i)^{s-1}}{1 + (1 - \gamma_i) + \dots + (1 - \gamma_i)^{f_B-1}} \quad (4.4)$$

In the above equation, s is the duration of header synchronization bits in mini-slots (6 for 802.11b). For this case, the RTS packet of A is captured (C_{XYx} denotes the probability that X 's control packet is captured over Y at x), whereas b receives collision. Furthermore, the CTS packet transmitted by a is received by B , and hence B will defer. The first row in Table 4.3 calculates the corresponding event probability with successful transmission.

Rows 2 and 3 in Table 4.3 correspond to second and third cases of Fig 4.9. In the second case, CTS is transmitted by b , however it is not received by B . Row 4 of the table corresponds to the third case, as if a CTS was transmitted by a but not received by A . Finally, the two overlapping packets can both be captured, resulting in simultaneous data packet transmission by both A and B where the probability they both succeed is denoted by the fifth row of the table. For cases 2 and 3, RTS packets can be retransmitted after a TIMEOUT. However, these RTS packets will not be captured by their receiver, since the other flow is transmitting a data packet and RTS packets arrive in the middle of its transmission. If $k - 1$ further RTS packets are transmitted by i , the final backoff stage of i at the end of the transmission of the other flow will increase by k . On the other hand, a successful reception of data packet

in these cases requires that the data packet is captured over the RTS packets. This capture probability is always in the power of unity function since it is only present when there is an RTS retransmission during data packet reception.

Thus, the last term that we need to calculate is the probability of k RTS retransmissions for the transmission duration of the other flow. If i is the transmitting node, this probability is equal to:

$$\begin{aligned}
P(k \text{ retransmissions}) &= \sum_{m_1=0}^{L-(k-1)r-1} (1 - \gamma_i)^{m_1} (\gamma_i) \times [\\
&\quad \sum_{m_2=0}^{L-(k-1)r-m_1-1} (1 - \gamma_{i+1})^{m_2} (\gamma_{i+1}) \times [\dots \times [\\
&\quad \sum_{m_k=0}^{L-(k-1)r-1-\sum_{i=0}^{k-1} m_i} (1 - \gamma_{i+k-1})^{m_k} (\gamma_{i+k-1}) \times \\
&\quad (1 - \gamma_{i+k})^{\frac{((L-kr-1-\sum_{i=1}^k m_i)+(L-kr-1-\sum_{i=1}^k m_i))}{2}}]]] \quad (4.5)
\end{aligned}$$

In the above expression, L is equal to the number of available transmission opportunities in mini-slots, and r is equal to RTS + TIMEOUT duration in mini-slots (RTS retransmissions only occur after a timeout). Eq. 4.5 calculates the retransmission probability based on the fact that the backoff stage of i would increase after each retransmission. The above expression divides the whole duration into k parts, each of size $m_j + 1$, where i is not transmitting in the first m_j slots and transmits in the last slot.

It's important to note that the states in Table 4.3 account for successful transmission by the flow winning RTS/CTS, and hence there must be a corresponding

state that accounts for unsuccessful data transmission. These probabilities can be calculated from their successful transmission counterparts. Moreover, we note that the backoff stage of the node with an unsuccessful data transmission will not change in the four-way access due to the MAC protocol in IEEE 802.11 [29]. In case that flow A 's RTS is leading, the successful and unsuccessful probabilities can be easily calculated from Table 4.3 due to topological symmetry. Finally, a last state should be considered to account for overlapping RTS transmissions where none is captured. The resulting backoff stage would be $(i + 1, j + 1)$, and its probability is equal to 1 minus the summation of all other states.

To State	Probability
$0, j + 1$	$P_\theta S_{AB}(1 - l_A)$
$0, 0$	$P_\theta S_{BA} C_{BAb,D} C_{Aba,D} (1 - l_A)(1 - l_B)$
$0, j + 1$	$P_\theta S_{BA} (1 - C_{BAb,D} + C_{BAb,D} l_B)(1 - l_A)$
$i + 1, 0$	$P_\theta S_{BA} C_{BAb,D} (1 - C_{Aba,D} + C_{Aba,D} l_A)(1 - l_B)$

Table 4.4 Two-way access with asymmetric connectivity.

Two-way Access with Asymmetric Connectivity

In this topology, a will not receive B 's transmissions, whereas b receives transmissions by A and a (refer to Fig. 4.8). The overlapping states' probabilities of this group are summarized in Table 4.4, while the non-overlapping states remain the same as the two-way symmetric scenario. As calculated in the first row, if A 's packet overlaps with B and arrives earlier, it will be received at its receiver, while the late packet

will be dropped. However, in case that B 's packet arrives earlier, three different scenarios can happen: simultaneous successful transmissions or successful transmission by either flow and loss by the other flow. To derive these probabilities, we assume that A 's data packet finishes after B 's data packet transmission. This assumption, though accurate for the same modulation rate, will be slightly inaccurate with different modulation rates. Any other overlapping will result in backoff stage increases by both flows, where the probability is equal to one minus the summation of all other probabilities. Finally, we assume that with overlapping packets the late packets arrive in the middle of the other flow's transmission and take the overall length as the state duration for each case.

To State	Probability
$0, j + k$	$\gamma_i(1 - \gamma_j)^{F_A} P_k(rts)(1 - l_A)$
$0, j + k$	$P_\theta S_{BA}(1 - C_{BAb})(1 - l_A)$
$0, j + k$	$P_\theta S_{AB}(1 - O_{AB}C_{BAb}C_{Bab})P_{k-1}(rts)(1 - l_A)$
$i + k, 0$	$P_\theta S_{BA}C_{BAb}(1 - C_{Aba})O'_{BA}C_{bAa}P_{k-1}(rts)C_{BAb,D}^{U(k-1)}(1 - l_A)$
$0, 0$	$P_\theta S_{BA}C_{BAb}(1 - C_{Aba})(1 - O'_{BA}C_{bAa})(1 - P_0(rts)) \cdot C_{BAb,D}C_{Bab,D}(1 - l_B)C_{Aba}(1 - l_A)$
$0, 0$	$P_\theta S_{BA}C_{BAb}C_{Aba}(1 - C_{abA})(1 - P_0(rts)) \cdot C_{BAb,D}C_{Bab,D}(1 - l_B)C_{Aba}(1 - l_A)$
$i + 1, 0$	$P_\theta S_{AB}O_{AB}C_{BAb}C_{Bab}C_{baA}(1 - l_B)$
$0, 0$	$P_\theta S_{AB}O_{AB}C_{BAb}C_{Bab}(1 - C_{baA})(1 - C_{baA})(1 - P_0(rts)) \cdot C_{BAb,D}C_{Bab,D}(1 - l_B)C_{Aba}(1 - l_A)$

Table 4.5 Four-way access with asymmetric connectivity.

Four-way Access with Asymmetric Connectivity

The main transition probabilities of this group are summarized in Table 4.5, and a sample of timeline graphs are plotted in Fig. 4.10. The first three rows of the table correspond to the case that A 's packets are captured, while the other flow retransmits RTS. This can happen in a single access state or overlapping states with either node transmitting earlier.

The fourth row of the table corresponds to the probability that the CTS packet transmitted by b is captured over A 's RTS at a . As plotted in the second case of Fig. 4.10, future RTS transmissions by A will not be replied by a since it is able to set its NAV timer correctly. Hence, the backoff stage of A will increase. Note that the CTS packet should arrive during the synchronization bits of A 's RTS, and this probability is denoted by O'_{BA} which can be similarly derived as Eq. 4.2. On the other hand, if the CTS packet is not captured at a , different states happen depending on further attempts of A and on successful or failed transmission of each flow. One such case is plotted in case 3 and calculated in the fifth row, where we have presented the probability of a further attempt which results in a simultaneous successful transmission. Other states include no further attempt by A and all combinations including loss of one or both flows. The next row of the table calculates the same probability if the first RTS attempt was replied with an unsuccessful CTS packet.

In the last case of Fig. 4.10, node A receives the CTS packet transmitted by b ,

and no further attempt is made. This probability is calculated in the seventh row of Table 4.5. If the CTS packet is not received, different outcomes can happen similar to the third case. The last row of the table calculates the probability if A makes another attempt which has a successful data transmission. With certain capture probabilities, two flows can have simultaneously winning RTS/CTS transmissions and hence data packet transmission. These probabilities can be calculated by plotting the timeline graphs.

We emphasize that all the probabilities presented in Table 4.5 assume a successful transmission by the flow with winning RTS/CTS, while unsuccessful transmissions can be derived from them. Finally, a collision state should be added to the system where the backoff stage and its probability are calculated similar to the symmetric scenario.

Throughput Calculation

By numerically solving the Markov chain for each access mechanism and topology, which is ergodic for any choice of parameters, we obtain the stationary distribution $\pi = \pi_i, \forall i$. Long-term performance metrics such as throughput can be obtained directly from the solution of the Markov chain. From renewal-reward theory, the throughput of either flow is given by:

$$T = \frac{\sum_n \pi_n P_{S_n}}{\Delta} \quad (4.6)$$

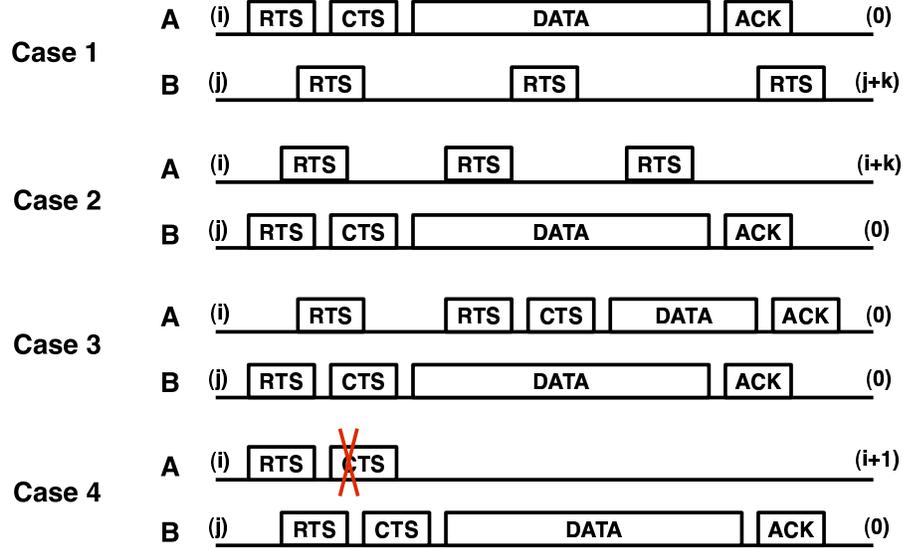


Figure 4.10 Four cases for asymmetric cross-flow connectivity based on timing and capture behaviors.

Here, P_{S_n} is the probability of successful transmission of either flow at state n , and Δ is the average duration of a step. Δ is computed from the average of the duration of all possible events in all states, weighted by their respective probabilities.

Modeling Sources with Uncoupled Backoff

If the uncoupled source is within the sensing range of the embedded flow, the senders sense each other's transmissions and defer when one is sending, which can be predicted by prior models [26]. Hence, we focus on a hidden, uncoupled source. Furthermore, as the backoff evolutions are uncoupled, we use a decoupling technique to model the behavior of the embedded link. Namely, we model the private channel view evolution of the sender of flow i as a renewal process with three different states: (i) idle channel; (ii) channel occupied by a successful transmission of the embedded

link, and (iii) channel occupied by the collision of the node.

Let p be the probability that the transmission of a station is not successful. Then, the occurrence of the channel states are: $\Pi_\sigma = 1 - \gamma_i$, $\Pi_c = \gamma_i p$, and $\Pi_s = \gamma_i (1 - p)$. Using standard renewal-reward theory, the throughput of the node is given by:

$$T_P = \frac{\gamma_i(1-p)}{\Pi_s T_s + \Pi_c T_c + \Pi_\sigma \sigma} \quad (4.7)$$

Now, the transmission attempt probability, γ_i , is a deterministic function of p given by [26] and is equal to:

$$\gamma_i = \frac{2q(1-p^{m+1})}{q(1-p^{m+1}) + W_0[1-p-p(2p)^{m'}(1+p^{m-m'}q)]} \quad (4.8)$$

where $q = 1 - 2p$, W_0 is the minimum window size, m is the *maximum retry limit*, and m' is the backoff stage at which the window size reaches its maximum value ($m' \leq m$). The average duration of successful transmission or collision can be computed *a priori* [26]. Thus, the only unknown variable in Eq. 4.8 is the conditional packet loss probability, p . To solve this problem, we assume that the transmission process of the hidden node is as an on-off process, where the on period is equal to packet transmission time which is fixed for a given modulation rate and packet size. The off process duration is an exponential random variable with an average duration of $\bar{T}_{OFF} = \frac{1}{\lambda_j}$, where λ_j is the packet arrival rate at the hidden node. We further assume that the transmission attempts of i happen randomly in the on-off process. Thus, we

have:

$$p = 1 - \frac{\bar{T}_{OFF}}{T_{ON} + \bar{T}_{OFF}} e^{\frac{-d}{\bar{T}_{OFF}}} - C_{ij} \frac{\bar{T}_{OFF}}{T_{ON} + \bar{T}_{OFF}} (1 - e^{\frac{-d}{\bar{T}_{OFF}}}) \quad (4.9)$$

where d is the duration of the packet transmitted by i . In the above equation, successful transmissions occur when the first packet arrives and fits completely into the idle period of the on-off process or arrives first, overlaps with its transmission and is captured.

Handling Non-Saturated Flows

So far in our analysis, we have assumed that when the backoff counter of a flow reaches 0, the transmitter always sends a data packet, i.e. the senders are fully backlogged. We now extend our analysis to the case that the packet arrival rate of each flow i is λ_i . We define a new probability ρ_i , which is the probability that the sender has a data packet to send when it is attempting to transmit a packet and replace γ_i in our prior equations with $\gamma_i \times \rho_i$. With saturated throughput, ρ_i is equal to 1, but with unsaturated throughput it must be calculated such that the achieved throughput of a flow i is less than or equal to λ_i . This value can be easily obtained when embedded flow i is competing with an uncoupled backoff source through Eq. 4.7. However with coupled sources a closed form expression for ρ_i that yields throughput equal to λ_i does not exist. Hence, we approximately set $\rho_i^{new} = \alpha \times \rho_i^{old} + (1 - \alpha) \times \frac{\lambda_i}{T_i^{old}}$ and adopt a global iterative procedure to update it where, during each iteration, we utilize the throughput analysis to update the variables of every node as a function of

its neighbors' ρ values (as computed in the previous iteration). The procedure ends when the throughput achieved by each flow is less than or equal to its demand.

4.4 Topological Profile Inversion and the Bi-Modal Effects of 1 dB

In this section, we perform thousands of measurements of embedded links in an urban mesh network to both validate our model and experimentally analyze the complex factors that contribute to topological profile inversion. With our validated model, we explore the full set of interdependencies that lead to this behavior and show that reverse capture plays a critical role. Further, we experimentally show that this inversion can be based on only 1 dB of link variation, causing a bi-modal shift as channels fade, and our model is able to reveal the reasons for such an empirical result.

4.4.1 Experimental Set-up and Measured Model Inputs

In our experiments, we activate two fully-backlogged UDP flows (Aa and Bb) with 1500B packets. We repeat the experiment in 120 second intervals for all combinations of 802.11b rates and for both access mechanisms. Each sender to receiver link is strong, enabling the highest modulation rate to have a high delivery ratio (i.e., 11 Mbps performs well on the flow's links). Before the experiment, we measure the data packet loss probability per modulation rate for each flow in isolation for our model. During the throughput experiments, we perform per-second SNR measurements and use the average relative SNR per link pair and our capture measurements from Section

4.2 to find the corresponding capture probability. Tables 4.6 and 4.7 describe the average relative SNR for each possible competing link pair for the topologies. In some cases, one of the two competing links lacks connectivity which results in a capture probability of 1 for the other link. We denote this as +12 dB or -12 dB, recalling from Section 4.2 that 12 dB is sufficient to completely capture regardless of the packet size and modulation rate.

4.4.2 Baseline Scenario: Symmetric Cross-flow Connectivity

As a baseline for our model validation, we first consider the throughput of an embedded link which has symmetric cross-flow connectivity with an interfering source which has been shown to fairly share bandwidth in idealized channel conditions and equal modulation rates [24, 27]. While this topology has *symmetric* cross-flow connectivity, there is vast heterogeneity in channel conditions between the flows which results in diverse capture characteristics based on the packet size and modulation rate. As an example, in Table 4.6, the SNR difference between A and B ($A - B$) is -3.2 dB at b , and $A - B$ is +0.6 dB at a . Hence, with overlapping control packets transmitted by A and B , the probability for B 's control packets to win capture at b is 0.98 whereas the same packets are likely to collide at a (A 's probability to win capture at a is only 0.25 even with control packets). We now consider how accurate our model is at predicting an embedded link's throughput with such topological complexities.

Fig. 4.11(a) and 4.11(b) depict the throughput achieved by each flow using the

RX	Relative SNR (dB)			
	$A - B$	$A - b$	$a - B$	$a - b$
A	-	-	+12	+1.0
a	+0.6	-1.4	-	-
B	-	-12	-	+3.9
b	-3.2	-	+1.0	-

Table 4.6 Symmetric cross-flow connectivity sub-topology where a positive value favors Aa and negative, Bb .

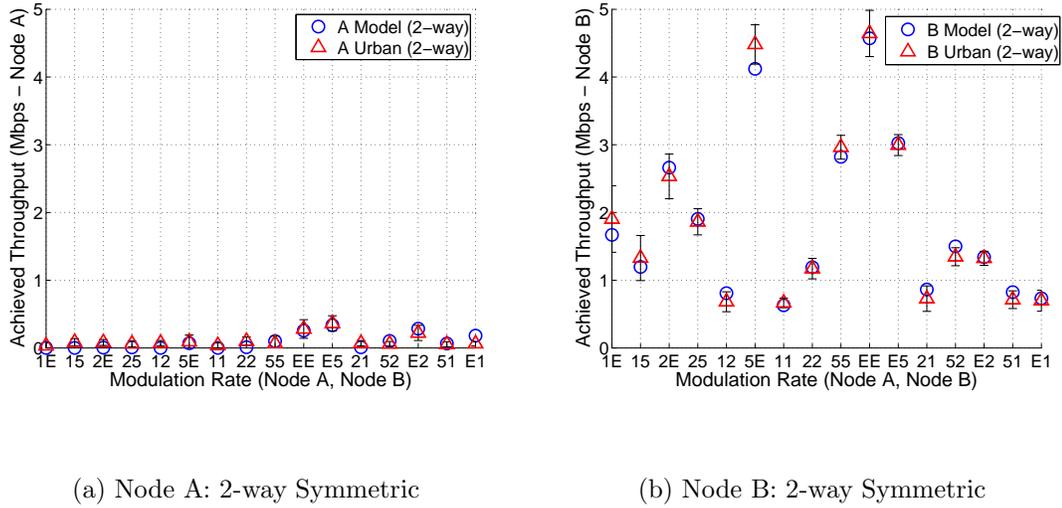
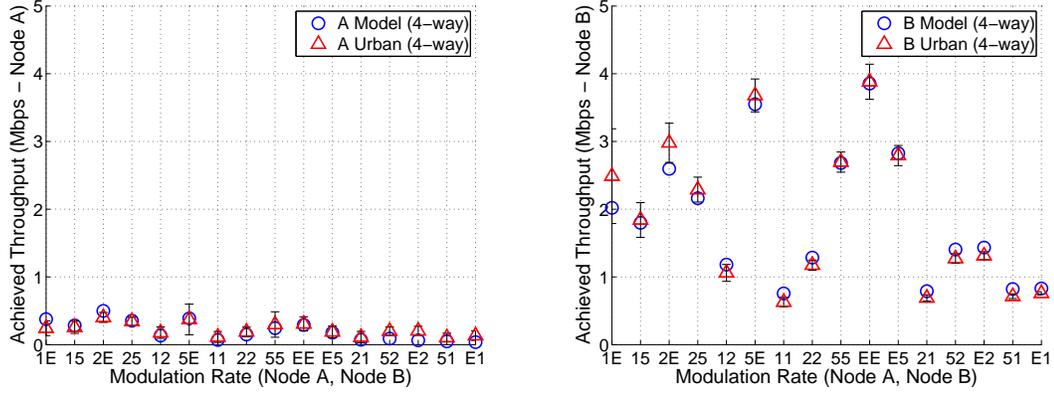


Figure 4.11 Symmetric cross-flow connectivity (from Table 4.6) with basic access.

two-way handshake. We observe that the embedded link model provides an excellent match with measurement results for all combinations of modulation rates. Furthermore, the throughput for flow Aa is near zero for all rate combinations. Our per-second throughput measurement ensured that Aa had a high-quality channel throughout the experiment. In fact, before and after the experiment, flow Aa is able to achieve over 4 Mbps with the highest modulation rate. We find that this low throughput for Aa is caused by physical layer capture of flow Bb over traffic from



(a) Node A: 4-way Symmetric

(b) Node B: 4-way Symmetric

Figure 4.12 Symmetric cross-flow connectivity (from Table 4.6) with the four-way handshake.

nodes A and a in the forward traffic direction. Since the four-way handshake reduces the time packets overlap, we now examine the performance of the same TFA topology and the respective prediction of our model.

Fig. 4.12(a) and 4.12(b) depict the same experiment as before performed an hour later with the four-way handshake enabled. We observe that the achieved throughput remains severely imbalanced. The reasons for this are revealed by our model. In order to have a successful transmission by flow Aa , its RTS transmission should not overlap with RTS transmissions of B . Moreover, if B 's RTS arrives earlier, it will be captured while if A 's RTS overlaps with B , it will be dropped. As a result, A 's backoff stage will continuously increase whereas B 's backoff stage remains close to zero. In summary, for both access mechanisms with symmetric cross-flow connectivity, forward traffic

capture causes a bi-modal shift in the traffic profile. Later, we discuss the additional effects of capture in the reverse direction.

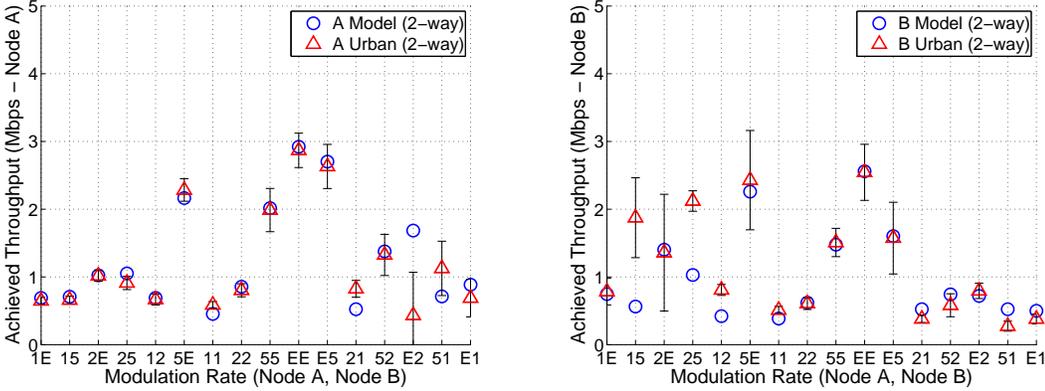
4.4.3 Inverted Traffic Profile for Asymmetric Cross-flow Connectivity

We now consider an embedded link that competes with an interfering source with asymmetric cross-flow connectivity. Under perfect channels, this case will yield one flow starving due to lack of information [24, 27]. Similar to the table for the symmetric sub-topology, Table 4.7 describes the competing links within the sub-topology. Since the cross-flow connectivity is asymmetric, the link between a and B is not able to carry packets and the value in the table for $a - b$ at B is less than -12 dB. We repeat the same experiment with this grouping of nodes and channel configuration.

RX	Relative SNR (dB)			
	$A - B$	$A - b$	$a - B$	$a - b$
A	-	-	+12	+3.8
a	-7.1	-2.8	-	-
B	-	-12	-	-12
b	-1.6	-	-2.6	-

Table 4.7 Asymmetric cross-flow connectivity sub-topology (positive value favors Aa and negative, Bb).

Fig. 4.13(a) and 4.13(b) report the model's prediction and corresponding throughput measurement from the two embedded flows in TFA when using the two-way handshake and asymmetric cross-flow connectivity. Surprisingly, the flow without information (Bb) is able to achieve equal throughput with the flow with information (Aa). Thus, the topology has an inverted traffic profile from [24, 27]. This can be



(a) Node A: 2-way Asymmetric

(b) Node B: 2-way Asymmetric

Figure 4.13 Asymmetric cross-flow connectivity (from Table 4.7) with basic access.

explained by the much larger ability to capture at b for B and A 's inability to capture at it's own receiver versus b .

Fig. 4.14(a) and 4.14(b) report the corresponding measurement and modeling throughput results for the four-way handshake and asymmetric cross-flow connectivity. Similar to the two-way handshake, we observe that the topology profile is inverted. However, in this case, there are many more dependencies that are required to allow equal sharing that we explore in the next section. In short, the joint presence of forward *and* reverse traffic is required to invert the imbalanced sharing of the topology, making it balanced. For example, this is seen here when the competing transmitters' modulation rates are equal. *Finding:* With asymmetric cross-flow connectivity, forward capture inverts the traffic profile for basic access. However, forward and reverse capture is required to invert the traffic profile for the four-way handshake.

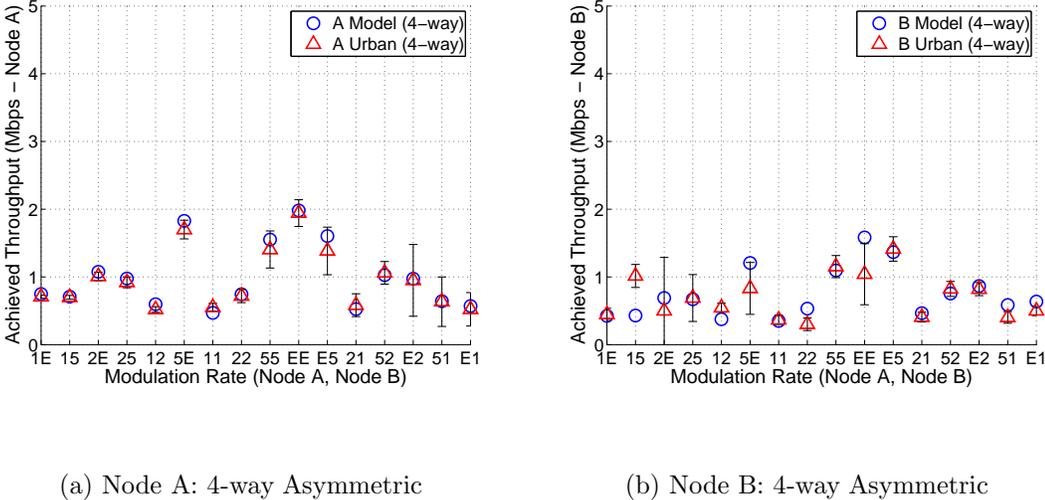


Figure 4.14 Asymmetric cross-flow connectivity (from Table 4.7) with the four-way handshake.

4.4.4 Asymmetric Profile Inversion: Effects of Forward and Reverse Capture

We now explore the full range of the aforementioned interdependencies to invert the traffic profile. To do so, we use the embedded link model and Jain’s Fairness Index, defined as $(\sum x_i)^2 / (n \cdot \sum x_i^2)$ where x_i is the achieved throughput of flow i and n is the total number of flows [30]. The fairness index of 1 corresponds to an equal throughput sharing whereas a fairness index of 0.5 corresponds to one flow starving and the other obtaining all the throughput.

For the two-way handshake, we showed that fairness occurred when the disadvantaged flow Bb (i.e., the flow which lacks information) is able to capture in the forward traffic direction (C_{BAb}). However, we have not yet considered the effect on the sharing when the advantaged flow Aa (i.e., the flow with full information) also has forward

traffic capture (C_{Aba}). Here, we present the results from our model where two coupled UDP, fully-backlogged flows compete with a modulation rate of 5.5 Mbps.

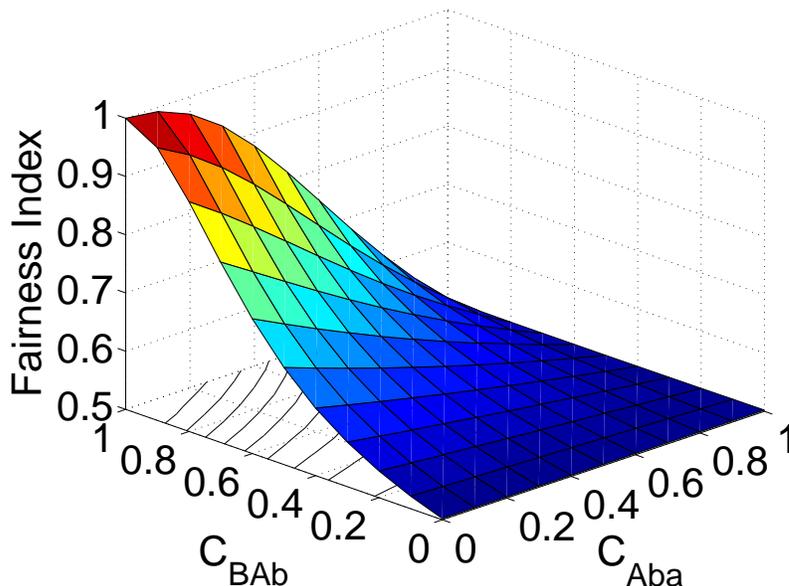


Figure 4.15 Asymmetric cross-flow connectivity with basic access.

Fig. 4.15 depicts a three-dimensional diagram of the fairness of two transmitters (A and B) according to their respective ability to capture at (a and b). In the left part of the figure, B (the disadvantaged node) is able to completely capture at b and A is unable to capture at a . This is the scenario that leads to perfect sharing for the two-way handshake. As A 's forward traffic capture (C_{Aba}) increases, the fairness index decreases rapidly and independent of B 's forward traffic capture value (C_{BAb}). *Finding:* With asymmetric cross-flow connectivity and basic access, inversion of the traffic profile primarily depends on the advantaged flow to lose forward traffic capture and secondarily depends on the disadvantaged flow to win forward traffic capture.

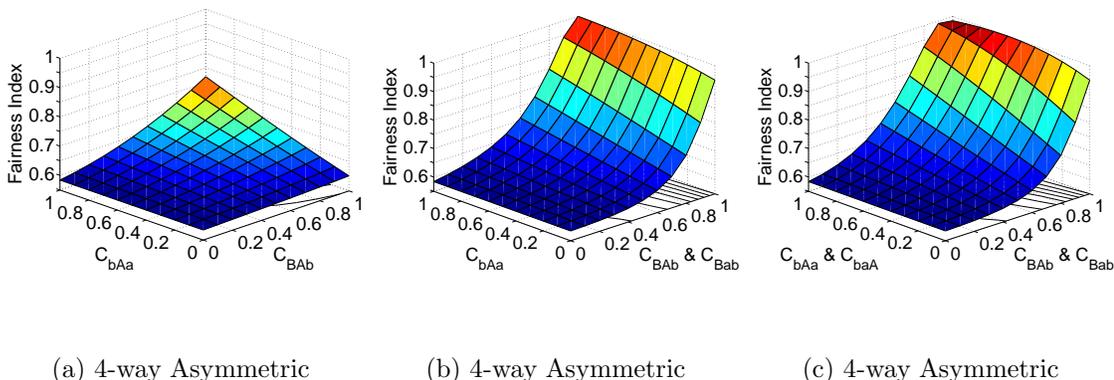


Figure 4.16 Asymmetric cross-flow connectivity with 2-way handshake (a) and 4-way handshake (b-d) where the disadvantaged flow has increased chance to win the contention due to favorable capture relationships.

Now, we consider the asymmetric case with the four-way handshake. Here, all four directions of capture (e.g., forward over forward, forward over reverse, reverse over forward, and reverse over reverse) must be considered since the RTS/CTS handshake preempts any data transmission. For the disadvantaged flow (Bb) the most important relationships for this embedded flow to equalize throughput sharing is the forward over forward traffic capture (C_{BAb}) and the reverse over forward traffic capture (C_{bAa}). We first present the results from the model for the same set-up with the four-way handshake with these two capture relationships (and later show other relationships that contribute to increased throughput of Bb).

Fig. 4.16(a) depicts the fairness index for asymmetric cross-flow connectivity with the four-way handshake based upon the ability of the disadvantaged flow Bb to capture in the forward direction versus competing forward traffic (C_{BAb}) and in the

reverse direction versus competing forward traffic (C_{bAa}). On the left and right side of the figure, near starvation of flow Bb occurs with the complete capture of the reverse or forward direction, respectively. However, in the middle of the figure, both relationships winning capture contribute to a much more equivalent throughput sharing. In Fig. 4.16(b), we add the ability of Bb to capture in the forward direction versus the reverse traffic (C_{Bab}). In the middle of the figure, we observe that a completely fair distribution of throughput (and complete inversion) can now be experienced. Finally, in Fig. 4.16(c), we add the ability of Bb to capture in the reverse direction versus reverse traffic (C_{baA}). Where these four capture relationships are 1 (middle of the figure), we observe that the flow Bb actually achieves greater throughput than Aa .

Finding: With asymmetric cross-flow connectivity and the four-way handshake, the disadvantaged flow requires a confluence of link capture relationships to cause profile inversion. Yet, when the profile does invert, the flow that is disadvantaged (in the information sense) can obtain even *higher* throughput than the advantaged flow, a behavior that does not occur with two-way handshake.

4.4.5 Reverse Capture Shifts the Traffic Profile for Symmetric Cross-flow Connectivity

Even when coupled flows are symmetrically connected, we showed that traffic profiles can be inverted for both types of access mechanisms. We now explore the impact that capture relationships have on the throughput sharing of the symmetric

cross-flow connectivity topology. We expect that the forward traffic capture would dominate the behavior of the two-way access, yet it is previously unstudied what role reverse traffic will have on the sharing, especially with four-way access. Here, we show the results for the model with the same traffic pattern as before but with symmetric cross-flow connectivity.

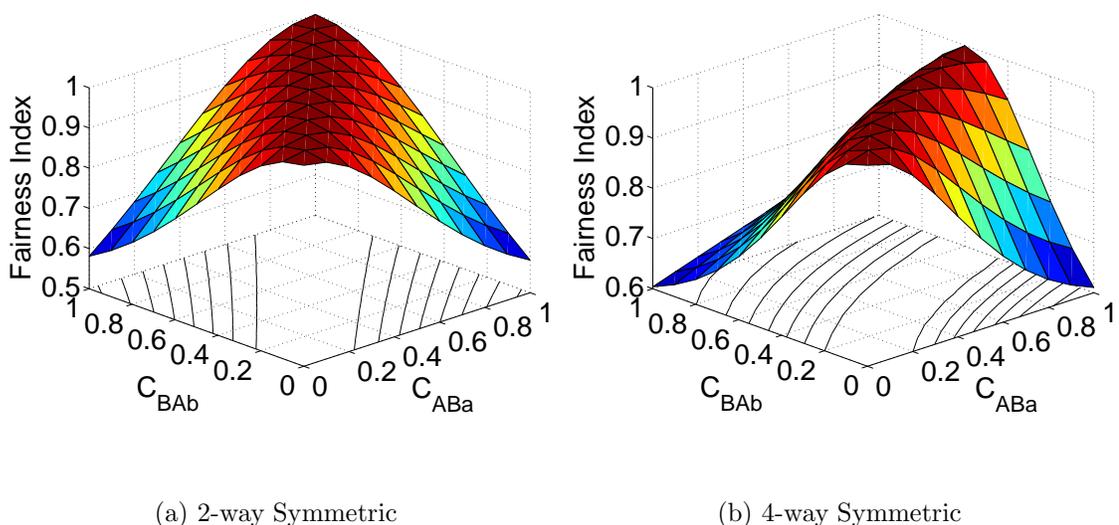


Figure 4.17 Symmetric cross-flow connectivity with (a) 2-way access and (b) 4-way access with both reverse captures of Bb to 1.

Fig. 4.17(a) depicts the fairness achieved by the two transmitters (A and B) based on their ability to capture at their own receivers (C_{BAb} and C_{ABa}) using the two-way handshake. The fairness property depends upon the symmetry of the forward traffic capture of the two flows. Prior work has predicted three points of this figure: [8] predicted the left and right corners of Fig. 4.17(a) (starvation mode) and [24, 27] predicted the result with no capture (fair-sharing mode).

With the four-way handshake, the fairness index remains nearly identical to Fig. 4.17(a). However, we emphasize that the capture requirements for each access mechanism is very different (i.e., for a given channel condition, RTS packets have much lower capture thresholds than data packets). *Finding:* For a given channel condition, use of two- vs. four-way handshake can yield a bi-modal shift because of both the lower modulation rate and smaller size of the RTS packet as compared to the data packet.

Fig. 4.17(b) depicts the throughput prediction from our model for the symmetric case with the four-way handshake where the reverse traffic is fully able to capture (C_{bAa} and C_{baA} are equal to 1). We observe that a shift in the sharing occurs favoring flow Bb (the flow which is able to capture in the reverse direction). For balanced throughput to be achieved in this case, flow Aa must have a greater forward traffic capture than flow Bb . *Finding:* Reverse traffic capture shifts the profile with symmetric cross-flow connectivity and the four-way handshake.

4.4.6 Link Variation of 1 dB Driving Bi-Modal Topological Profile Inversion

From our thousands of measurements over the course of a month on multiple topologies, we found many topologies to have highly varying throughput sharing. The vast differences are despite the use of off-peak times for our experiments and limited activity of other nodes in the mesh network. In a particular grouping of four

nodes with asymmetric cross-flow connectivity (described in Table 4.7), we found that the throughput sharing over a month's time period went from a starvation mode to a fair-sharing mode. While this is a larger time-scale (over the course of a month), it exposes many of the fluctuations that are happening on smaller time-scales (per-second).

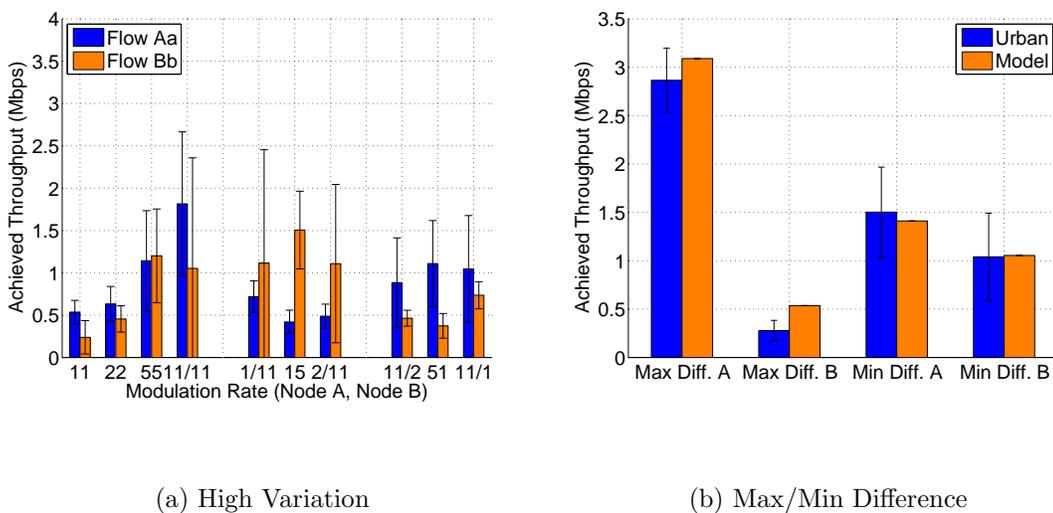


Figure 4.18 Asymmetric topology (from Table 4.7) over a month of measurements and two days' measurements where the maximum and minimum difference between sharing is achieved.

Fig. 4.18(a) depicts each flow's throughput over the course of a month where the average is represented by a bar, and the standard deviation is represented by error bars above and below the average.* We observe that across many different modulation rate combinations, both flows have highly varying achieved throughput. Fig. 4.18(b)(b)

*For all experiments, we measure all combinations of modulation rates, but for ease of exposition, we present only seven cases here.

shows single-day measurements where both the minimum and maximum difference between throughput sharing were achieved. With the minimum difference in sharing, nearly fair throughput is achieved, and the case with the maximum difference is highly imbalanced. By examining the differences in average SNR values between the two experiments, we observed a 1 dB relative difference in the pair of competing links in the forward traffic direction. Namely, when flow Bb is able to win forward traffic capture, it can achieve approximately equal throughput with flow Aa (as described with both access mechanisms in the validation experiment above). However, when B is unable to do so (e.g., when $A - B$ at b in Table 4.7 goes from -1.6 dB to -0.6 dB), there is a large throughput difference between flows Bb and Aa . Recall that the reverse traffic capture is present in this topology allowing the forward traffic capture relationship to make a difference. Therefore, 1 dB of change allows switching between the fair-sharing mode and the starvation mode. *Finding:* Even 1 dB of channel fluctuation can cause topological profile inversion.

4.5 Applications of Embedded Link Model and Experimentation

In this section, we apply our experimental analysis and model in two ways. First, we consider how modulation rate can be selected according to joint properties of channel condition, topology, and capture. Second, we predict and explain the disproportionate effect of low-rate control traffic on embedded data flows within a mesh

network.

4.5.1 Altered Rate Selection Problem Due to Capture

Prior work has considered the problem of choosing the modulation rate that achieves the highest throughput based on the channel condition from the sender to receiver [28]. However, no prior work has considered the interdependence of physical layer capture and modulation rate selection. We now apply our experimental analysis of different capture behaviors and our model to a scenario in which the embedded link competes with an uncoupled, hidden source which saturates the channel. Using the embedded link model, we fix the modulation rate (11 Mbps) and packet size (1500B) for the interfering transmitter. For the embedded flow, we vary the modulation rate (2, 5.5, and 11 Mbps) and packet size (100, 500, 1000, and 1500 bytes) with fully-backlogged, UDP traffic over a range of relative SNR.

Fig. 4.19 depicts the throughput for the embedded flow based upon its choice of modulation rate, packet size, and relative SNR to the interfering flow. At low relative SNR, no throughput is achieved as all packets are unable to capture against the interfering transmitter. As relative SNR increases, the throughput increases based upon the capture probability. Observe that the 5.5 Mbps rate is able to achieve the highest throughput out of all modulation rates for a relative SNR of at least 6 dB. This contrasts the throughput-maximizing modulation rate whenever the interfering source is off which would be 11 Mbps based on the link quality from sender to receiver.

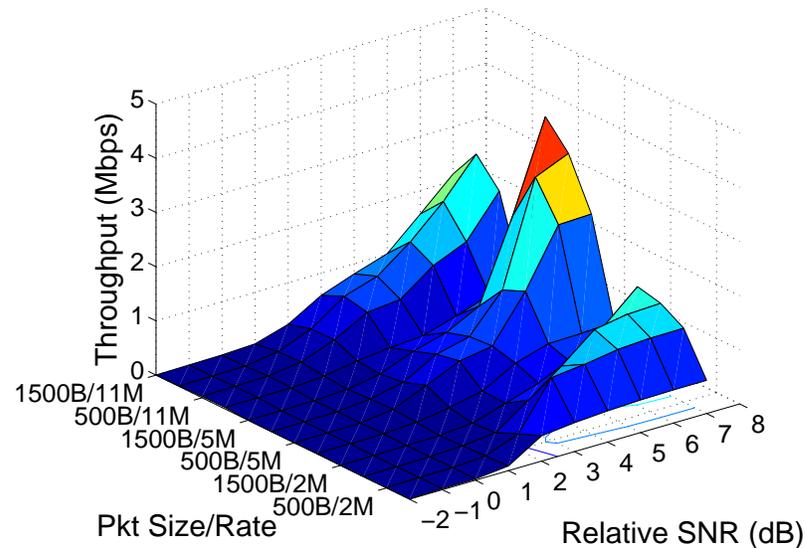


Figure 4.19 Embedded link throughput when hidden source with uncoupled backoff saturates the channel.

Further, it contrasts the throughput-maximizing modulation rate if the interfering packet was always on (i.e., the difference between the embedded flow’s packet and the interfering packet becomes the new channel condition) which would be the 1 Mbps [31]. The difference between both of these cases is that capture allows a temporarily “clean” channel until the receiver locks on and a “noisy” channel (up to the capture threshold) thereafter. Thus, the throughput-maximizing modulation rate is not 11 Mbps nor 1 Mbps but 5.5 Mbps due to its increased capture probability (from 11 Mbps) and increased rate (from 1 Mbps).

While this is a relatively small region of the graph, consider the relative capture thresholds for 802.11a as shown in [5]. Most of the delivery ratio curves for different modulation rates are completely orthogonal, meaning that the delivery ratio goes from

0 to 1 over orthogonal SNR regions. Thus, the rate selection problem with 802.11a would be based heavily upon this capture-induced dimension when competing with other transmitters as opposed to the channel condition between sender and receiver alone. *Finding:* With interfering flows, the throughput-maximizing modulation rate may be *lower* than the throughput-maximizing modulation rate allowed for the flow in isolation and *higher* than the throughput-maximizing modulation rate for a channel with constant noise.

4.5.2 Experimentally Discovering the Wireless Overhead Multiplier

Control messages inherently must be exchanged for network management protocols such as routing, client association, and backhaul link maintenance. We define all such messages as “overhead” and omit per-data-packet overhead such as packet headers, RTS/CTS messages, etc. Overhead can be generated from devices within a network (controllable by the network operator) or from external devices such as residential APs and clients (uncontrollable by the network operator).

While a well-understood tenet of protocol design is to restrict overhead traffic to a minimum, typically via use of low-rate periodic or on-demand small-sized messages, in this section, we show that despite having low rate, overhead can profoundly degrade network performance. Specifically, an overhead rate of λ can reduce the data throughput on a nearby link by up to 50 times λ . Here, we measure an initial scenario, characterize the overhead, and present our measurement methodology to show

the factors driving such effects.

Diverse Overhead Effects

To quantify the impact of overhead traffic on data throughput, we design an experiment in which we compare the throughput of a single link with and without the overhead induced by the surrounding nodes. In the experiment, we measure the *achievable throughput* defined as follows. Given a network N , and a sender-receiver pair $s, r \in N$, consider a fully backlogged flow $f_{s \rightarrow r}$ from node s to node r . The *achievable throughput* of the flow $f_{s \rightarrow r}$ is the throughput $t_{s \rightarrow r}^N$ achieved when all nodes in $N \setminus \{r, s\}$ only transmit overhead. Achievable throughput is defined for a particular protocol set (e.g., long-lived UDP flows with 1500 byte packets, over 802.11 with no RTS/CTS, and autorate enabled).

To eliminate known throughput degradation effects such as [2, 32], we first measure only single-active, one-hop flows where the user activity of the system is negligible. Furthermore, we concentrate on the effects of overhead only on high quality links (i.e., links that can send at the highest modulation rate). Thus, we measure the throughput degradation of each link's achievable throughput due to the injected backhaul overhead traffic.

In particular, we select a single one-hop backhaul node near the gateway (see Fig. 2.1) to send backlogged UDP traffic when all surrounding nodes are disabled and measure the UDP achievable throughput. We then measure the achievable throughput

of the same sender-receiver pair in the presence of overhead from surrounding nodes, i.e., neighboring nodes in the network are enabled but allowed to transmit only control traffic. For both measurements, we have identical hardware configurations for all nodes (200 mw transmit power, RTS disabled, autorate enabled), and hold the traffic type constant (1500 byte, constant bit rate, UDP traffic). We repeat the three-node experiment sequentially for each node that is one hop from the gateway.

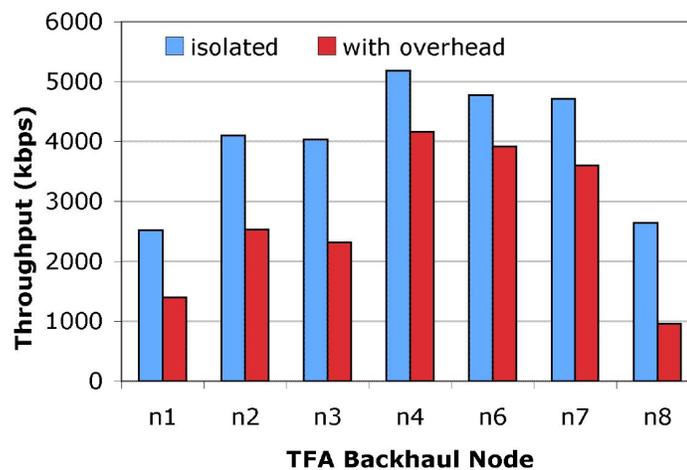


Figure 4.20 Achieved throughput with and without overhead (isolated) injected from the TFA network.

Fig. 4.20 shows the throughput degradation that each node experiences where the x-axis is the backhaul node and the y-axis is the achievable throughput. For each node, the left bar represents the achievable throughput in isolation (when no other nodes are transmitting overhead) and the right bar represents the achievable throughput when the network overhead is injected.

As can be seen in Fig. 4.20, there are two dramatic effects from the overhead.

First, there is a sizable portion of each achievable throughput in isolation that is lost on each link (at least $1/5$ of the throughput is lost on each link). Second, the throughput degradation is vastly different among links caused by the presence of overhead within the network. Specifically, the throughput degradation ranges from 850 kbps in the best-case to over 1700 kbps in the worst-case. Since the only difference between the setup of the measurements taken for the two bars for each node is the presence of TFA network overhead, the throughput degradation must be associated with the overhead injected by TFA. In order to verify these results, we repeated the same experiment on all channels. Indeed, in all channels, we observed the same trend which verifies the cause of the effect is the network overhead and eliminates the possibility that the two effects perceived in Fig. 4.20 are exclusive to the channel used by TFA or due to interaction with external networks operating on the same channel.

As previously explained, the throughput degradation is solely related to the overhead injected by the nodes of the TFA network. Furthermore, in all tested pairs the receiver is the gateway which sees the same number of transmitters (overhead-injectors) and the same environment (other noise-injectors, etc.) across all measurements. Also, since the hardware platform of all senders tested is identical (transmission power, autorate policy, RTS/CTS mechanism, etc.), the differences in throughput degradation caused by the overhead must be due to the location of each transmitter, i.e., topological differences seen by each transmitter. More specifically, the throughput degradation

experienced by each sender is correlated to the quality of the bidirectional links that are formed between the sender and the other nodes (receiver and overhead-injectors).

Heterogeneous Non-Forwarding Links

The only difference driving the heterogeneity in overhead effect is the varying spatial location of overhead-injecting nodes to the data-sender. These links between transmitters which are not intended to communicate directly (non-forwarding links) are inherent within the topology (i.e., not planned within the design of the forwarding links of the network). Thus, these non-forwarding links vary greatly in quality compared to the data-carrying or forwarding links. Such non-forwarding links impact the data transmission whether causing the node to defer at the transmitter or yielding simultaneous transmissions resulting in collisions or capture effects. For example, node n7 can cause node n4 to defer since the two nodes are able to decode one another's packets. On the other hand, n4 and n8 are unable to decode each other's packets or even sense each other on the medium and hence collide.

Thus, the difference in overhead effects is caused by the differing nature of these links between sender and non-receiver neighbors. We define the resulting connectivity matrix of vastly heterogeneous non-forwarding and forwarding links within a mesh topology as the *heterogeneous backhaul connectivity*. We now define a term to quantify the multiplicative overhead effects caused by the heterogeneous backhaul connectivity.

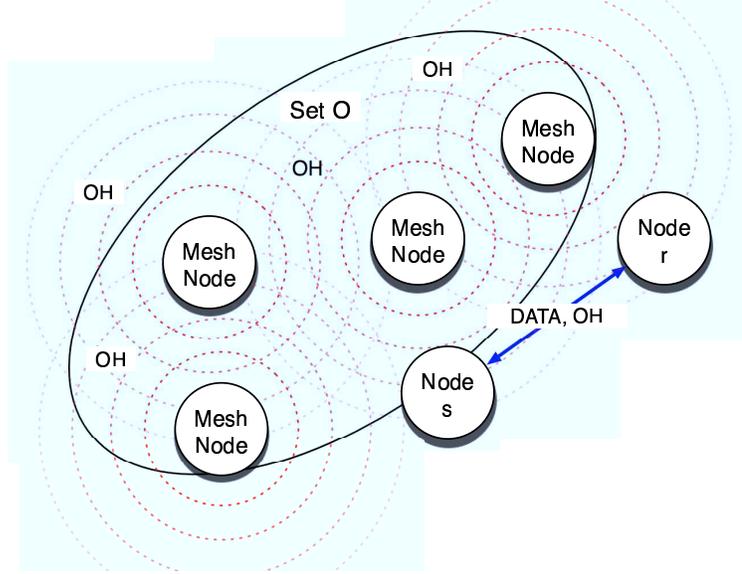


Figure 4.21 Node s sends data and overhead (OH) to node r while all other nodes within set O transmit only OH.

Wireless Overhead Multiplier

To formally define WOM, consider a sender-receiver pair s, r and a set O consisting of nodes which are primary interferers to s and r , as depicted in Fig.4.21. Let λ_O denote the cumulative mean rate of the overhead transmitted by the nodes in O . Of course, $t_{s \rightarrow r}^N$ is related to the set of active interfering nodes O , i.e., the more interfering nodes are active, hence transmitting overhead, the lower the expected $t_{s \rightarrow r}^N$. Accordingly, we evaluate the impact of the overhead due to O on flow $f_{s \rightarrow r}$ by comparing the achievable throughput $t_{s \rightarrow r}^N$ in isolation (i.e., $N = \{s, r\}$ and nodes in O are not active), with the value of $t_{s \rightarrow r}^N$ when the interferers are active (i.e., $N = \{s, r\} \cup O$):

$$W_{s \rightarrow r} = \frac{t_{s \rightarrow r}^{\{s,r\}} - t_{s \rightarrow r}^{\{s,r\} \cup O}}{\lambda_O} \quad (4.10)$$

Eq. (4.10) gives a measure of the achievable throughput degradation normalized to the injected overhead. Notice that, since overhead can be due to pure MAC frames (e.g., beacons) as well as to IP packets (e.g., routing messages), we include in λ only the MAC throughput, i.e., we take into account only the payload of overhead frames. Hence, ideally, the protocol overhead causes a degradation of the achievable throughput equivalent to the air-time utilization of overhead traffic, which is greater than λ_O . For example, a short unicast (90-byte) IP message sent at maximum modulation rate (11 Mbps lasting 58 μ s) incurs a per-packet overhead of a preamble (at 2 Mbps lasting 192 μ s) and a 30-byte MAC header (at 11 Mbps lasting 22 μ s) plus the 14-byte ACK (192 μ s for the PHY preamble, and 10 μ s for the ACK MAC frame at 11 Mbps). Hence, also considering the spacing between frames (SIFS and DIFS), a 90-byte packet flow uses a gross bandwidth of 11 Mbps to carry 1.3 Mbps, i.e., the overhead consumes 11/1.3 times its nominal bandwidth λ_O . Analogously, a 1500-byte payload transmitted at 11 Mbps yields an average transmission rate of approximately 7.9 Mbps, i.e., an actual overhead rate of $1.4 \cdot \lambda_O$. Thus, the ideal expected WOM value caused by acknowledged frames ranges from 1.4 to 8.5, depending on the size of the overhead payload.

However, we find that the WOM value can range from near 0 to over 50. Further,

we show that the WOM is controlled by effects due to the heterogeneity of the quality of all links formed between nodes s, r , and the interferers in O . In order to understand the basic interaction of links, we first investigate the WOM effect within a topology of three nodes: the data-sender (s), the data-receiver (r), and the overhead-injector (o). Thereafter, we show the compounding effects of more complex topologies.

4.5.3 Isolating Link Effects for the Wireless Overhead Multiplier

We now explore the effect of a single overhead-injecting node as a function of the link quality to a data-sender and a data-receiver. Following the 802.11 standard, nodes behave differently according to differing link qualities with respect to other transmitters. Thus, we classify links according to the transmitter behavior specified in IEEE 802.11 and isolate the overhead effects due to different node behaviors using a three-node topology as shown in Fig. 4.22. Node s represents the data-sender, node r represents the data receiver, and node o represents the overhead-injecting node. Links s, r and o, r both are able to achieve transmissions at the maximum rate, while link o, s can vary.

IEEE 802.11 Node Behaviors

The standard describes three different behaviors within for medium access: (i) if a node is able to decode a transmission of another sender, it NAVs (according to the physical or virtual carrier sense mechanism), (ii) if a node is able to detect channel

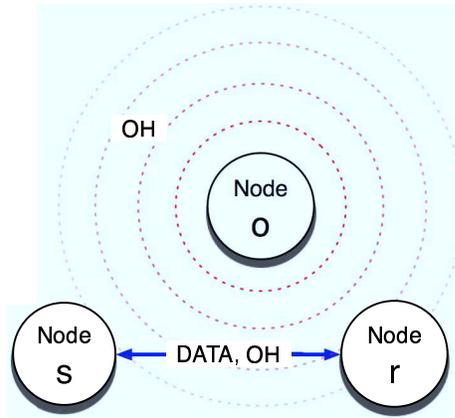


Figure 4.22 Node s sends data and overhead (OH) to node r while node o transmits only OH.

activity, it defers until the channel is free and additionally defers its transmission for an Extended Inter Frame Space (EIFS) period which covers the longest possible ACK duration, and (iii) if a node is unable to detect channel activity, it transmits according to the normal backoff mechanism. Correspondingly, we classify each node pair (i.e., the link between the two nodes) according to their degree of connectivity as defined by the standard: (i) transmission range, (ii) carrier sense range, and (iii) out of range. While variation in channel quality can cause links to change their class over time, each individual packet is within a single class according to the MAC behavior.*

IEEE 802.11 Off-the-Shelf Card Behavior

We begin our investigation by testing the off-the-shelf hardware for the node behavior described in the 802.11 standard to enable detection of a particular TFA link class. Determining that nodes are in transmission range can be achieved simply by

*For a statistical description of the links in TFA, refer to [33].

ensuring that beacons (sent at the base rate) are successfully received. However, distinguishing between carrier sense range and out of range classes requires experimentation since the MAC state machine is not directly observable. Thus, we next design an experiment to distinguish between these two classes.

If s adopts the energy detection behavior described in the standard, when the energy level is above a given threshold, it will defer transmission via physical carrier sensing. To experimentally find the energy detection threshold, we use the configuration depicted in Fig. 4.23 where the sender-receiver pair s, r of the data flow communicate over the air, and the noise generator n is hard-wired to the sender.

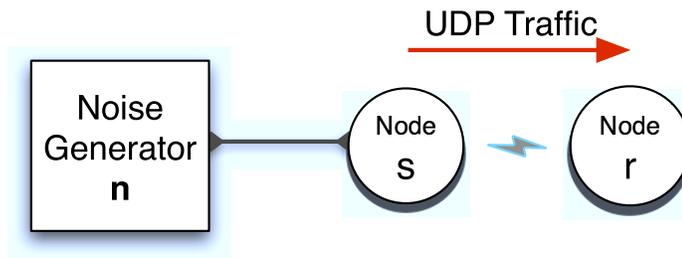


Figure 4.23 802.11-behavior experiment set-up for the off-the-shelf wireless cards.

In the experiment, we send a UDP flow from s to r at a constant physical layer rate of 2 Mbps. The link from s, r is held constant. We observe the behavior of the achievable throughput of s, r as a function of the noise level generated by n . Thus, any change in the throughput at r is caused by the behavior of s , i.e., if the noise is above an energy detection threshold, s defers, driving the achievable throughput to 0.

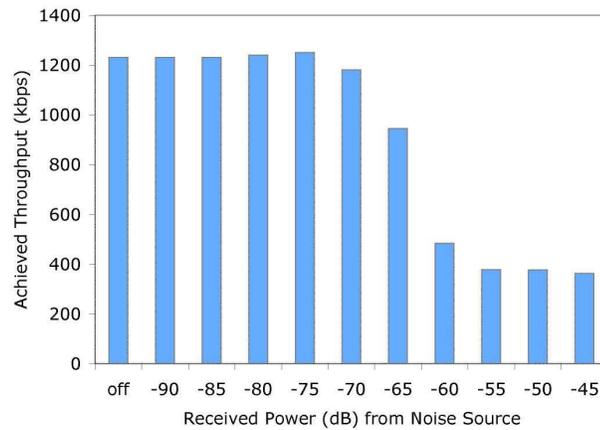


Figure 4.24 802.11 card behavior when noise is injected at the transmitter only.

Non-Existence of Carrier Sense Range

Fig. 4.24 depicts the throughput of the data flow from s to r where the x-axis is the level of generated noise. The noise source is a modulated sine wave within the spectrum of the 802.11 channel used in the experiment. We observe a dip from 1230 kbps (the achievable throughput when the noise source is disabled) to 360 kbps at -60 dBm. We observe with *Kismet* that the throughput decrease is due to the deafness of the transmitter to hear the ACK, leading to excessive backoffs and retries of the same application layer packet. Thus, there is no energy detection threshold.

Therefore, we find that the chipset/driver used in TFA (Prism/HostAP) defers only when another packet in the air is able to be decoded and will not defer due to noise alone. Note that this behavior is compliant with one of the modes available for CCA procedure described in the IEEE 802.11 standard [29], i.e., carrier sense

without energy detection. Furthermore, this is a common choice for vendors, e.g., another well-known chipset/driver, the Atheros/MadWiFi, operates in the same way. Hence, in TFA there are only two link classes.

WOM in Three-Node Topologies within TFA

In this section, we experimentally study the WOM effect defined by Eq. (4.10) on the TFA network in accordance with the TFA link classes. To achieve this, we systematically isolate three node topologies from the same nodes involved with the experiment shown in Fig. 4.20, all other nodes are disabled.

Specifically, we perform extensive measurements to form a data set from these three-node topologies consisting of both TCP and UDP data traffic of 1500 bytes from the data-sender s to the data-receiver r . As observed via *tcpdump* and *kismet*, the overhead traffic sent from the third node o consists of 90-byte packets (on average) at approximately $\lambda = 10$ kbps. More than 90% of the overhead traffic consists of unicast frames, and user traffic is negligible. AutoRate Fallback is enabled in all experiments unless otherwise specified.

Fig. 4.25 shows the average WOM values with error bars representing one standard deviation above the value for our data set where nodes are within transmission range and out of range with TCP data traffic. We find that the average WOM induced by an overhead-injecting node within transmission range is 4.6. Further, we observe that the nodes out of range exhibit a much larger average (11.8) and variance in the

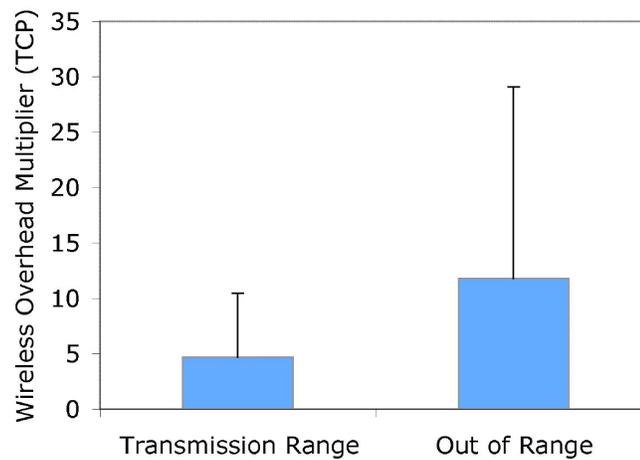


Figure 4.25 WOM considering the link class (transmission range or out of range) of the link between the WOM-inducing node o and the data-sender s .

WOM values as compared to the transmission range.

The results in the transmission range case are not surprising. In fact, due to the perfect coordination between the nodes (beyond the negligible propagation delay), the WOM value is in the range predicted for the ideal case discussed above. In contrast, when the data-sender and the overhead-injecting nodes are out of range, there is a lack of coordination that yields significantly larger WOM values as we now explore.

WOM in Out of Range Class

Within the out of range class, simultaneous transmissions occur causing various effects: (i) collisions resulting in loss, (ii) retransmissions, and (iii) the physical layer capture effect. We now describe these effects within the context of the MAC (i and ii) and PHY (iii) layers.

MAC Effects

In a CSMA MAC, simultaneous transmissions can collide at a mutual receiver, resulting in loss and retransmissions. Since the optional RTS mechanism was designed to avoid such collisions, we investigate the WOM effects with and without this collision avoidance mechanism.

With RTS disabled, the cost of a single retransmission is approximately one doubled backoff period plus the packet period. In the case of low-rate overhead, each data packet from s is unlikely to collide with more than one overhead packet. Hence, the collision rate of the system is approximately equivalent to the overhead packet injection rate. For example, to retransmit a 1500 byte packet at 11 Mbps, it takes approximately 2.2 ms, on average, including DIFS, SIFS, ACK, backoff and PHY overhead. Hence, an overhead of 90 byte packets at 10 packets per second (i.e., λ equals 7.2 kbps) reduces the rate of successful transmissions of s to r and yields a WOM value of over 20 for UDP traffic. However, because Fig. 4.25 indicates substantial variation from this point, hidden terminal collision effects alone are insufficient to characterize the WOM value.

With RTS enabled, the cost of collision is reduced to one doubled backoff period plus the RTS/CTS exchange duration. We now compare the aforementioned cost of collision to the cost of the additional signaling imposed by the use of the RTS/CTS mechanism to the gains of the reduced cost. To compare this, we show the case where

two nodes (n4 and n8) are out of range. We use n4 as the data-sender s and n8 as overhead-injector o . We measure the induced WOM with and without the use of the RTS/CTS mechanism with TCP data traffic.

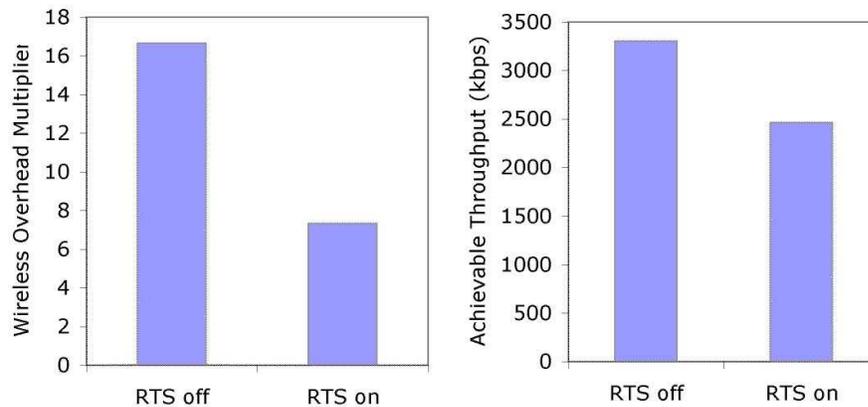


Figure 4.26 WOM (left) and aggregate TCP (right) considering use of RTS mechanism in an out of range scenario.

Fig. 4.26 (left) depicts measurements of WOM over multiple trials for node n4, with an out of range overhead-injector, n8. Indeed, the WOM is reduced by the use of the RTS/CTS mechanism. However, note that the protocol set for a given WOM has changed, thereby altering the achievable throughput used for reference flow. Since the RTS/CTS mechanism induces *per-packet* overhead, the use of the protocol set here with TCP traffic with RTS enabled has lower achievable throughput than TCP traffic with RTS disabled. The induced per-packet overhead of RTS used to reduce the WOM produces a net loss of aggregate throughput. Namely, the achievable throughput of n4 is 2.5 Mbps with RTS enabled and 3.3 with RTS disabled, *after* WOM is taken

into account. In summary, our measurements indicate that while use of RTS/CTS reduces WOM, its increased per-packet overhead yields a net throughput reduction for data traffic.

Joint PHY/MAC Effects

Throughput and MAC behavior are strongly influenced by physical layer capture [4]. Thus, we next establish the existence of capture in the TFA network and explore its impact on WOM.

First, since it has been shown that ARF causes throughput imbalances in the hidden terminal scenario [17], we fix the physical layer rate of the transmission to the base rate (2 Mbps) to eliminate these effects. Next, we measure the achievable throughput of each one-hop backhaul node s from the gateway r in isolation and in the presence of one out of range overhead-injector o . We also record the differences in SNR at the gateway between the two transmitters.

Fig. 4.27 shows the WOM value for each of the differences in SNR where a positive value indicates s has a more powerful SNR at r than o . The results indicate a bimodal relationship in the WOM values for the positive and negative SNR differences. More precisely, when the SNR difference is positive, the WOM value is approximately 1, indicating that the overhead losses experienced by the data sender are less than the actual injected overhead. However, when the SNR difference is negative, the WOM value ranges from 6 to 12. We conclude that capture effect occurs with a difference

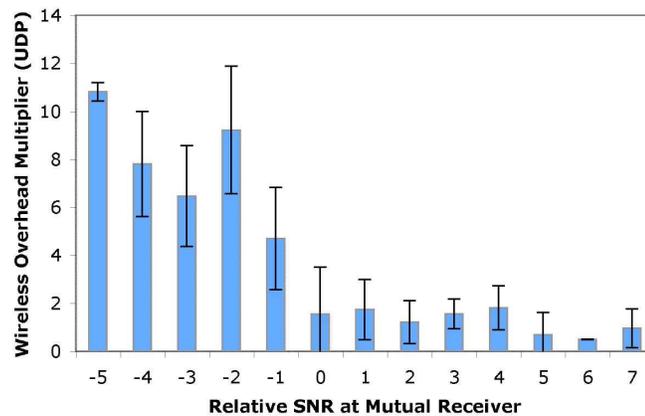


Figure 4.27 WOM of out of range links considering relative RSSI at data-receiver r from data-sender s and WOM-inducing node o .

in SNR of greater than 0.

To show that this bimodal behavior is due to the capture effect, we now evaluate the WOM associated with two specific out of range nodes. We use the first node n7 as a data-sender, and the second node n2 as an overhead-injector and measure the WOM. The SNR from n7 is 3 dB greater than from n2 at GW. We then repeat the experiment after switching the roles of the nodes.

Fig. 4.28 shows the WOM for the two experiments for both TCP and UDP fully-backlogged traffic and physical layer rate of 2 Mbps. Node n2 has a WOM value of 9.2 and 7.6 for UDP and TCP, respectively; while n7 has a WOM value of 0.9 and 0.6, respectively. Hence, we find that the severe asymmetry exists across both traffic types. Regardless of the traffic type, the out of range class must be split into two subclasses to characterize the WOM behavior.

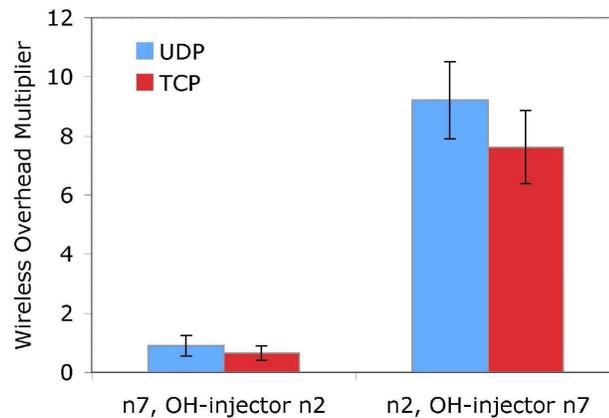


Figure 4.28 Asymmetry of WOM of two nodes with respect to one another.

Discussion

We now post-process Fig. 4.25 considering a positive SNR difference (or capture win) or negative SNR difference (or capture lose) by the data-sender s . Fig. 4.29 illustrates the net effect of capture and depicts the WOM values for the two cases as to whether the data transmitter s wins or loses the capture. The figure indicates that despite node o being in the same out-of-range class, the WOM value can be as small as 1 (capture win) or as large as 25 (capture lose). The observed physical layer capture effect on WOM explains the asymmetry shown in both Fig. 4.26 and Fig. 4.28.

In summary, the primary factor that controls the aforementioned WOM classes/subclasses is the non-forwarding links, i.e., the level of coordination the data-sender has with the overhead-injector. If the two transmitters are out of range, the secondary factor

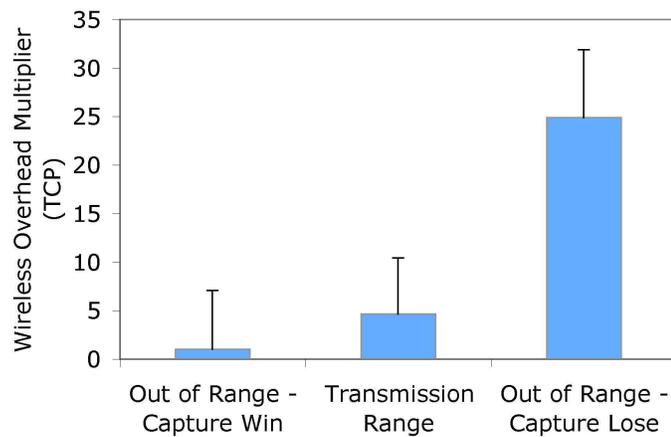


Figure 4.29 WOM considering the different effects of physical layer capture effect within the out of range link class.

is the relative quality of the forwarding links, i.e., the relative SNR at the mutual receiver which drives the capture effect. From these two factors, clear WOM modes can be established. Finally, the behavior within the modes is driven by a tertiary effect, the protocol set, consisting of the traffic type (e.g., TCP or UDP) and protocol parameters (e.g., the usage of RTS/CTS).

For example, we can reconsider the RTS results presented in Fig. 4.26 according to the discussion above: The RTS/CTS mechanism (tertiary effect) is unable to completely reduce WOM to the values associated with the transmission range scenario since the RTS messages are also captured (secondary effect) at the mutual receiver, thereby reducing the ability of the collision avoidance mechanism to counter the hidden terminal problem (primary effect).

Finally, although we cannot show the carrier sense WOM behavior within the

TFA hardware and environment, the expected values are similar to the transmission range class as verified by ns-2 simulation.

4.5.4 Predicting the Throughput Reduction Factor for all Traffic Types and Rates

We now apply our model and experimentation to a second domain in which hidden nodes transmit small-sized, low-rate control overhead, causing a disproportionately large effect on large-sized, high-rate data flows. While in the previous section, we showed the existence of such an effect in a mesh network [21], the reasons for the losses have not clearly been identified nor have the implications been explored for other packet types and modulation rates often used throughout a network. Since this effect can exist for all traffic types, we term it the *throughput reduction factor*. The throughput reduction factor is the achieved throughput without the presence of interfering traffic minus the achieved throughput with interfering traffic over the injected traffic rate of the interfering sources.

We begin by understanding the throughput reduction factor within TFA for different modulation rates for the embedded flow. We perform the experiments at off-peak times and generate an interfering traffic profile presented in the prior section of 100 byte packets at a rate of 10 kbps from a hidden interferer. We measure the throughput reduction factor for embedded flows for each of the modulation rates for flows across the TFA network.

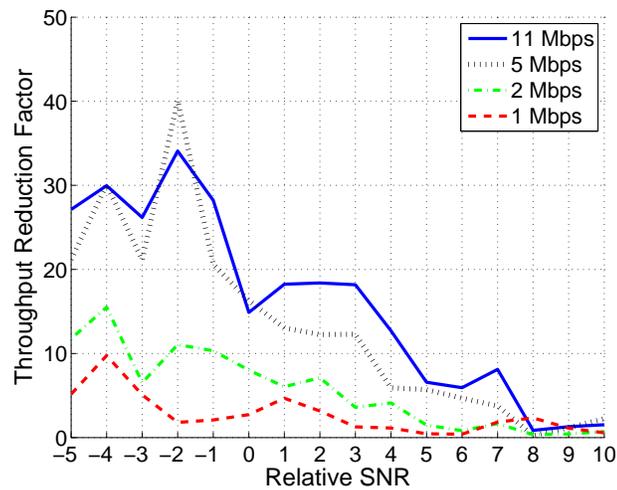


Figure 4.30 Range of throughput reduction multipliers based upon the data rate of the embedded link.

Fig. 4.30 depicts the throughput reduction factor for each relative SNR from -5 to +10 dB over each modulation rate. In the left of the graph, the data flow loses capture to the interfering source, experiencing the highest throughput reduction factor. In the right of the graph, the data flow has minimal reduction due to higher capture probability at the receiver. Clearly, the highest modulation rates have higher penalties (in terms of air-time) as the interfering source collides and forces the embedded flow to backoff. However, the throughput reduction factor is also affected by the ability of the modulation rate to capture. For example, 1 Mbps reaches a near-zero value at a relative SNR of 5 dB versus 8 dB for 11 Mbps). *Finding:* For a given relative SNR in relation to the interfering traffic, the modulation rate choice of the embedded flow can determine a throughput reduction factor of nearly 20 to a value close to zero.

We now use our measurements and model to reveal the key system properties

that yield the throughput reduction factor. In our capture experiments, neither modulation rate nor packet size of the interfering source affected the performance of the data flow, i.e., if the data flow was stronger, the capture threshold did not depend on the traffic characteristics of the interfering traffic (given that it was overlapping). Thus, for the throughput reduction factor, neither packet size nor modulation rate of the interfering flow determines the capture behavior of overlapping data packets.* Hence, the throughput reduction factor is primarily driven by the traffic characteristics of the data flow since choices of modulation rate and packet size affect the ability to capture. Therefore, we now use our model to predict the throughput reduction factor for different packet sizes and modulation rates of the data flow with the same interfering traffic profile as before.

Fig. 4.31 depicts the throughput reduction factor for the data flow based on its packet size and modulation rate over different relative SNR values versus the interfering source. The highest values of the throughput reduction factor are for large-sized, high-rate packets and lowest for small-sized, low-rate packets. While this is somewhat expected due to the inability of small-sized packets to achieve high throughput, the crossover point of when the factor approaches zero is interesting. For example, consider a modulation rate of 2 Mbps. For the packet size of 100 bytes, the throughput reduction factor goes to a near-zero value at +2 dB versus +7 dB for

*The interferer's offered load would affect the throughput reduction due to the increased probability of overlapping packets as the load increased.

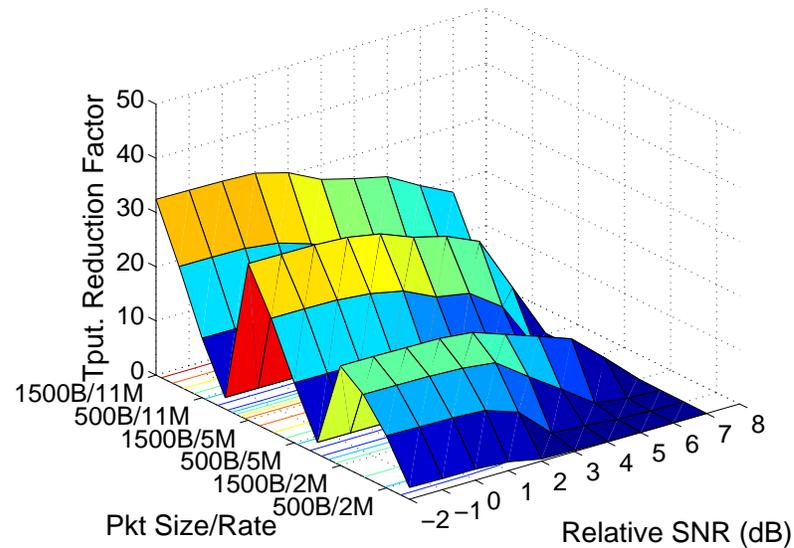


Figure 4.31 Packet size and physical layer rate effect on throughput reduction factor. 1500-byte packets at the same rate. *Finding:* Our model and analysis shows that the largest throughput reductions are due to the joint factors of control traffic originating from a hidden terminal and the data traffic’s inability to win capture over the control traffic. Note that since the control source has an uncoupled backoff with broadcast traffic, the control packet’s ability to win capture does not affect the throughput of the data flow as with the interfering source with coupled backoff behavior.

4.6 Summary

In summary, we perform extensive measurements on embedded links within an urban mesh network and analytically model the complex factors that exist within these scenarios such as topology, modulation rate, packet size, channel conditions, and physical layer capture. Our experimental analysis and model reveal that only 1 dB

of channel fluctuation is able to cause topological profile inversion where, though the connectivity remains, the performance mimics a completely different sub-topology. Using our model, we explore the interdependencies of these complex factors and find that reverse capture plays a critical role in the topological profile inversion. Finally, we apply our model and experimentation to two different problem domains: modulation rate selection and the interaction of control and data traffic. When these aforementioned complex interdependencies are understood and applied, we show that embedded links operate in improved performance regimes.

Chapter 5

Related Work

5.1 Modulation Rate Adaptation Rate Adaptation Protocol Design

There are two classes of rate adaptation mechanisms that have been developed which differ in how they determine the appropriate physical layer rate according to the perceived channel state. These schemes can be classified into loss-triggered multirate protocols [12, 13, 14, 15, 16, 17] and SNR-triggered multirate protocols [18, 19]. *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 [14] or delivery statistics over a window of time based upon historical performance of the modulation rates [12, 13, 17]. *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 [18] or sends a burst of data packets in proportion to the modulation rate over the base rate for time-share fairness [19]. Reference [34] implements the protocol proposed in [18] with a software-defined MAC and PHY resulting in a turn-around 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 μs turn-around 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.

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 [35, 18], collision/fading differentiation with hidden terminals indoors with off-the-shelf hardware [17], equal air-time assurance for SNR-triggered rate adaptation simulations where the coherence time is assumed to be much greater than the packet period [19], or single-active, one-hop flow performance of loss-triggered protocols compared on a mesh deployment [13]. 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 of the cross-layer implementation, 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).

Physical Layer 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 [13], non-vehicular mobility in a simulator [18, 19], or indoor environment [28, 17]. 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 city-wide wireless networks and other large-scale mesh networks such as [2], bring to 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: (i) a weaker (i.e., more distant) transmitter consumes a majority of the air-time and causes the stronger transmitter to have reduced rate [22], and (ii) 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 [17]. 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 [21]. Therefore, we evaluate how accurate the rate adaptation mechanisms are with heterogeneous forwarding links within a hidden terminal scenario.

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 [4], the effect is common within deployed networks [21] 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 3.4.

Residential Urban and Downtown

Reference [23] 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 [2]. We show that while the effect of multipath is severe in the downtown scenario, it is far less severe in the residential urban scenario. Further, we find that the assumption of coherence time being much greater than the packet period made in [19] does not hold even in static topologies within downtown scenarios

due to effects such as the mobility of scatterers.

5.2 Embedded Link Performance

Analytical Models of 802.11 and CSMA

There is a rich body of work on modeling CSMA, dating back to the seminal work by Kleinrock and Tobagi [36]. Other models include a perfect capture assumption which was based upon the timing of the packet as opposed to the channel condition [8, 37]. With the introduction of 802.11, Bianchi presented a simplified model that used the assumptions of single rate, single clique, and fully-backlogged traffic with fixed packet size [26]. Reference [38] considered physical layer features such as hidden terminals and capture without topological asymmetries. More general topologies and scenarios were later explored with an idealized channel and interference model [27, 39]. Recent measurement-based models use $O(n)$ measurements of n nodes to predict throughput for use in applications such as online network management [40, 41, 42].

Measurement Studies of Multi-Hop 802.11 Networks

A number of works have identified the channel conditions and timing under which physical layer capture occurs for pairs of nodes [4, 5, 6]. In such a scenario, others have proposed modifying the physical layer properties to address the lack of fairness that results [43, 44]. Other work has performed measurements on indoor multi-hop wireless networks and topologies to characterize interactions of flows [45]. Finally, a number of measurement studies have been performed in mesh networks to explore the

link behavior [23], flow performance [46], rate adaptation [13], routing metrics [47], and overhead effects [21].

In contrast, our work is the first to show via modeling and experimentation that small-scale channel fluctuation can yield bi-modal performance shifts. Namely, embedded link interactions affected by topology, channel conditions, modulation rate, packet size, and physical layer capture can yield topological profile inversions that emulate changes in connectivity.

Chapter 6

Conclusion

In conclusion, embedded links have complex and time-varying interacting factors including topology, channel conditions, modulation rate, packet size, and physical layer capture. To understand these complex factors that contribute to embedded link performance, I perform 100's of thousands of urban measurements and develop an embedded link model that is able to predict throughput. First, I design a custom, cross-layer framework to evaluate multiple and previously un-implemented modulation rate adaptation mechanisms and show that no existing protocol is able to successfully track a simple urban scenario. By performing experiments in both in-lab, controlled environments and urban vehicular and non-mobile environments, I identify the reasons for the protocols' inaccuracies and show that large gains can be achieved by the joint consideration of coherence time and SNR. Second, embedded link measurements in the TFA Network reveal that a topological profile inversion exists where though only 1 dB of channel fluctuation occurs, there is a bi-modal throughput shift that mimics a change in node connectivity. I show that the inversion depends on the physical layer capture of reverse traffic, an effect which has not been studied. By applying understanding from both our model and experimentation, a new aspect of modulation rate selection is revealed dealing with capture, topology, and packet size.

Further, I experimentally show that the low-rate control traffic can have a disproportionately large effect on the high-rate data throughput of embedded links and reveal the reasons for such an effect via the model.

There are immediate implications from my thesis. First, since the control traffic can have a multiplicative effect on the data, the tradeoff of injected traffic versus the gains of network management protocols must be reconsidered. Ideally, zero- or near-zero-overhead mechanisms would be used for such functions as rate limiting, routing, fault-detection, and link establishment. Second, since coherence time and channel conditions leads to accurate modulation rate selection, there is a need for knowledge of the coherence time of the channel at the MAC layer. There are a couple of different solutions to this problem. A more intelligent physical layer design can have on-going updates on a per sub-carrier basis using existing traffic. However, even without physical layer re-design a solution could exist in the form of using context information such as cellular signals (and resulting location information) combined with motion changes to infer the coherence time. Another implication is that in order to maximize the throughput with modulation rate selection, capture, topology, and channel conditions must be considered when contending with other devices, presenting a challenging problem to solve for a particular embedded link. Finally, the work performed for this thesis has already contributed to improved performance of the TFA Network. As outlined here, the understanding of embedded links has future

implications on performance improvements on both the TFA Network and wireless networks of all types.

References

1. E. Guizzo, “Networking from the rooftop,” in *Technology Review*, Aug. 2003.
2. J. Camp, J. Robinson, C. Steger, and E. Knightly, “Measurement driven deployment of a two-tier urban mesh access network,” in *ACM MobiSys*, Uppsala, Sweden, June 2006.
3. IEEE, “Wireless LAN medium access control and physical layer specification,” ANSI/IEEE Standard 802.11, 1999.
4. A. Kochut, A. Vasani, A. U. Shankar, and A. Agrawala, “Sniffing out the correct physical layer capture model in 802.11b,” in *IEEE ICNP*, Berlin, Germany, Oct. 2004.
5. J. Lee, W. Kim, S.-J. Lee, D. Jo, J. Ryu, T. Kwon, and Y. Choi, “An experimental study on the capture effect in 802.11a networks,” in *ACM WiNTECH*, Sept. 2007.
6. G. Judd and P. Steenkiste, “Understanding link-level 802.11 behavior: Replacing convention with measurement,” in *IEEE WICON*, Oct. 2007.
7. F. A. Tobagi and L. Kleinrock, “Packet switching in radio channels: Part II - the hidden terminal problem in carrier sense multiple access and the busy tone solution,” *IEEE Transactions on Communications*, vol. 23, no. 12, pp. 1417–1433, 1975.
8. F. Tobagi, “Modeling and performance analysis of multihop packet radio networks,” *Proceedings of the IEEE*, vol. 75, no. 1, pp. 135–155, Jan. 1987.
9. J. Camp, “Measurement driven deployment of a two-tier urban mesh access network,” M.S. Thesis, Rice University, January 2006.
10. P. Murphy, A. Sabharwal, and B. Aazhang, “Design of WARP: Wireless open-access research platform,” in *European Signal Processing Conference*, June 2006.
11. 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.
12. “MadWifi Project,” <http://madwifi.org>, September 2005.

13. J. C. Bicket, "Bit-rate selection in wireless networks," M.S. Thesis, MIT, February 2005.
14. A. Kamerman and L. Monteban, "WaveLAN II: A high-performance wireless LAN for the unlicensed band," *Bell Labs Technical Journal*, pp. 118–133, Summer 1997.
15. J. Kim, S. Kim, S. Choi, and D. Qiao, "CARA: Collision-aware rate adaptation for IEEE 802.11 WLANs," in *IEEE INFOCOM*, 2006.
16. M. Lacage, M. Hossein, and T. Turletti, "IEEE 802.11 rate adaptation: A practical approach," in *MSWiM*, October 2004.
17. S. Wong, S. Lu, H. Yang, and V. Bharghavan, "Robust rate adaptation for 802.11 wireless networks," in *ACM MobiCom*, 2006.
18. G. Holland, N. Vaidya, and P. Bahl, "A rate-adaptive MAC protocol for multi-hop wireless networks," in *ACM MobiCom*, Rome, Italy, July 2001.
19. B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic media access for multirate ad hoc networks," in *ACM MobiCom*, Atlanta, GA, Sept. 2002.
20. "Technical report on RF channel characterization and system deployment modeling," JTC (Air) Standards Contribution, Tech. Rep. JTC(AIR)/94.09.23- 065R6, September 1994.
21. J. Camp, V. Mancuso, O. Gurewitz, and E. Knightly, "A measurement study of multiplicative overhead effects in wireless networks," in *IEEE INFOCOM*, Apr. 2008.
22. M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *IEEE INFOCOM*, San Francisco, CA, Apr. 2003.
23. D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11 mesh network," in *ACM SIGCOMM*, Portland, OR, 2004.
24. V. Bharghavan, S. Demers, S. Shenker, and L. Zhang, "MACAW: A media access protocol for wireless LANs," in *Proceedings of ACM SIGCOMM '94*, London, UK, Aug. 1994.
25. R. Sinha, C. Papadopoulos, and J. Heidemann, "Internet packet size distributions: Some observations," University of Southern California, Los Angeles, Tech. Rep. ISI-TR-2007-643, May 2007.

26. G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE JSAC*, vol. 18, no. 3, pp. 535–547, Mar. 2000.
27. M. Garetto, J. Shi, and E. Knightly, "Modeling media access in embedded two-flow topologies of multi-hop wireless networks," in *ACM MobiCom*, Aug. 2005.
28. G. Judd, X. Wang, and P. Steenkiste, "Efficient channel-aware rate adaptation in dynamic environments," in *ACM MobiSys*, June 2008.
29. IEEE, "Wireless LAN medium access control and physical layer specification," ANSI/IEEE Standard 802.11, 2003.
30. R. Jain, W. Hawe, and D. Chiu, "A quantitative measure of fairness and discrimination for resource allocation in shared computer systems," Ohio State University, Columbus, Tech. Rep. DEC-TR-301, September 1984.
31. T. Rappaport, *Wireless Communications, Principles & Practice*, ser. Emerging Technologies Series, T. Rappaport, Ed. Upper Saddle River, New Jersey: Prentice Hall, 1996.
32. J. Shi, O. Gurewitz, J. Camp, and E. Knightly, "Measurement and modeling of the origins of starvation in congestion-controlled mesh networks," in *IEEE INFOCOM*, Apr. 2008.
33. J. Camp, V. Mancuso, O. Gurewitz, and E. Knightly, "Effects of heterogeneity in backhaul connectivity of mesh networks: A measurement study," Rice University, Electrical and Computer Engineering, Houston, TX, Tech. Rep. TREE0801, Jan. 2008.
34. 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 *IEEE Vehicular Technology Conference*, Dublin, Ireland, Apr. 2007.
35. P. Chevillat, J. Jelitto, A. N. Barreto, and H. L. Truong, "A dynamic link adaptation algorithm for IEEE 802.11a wireless LANs," in *IEEE International Conference on Communications*, Anchorage, Alaska, May 2003.
36. L. Kleinrock and F. A. Tobagi, "Packet switching in radio channels: Part I - carrier sense multiple-access modes and their throughput-delay characteristics," *IEEE Transactions on Communications*, vol. 23, no. 12, pp. 1400–1416, 1975.
37. R. R. Boorstyn, A. Kershenbaum, B. Maglaris, and V. Sahin, "Throughput analysis in multihop csma packet radio networks," *IEEE Transactions on Communications*, vol. 35, no. 3, pp. 267–274, mar 1987.

38. M. Carvalho and J. J. Garcia-Luna-Aceves, "A scalable model for channel access protocols in multihop ad hoc networks," in *Proceedings of ACM MobiCom '04*, Philadelphia, PA, Sept. 2004.
39. M. Garetto, T. Salonidis, and E. Knightly, "Modeling per-flow throughput and capturing starvation in csma multi-hop wireless networks," *IEEE/ACM Transactions on Networking*, vol. 16, no. 10, pp. 864–877, Aug. 2008.
40. C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Measurement-based models of delivery and interference in static wireless networks," in *ACM SIGCOMM*, Pisa, Italy, Sept. 2006.
41. L. Qiu, Y. Zhang, F. Wang, M.-K. Han, and R. Mahajan, "A general model of wireless interference," in *ACM MobiCom*, Sept. 2007.
42. Y. Li, L. Qiu, Y. Zhang, R. Mahajan, and E. Rozner, "Predictable performance optimization for wireless networks," in *ACM SIGCOMM*, Aug. 2008.
43. S. Ganu, K. Ramachandran, M. Gruteser, I. Seskar, and J. Deng, "Methods for restoring MAC layer fairness in IEEE 802.11 networks with physical layer capture," in *ACM REALMAN*, May 2006.
44. T.-Y. Lin and J. Hou, "Interplay of spatial reuse and SINR-determined data rates in CSMA/CA-based, multi-hop, multi-rate wireless networks," in *IEEE INFOCOM*, May 2007.
45. J. Lee, S.-J. Lee, W. Kim, D. Jo, T. Kwon, and Y. Choi, "Understanding interference and carrier sensing in wireless mesh networks," *IEEE Communications Magazine*, to appear 2008.
46. J. Bicket, S. Biswas, D. Aguayo, and R. Morris, "Architecture and evaluation of the MIT Roofnet mesh network," in *Proceedings of ACM MobiCom*, Cologne, Germany, August 2005.
47. D. De Couto, D. Aguayo, J. Bicket, and R. Morris, "A high-throughput path metric for multi-hop wireless routing," in *ACM MobiCom*, September 2003.