AGILE AND ROBUST RESOURCE MANAGEMENT IN CSMA WIRELESS NETWORKS

BY KISHORE RAMACHANDRAN

A dissertation submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Doctor of Philosophy

Graduate Program in Electrical and Computer Engineering

Written under the direction of

Marco Gruteser

and approved by

_____

_____

_____

_____

New Brunswick, New Jersey

January, 2009

ABSTRACT OF THE DISSERTATION

Agile and Robust Resource Management in CSMA Wireless Networks

by Kishore Ramachandran

Dissertation Director: Marco Gruteser

With the recent push towards wireless broadband, and user migration towards mobile devices, it has become imperative for wireless networks to support higher network capacity, increased battery life, and greater mobility. While a combination of both local- and wide-area wireless technologies will be needed to meet these requirements, we believe that *carrier-sense multiple-access (CSMA)* wireless networks, if designed well, can play an important role in the future. Traditional CSMA wireless networks, which form the basis for today's ubiquitous wireless LAN technology—IEEE 802.11 or WiFi, cannot satisfy the stated requirements mainly because they lack efficient resource management (or resource parameter adaptation) techniques. Several fundamental characteristics of these networks, and practical implementation challenges, also limit the applicability of solutions from other domains. Taken together, these considerations force us to fundamentally re-think the design of resource parameter adaptation for CSMA wireless networks.

In this dissertation, we first identify that, to satisfy the requirements, resource management techniques for CSMA wireless networks should possess two fundamental, but conflicting properties, *agility*, and *robustness*. Briefly, to provide increased bandwidth to mobile users, *agile* solutions are required that exploit opportunities for improved performance; at the same time, solutions cannot afford to compromise on *link robustness*. In addition, we realize that striking this trade-off effectively calls for different solutions in indoor and outdoor environments. To

prove these hypotheses, we present the design and implementation of two resource management frameworks, *Symphony* and *Sonata*, for indoor and outdoor environments respectively. Indoors, *Symphony* increases network capacity and battery life for mobile clients by addressing the classical problem of joint, per-link, transmit power control and rate adaptation. For improved robustness, *Symphony* uses novel mechanisms based on measuring the expected transmission time (ETT), and the *utility* of RTS/CTS frames, while relying on a *learning* approach to converge quickly to the right resource parameter choice. Outdoors, the Sonata framework introduces a novel and fundamental tradeoff between directionality and base station diversity for uplink transmissions. Using a new location-based approach for improved parameter convergence, *Sonata* is able strike the agility-robustness tradeoff effectively.

Together, these frameworks prove that, achieving the right balance between agility and robustness can enable CSMA wireless networks to transition to the wireless broadband era.

Acknowledgements

Looking back, graduate school, and the process of earning my Ph.D., have fundamentally changed the way I think. I strongly believe that this positive change has been brought about by the people I have interacted with, during this period. More than most other people, I have been influenced by my advisor, Marco Gruteser. I would like to thank him for spending a significant amount of his time, especially during my first few years, in helping me become a better researcher. I will always remember our incredibly long meetings on Friday evenings, and how we did not go home till we achieved our weekly goals. I will also not forget our mantra of not giving up till about 15 minutes *after* the paper deadline! His commitment and work ethic have been, and will be, an important source of inspiration for me.

Over the last two years, I have also been extremely fortunate to work with Ravi Kokku, who is not only a wonderful human being, but also a great mentor. I am thankful to Ravi for believing in me so strongly, for encouraging me to think differently, for providing me with a living example of how to effectively communicate with people, and finally, for making research fun again. In short, thanks a lot Ravi, for going above and beyond your call of duty as a mentor.

Just like it takes a village to raise a child, in my case, it took a lab to mold a graduate student. I was lucky to interact and work with Dipankar Raychaudhuri, Ivan Seskar, Predrag Spasojevic, and Wade Trappe — all great researchers in their own right. I am deeply indebted to Dr. Raychaudhuri for giving me a chance to pursue a career in research, especially at a time when most other avenues were closed for me. For involving me in the ORBIT project, for ensuring that I was never short of funding, for taking an active interest in my professional career, and for never letting me lose sight of the big picture, I am grateful to him.

Ivan is the most skilled, and knowledgeable systems designer I have ever interacted with. I am thankful to him for being a patient manager during my initial years, and for helping me become more resilient to criticism. I would also like to thank Predrag for serving on my committee, and for providing insightful comments on my work at various stages.

While Dr. Trappe was not directly involved in my research, his constant interest and encouragement were strong factors in keeping me motivated, even during the bad times. Thanks a lot Dr. Trappe for always treating me with a lot of respect, and for keeping the door open for me. Its a pity we could not work together more, but I hope to change that in the future.

My collaborators, Sachin Ganu, Mesut Ergin, Sanjit Kaul, Sampath Rangarajan, Karthik Sundaresan, and Haris Kremo, have contributed in different capacities to my success. Also, many thanks to all WINLAB students for the fruitful discussions at student seminars.

On the personal front, the years through graduate school would not have been so much fun if it were not for friends like Pravin Shankar, Ramya Rao, Prabhu Shankar, Bhuvnesh Parikh, Suhas Mathur, Joydeep Acharya, Chandru Raman, and Pandurang Kamat.

The most important people in my life—my family—have always been there for me. More than anything else, I am glad that these personal relationships are still strong even after phases where I could hardly stay in touch. Arun, although I lost the race to be a doctor before you, I am so glad that at least one of us can fix a cold! Radhika, and Mony, thank you for your love and support, especially during the toughest phase of my life so far (and yes, you need not pay for dinner anymore). More importantly, thanks a lot for Sandhya and Gitu, the ultimate stress relievers. To Vidya, the love of my life, thank you for being so patient and understanding especially during paper deadlines, and the last few months when the future was so uncertain. Finally, I would like to thank my parents for giving me everything that I could ask for, even during the hardest of times. Amma and Appa, this one's for you.

Dedication

To My Parents

Mrs. Karpagam Ramachandran & Mr. S. Ramachandran

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

This dissertation describes the design and implementation of *Symphony* and *Sonata* — two distributed frameworks that enable efficient channel utilization, in indoor and outdoor environments, for the general class of *carrier-sense multiple access (CSMA)* wireless networks.

1.1 Technology trends: Wireless Broadband on Mobile Devices

Over the last century, the use of electronic devices as a means of *communication* has elevated the quote, *the world is a small place*, to a fact. While copper wires were the primary means to connect these devices for a substantial period of this time, the ability to communicate without wires was possible as early as 1895 [3]. However, it took another 70 years before core concepts, collectively referred to as cellular telephony, were developed that would make mass-market wireless communications a reality. In the four decades since then, the success of cellular wireless networks, engineered to carry voice traffic, has been unparalleled — currently, there are over 2 billion cellular telephone (or cell phone) subscribers worldwide [4, 5]. Another simultaneous development in the wired network domain has been the phenomenal growth of the Internet with over a billion hosts today [6, 7]. It is but natural for users to *demand for* wireless Internet connectivity.

Users are increasingly viewing their hand-held wireless devices as mobile PCs that can host a wide range of applications in addition to voice. This is evident from today's cell phone units that routinely come bundled with digital cameras and mp3 players [8, 9], in addition to software that includes web browsers, media players, and email clients. This is also reflected in the *significant increase in wireless data traffic* on existing cellular networks [10] — between 2006-2007, wireless data traffic quadrupled for a major US cellular network service provider, and is expected to grow exponentially in the near future (see Figure 1.1).

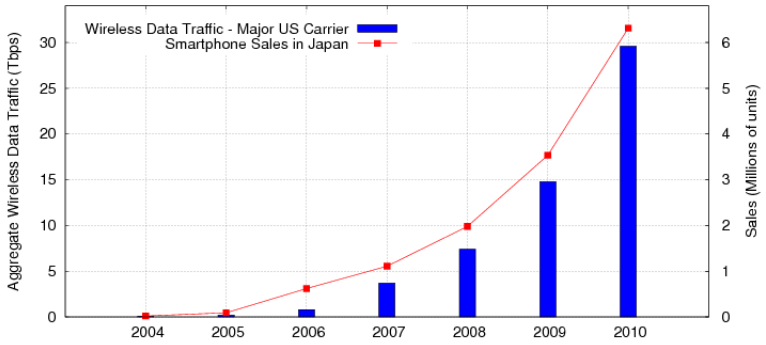

Figure 1.1: Market Forecast: Exponential increase in demand for network capacity and mobile device sales [2].

Additionally, computing firms that have traditionally developed hardware platforms, operating systems, and applications for desktop PCs (e.g. Apple, Google and Microsoft) are now shifting their focus to wireless hand-held devices [11, 12]. With the release and widespread adoption of platforms such as Apple's iPhone [13] and Google's gPhone [14], large-scale user migration to mobile devices in the next decade is a foregone conclusion. In light of these trends, wireless networks will have to support a set of basic requirements, which we discuss next.

1.1.1 Emerging Requirements

To support wireless broadband on mobile devices, wireless networks will have to provide:

1. **Support for increased network capacity:** Unlike the dominant usage of wireless devices for voice calls in the past, future wireless devices are expected to run a variety of applications ranging from low bandwidth web browsing to high bandwidth video, in addition to delay sensitive voice. Moreover, networks will be required to provide service guarantees to users, while also having to constantly monitor the minimum level of service that can be provided on the time-varying wireless medium. Finally, just like their wired counterparts, we expect wireless ISPs to also employ the deployment guideline of capacity over-provisioning to deal with unanticipated increases in network traffic. Taken together, increased network capacity emerges as a fundamental requirement.

2. **Support for increased battery lifetimes:** With the rapid migration of users to mobile platforms, protocols and algorithms have to be designed with energy consumption in mind. In fact, this is one requirement that directly affects user adoption.

3. **Support for high mobility:** With the increasing use of computing devices by mobile users, another intrinsic requirement is support for high-speed mobility. Unlike today's wide-area cellular networks that use *low-bandwidth macro-cells* to support highly mobile or vehicular users (so as to reduce the number of handoffs), future networks will have to support high bandwidths even at vehicular speeds.

4. **Support for high user densities:** Given the high rate of penetration of cell phones worldwide, and the expectation that today's cell phones will be replaced by tomorrow's smart phones in a user-transparent manner, future wireless networks will have to be designed with high user densities in mind.

1.2 The Case for CSMA Wireless Networks

To satisfy these requirements, a number of alternatives have been proposed over the last decade including local-area IEEE 802.11 (WiFi) [15], metropolitan-area IEEE 802.16 (WiMax) [16], and wide-area GSM HSPA [17], CDMA 1xEVDO [18], and 3G LTE [19]. While these technologies are in various stages of deployment — WiFi has already been deployed in over 200 million homes worldwide [20]; WiMax, GSM HSPA, and CDMA 1xEVDO are in the process of being deployed; 3G LTE deployments are not expected to begin before 2010 — we believe that a combination of technologies will be necessary to satisfy the design requirements. This is primarily because it is hard for any one of these technologies to satisfy all the requirements alone. Since capacity and coverage are at odds with each other, especially for wireless networks, capacities on the order of 100s of Mbps will only realistically be supported by local-area networks (that use small cells) such as WiFi. However, it has always been easier for wide-area networks to support vehicular mobility, through the use of larger cells (and reduced number of cell transitions or handoffs). We further believe that *carrier-sense, multiple-access* (CSMA) wireless networks, which form the basis for today's widely deployed WiFi networks, will be an integral part of the wireless broadband future. This belief is backed by

1. the widespread deployment of WiFi networks in homes and offices,

2. the inclusion of WiFi interfaces in a majority of smart phones, dual-mode cell phones [21, 22], and portable media players [23, 24].

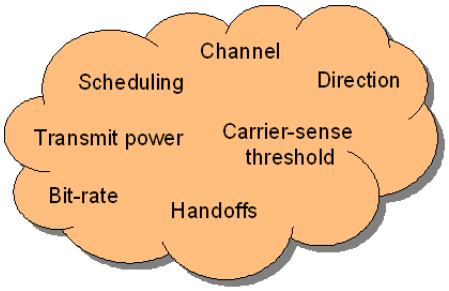

Figure 1.2: CSMA Wireless Networks: Degrees of Freedom.

3. the standardization of interactions between cellular and WiFi networks via Unlicensed Mobile Access (UMA) [25], and

4. the adoption of the IEEE 802.11p draft standard [26–28], which is heavily influenced by the existing IEEE 802.11 and IEEE 802.11a standards, by leading car manufacturers for increasing driver safety and for in-vehicle infotainment applications.

Since WiFi is a local area networking technology with limited coverage, the two usage models that are envisioned include (a) the opportunistic use of WiFi networks for last-mile network access, and (b) their use in an *ad-hoc* manner to serve the growing base of peer-to-peer (p2p) applications (e.g. p2p gaming [29]). Of these, we focus on the former model given that its more closely aligned with our vision for ubiquitous, wireless broadband connectivity. In general, WiFi's relatively certain position in future wireless broadband networks serves as our primary motivation to study the more general class of CSMA wireless networks.

1.3 The Research Problem

In light of the design requirements (in Section 1.1.1), efficient management of available resources, both from a individual user's perspective (e.g. battery life), and from a network-wide perspective (e.g. network capacity), becomes critical. However, today's CSMA wireless networks lack efficient resource management or resource parameter adaptation techniques.

More specifically, resource management in wireless networks is affected by *tuning* multiple resource parameters that are designed for complementary purposes. Fig. 1.2 shows the key resource parameters for CSMA wireless networks. Briefly, these include: *schedule* that corresponds to transmission isolation in time, *transmit power* and *direction* that correspond

to transmission isolation in space, *channel* that corresponds to transmission isolation in frequency, *bit-rate* that corresponds to the amount of forward error correction (modulation and coding) added to each frame, *carrier-sense (CS) threshold* that corresponds to controlling the space over which transmitters defer to each other, and *handoff* that corresponds to the cell transition policy adopted. Note that CSMA wireless networks employ CSMA with collision avoidance (CSMA/CA) scheduling, and for backwards compatibility purposes, this degree of freedom cannot change significantly. As for the rest of the resource parameters, the main criticism is that they are either static or coarsely dynamic in today's CSMA wireless networks. For example, a majority of deployments rely on manual channel assignment that is routinely suboptimal [30], and commercial vendors have only recently started using dynamic channel assignment [31]. Similarly, transmit power control and CS threshold adaptation are performed on a *per-cell* basis rather than on a *per-link* basis [32, 33]. In addition, even existing dynamic parameter adaptation techniques (such as bit-rate adaptation and handoffs), as we will show later, cannot cope well with high node density and/or mobility. The main reason for these inefficiencies in resource parameter adaptation stem from the different set of design guidelines used by today's CSMA wireless networks. Existing incarnations of these networks were designed for low mobility scenarios and expected CSMA/CA scheduling to prevent interference at high node densities. As we will show later, CSMA/CA scheduling is able to only partially mitigate interference, even in low node density scenarios. The use of license-free spectrum further complicates matters by introducing the possibility of CSMA-incompatible interference from cordless phones and microwave ovens. All these factors motivate us to fundamentally re-think the design of resource parameter adaptation for CSMA wireless networks.

In this dissertation, we focus on the following open questions:

1. What fundamental properties should resource parameter adaptation techniques in CSMA wireless networks possess, in order to satisfy the new requirements?

2. What are the primary challenges in incorporating these properties? Can these challenges be overcome in a backwards-compatible manner, and if so, how?

1.3.1 The Agility-Robustness Tradeoff

To answer these questions, we focus on three resource parameters, bit-rate, transmit power and directionality, in depth. We select these parameters because we believe that they are fundamentally required to satisfy all the requirements. We also believe that these parameters have the potential to deliver the maximum performance gain. We use these parameters to show that, in general, resource parameter adaptation has to possess two fundamental, but conflicting, properties: *agility*, and *robustness*.

On the time-varying wireless medium, to avoid performance degradation due to a variety of reasons, such as interference, mobility, and fading, and to exploit opportunities for performance enhancement, solutions need to be *agile*. They need to be able to rapidly adapt to both extremes, while converging quickly to the right parameter choice. At the same time, solutions also need to keep *link robustness* and *network stability* in mind — the agility or responsiveness of solutions to channel dynamics should not (a) come at the cost of link reliability, and (b) drive the network into an unstable or suboptimal state. For this purpose, adaptation techniques need to accurately identify the exact cause for performance degradation or enhancement, so as to react appropriately. Thus, the main challenge for resource parameter adaptation is to strike this *agility-robustness tradeoff* effectively.

We now look at several unique fundamental characteristics and practical issues that make it challenging to incorporate these properties in the design of parameter adaptation techniques. These challenges include:

- *The use of unlicensed spectrum:* As today's CSMA wireless networks use license-free spectrum, the sources of co-channel interference can be varied (e.g. cordless phones and microwave ovens in the 2.4 GHz band). Even in the absence of such sources that are not protocol-compliant, the use of unlicensed spectrum implies that multiple CSMA wireless networks belonging to different domains (e.g. neighbor's WiFi network in homes or enterprises) can co-exist in the same location [30]. While these issues can be addressed if there were enough orthogonal channels or frequencies, this is understandably not the case given the scarcity of wireless spectrum.

  The IEEE 802.11 standard does provide orthogonal frequencies or channels but the number of such channels (three in IEEE 802.11b/g and twelve in IEEE 802.11a) is insufficient for intelligent channel assignment to prevent interference on its own. While the recent

addition of twelve channels in the 5 GHz band does provide reason to believe that channel assignment may be able to mitigate interference significantly, this is not the case because: (a) the use of twice the channel bandwidth (40 MHz) by the upcoming IEEE 802.11n standard with multiple-input, multiple-output (MIMO) enhancements, (b) the existence of military radar on some of these 5 GHz channels, and (c) the simultaneous use of some of these channels by other technologies such as WiMax.

Hence, resource parameter adaptation cannot depend of channel assignment alone to prevent interference, which can come from a variety of unpredictable sources.

- *The inability of CSMA/CA scheduling to prevent interference:* Having ruled out channel assignment as a means to prevent interference, we next look at CSMA/CA scheduling. As we will show later, CSMA/CA scheduling is unable to prevent the adverse effects of interference in (a) dense scenarios where all transmitters can hear each other, (b) hidden-terminal scenarios where some transmitters cannot sense each other's carrier signals, and (c) asymmetric channel access scenarios where some transmitters can sense others' carrier signals but not vice-versa. While this problem has prompted a number of recent studies [34–36] on measurement-based modelling of interference in static WiFi networks, extending these techniques to the general case with mobile users is an unsolved problem at present. It has also been shown that, even for the case of static networks with nodes using multiple cards and operating on multiple channels, the complexity of measurement is prohibitive even for moderate scales of network density [36].

  Thus, resource parameter adaptation cannot depend on CSMA/CA scheduling either, and techniques need to be *agile* and *robust* in the face of co-channel interference.

- *The lack of mechanisms for accurate interference measurement and feedback:* Even in the presence of hard-to-predict interference, mechanisms for accurate measurement and feedback, can go a long way in mitigating its impact. In fact, such signal-to-interference-and-noise-ratio (SINR) mechanisms are routinely used in cellular networks. Unfortunately, interference measurement mechanisms are inaccurate in today's CSMA wireless networks. This is because existing incarnations of these networks were initially designed for low levels of node density, in which, CSMA/CA scheduling was sufficient to prevent interfering transmissions from taking place. In addition, they were designed to be built with inexpensive hardware components. The end result was that little attention was

paid to the accurate estimation of interference at the physical layer. Further, existing incarnations lack critical feedback mechanisms between receivers and transmitters. These mechanisms are needed because they can help transmitters measure channel conditions, and interference at the receiver. They can also help solutions to rapidly converge to the right parameter settings in the presence of user mobility.

Thus, resource parameter adaptation in these networks cannot rely on existing SINR-based solutions to detect, and mitigate interference or to converge quickly in the presence of mobility. New techniques are required to address both issues that are ideally backwards-compatible.

- *The requirement for distributed resource parameter adaptation:* As opposed to centralized approaches used in cellular networks today, where the network infrastructure is responsible for resource parameter tuning on both clients and base stations, the network architecture of existing CSMA wireless networks demands more distributed solutions. Specifically, WiFi networks, from their instantiation, have lacked protocols for explicit co-ordination between cells (or basic service sets (BSS)). WiFi clients are also far more independent (relative to their cell phone counterparts) in terms of being able to make decisions on resource management. These factors have actually been responsible for their quick uptake among users and standards bodies are actually going to great lengths to preserve this model. Ensuring the optimality, and convergence of distributed parameter adaptation is, however, known to be a hard problem.

We further realize the need for different solutions in indoor and outdoor environments. This is because of the different characteristics of these two environments. Indoor environments are traditionally characterized by low mobility (walking speeds), multipath effects, and the potential for sustained co-channel interference. Outdoor environments, on the other hand, are mainly characterized by high-speed mobility.

1.3.2 Approach

Keeping these issues in mind, and since no single resource parameter is sufficient in isolation to satisfy all the requirements, we concentrate on joint adaptation of resource parameters. In particular, we study joint bit-rate (or rate) adaptation and transmit power control in indoor environments, mainly because both parameters are necessary to satisfy the requirements of

increased capacity and reduced energy consumption. Among the available resource parameter choices, while carrier-sense threshold adaptation [37,38] and dynamic channel assignment [39] can contribute to increased network capacity, they do not directly contribute to client battery life. Our experimental study in [40] also reveals that directionality is hard to realize with existing hardware in indoor, non-line-of-sight (NLOS) environments [40].

Outdoors, our work focuses on resource parameter tuning in the presence of high-speed mobility, and targets clients inside vehicles. Keeping in mind the rapid migration of users to mobile platforms [41, 42], and the recent emergence of uplink-intensive applications such as Web 2.0, image and video backups from cameras, video calls from phones, etc., we focus on uplink connectivity from vehicular users. While several techniques will be required to collectively meet the emerging requirements, we believe that three fundamental mechanisms—directionality, bit-rate, and base station diversity—will form an integral part of future mobile networks. In particular, while line-of-sight (LOS) conditions in outdoor environments favor directionality, we realize that directionality alone is insufficient to ensure robustness to mobility-induced channel fluctuations, and that receive diversity across multiple base stations (used for soft handoffs in cellular networks), along with robust bit-rate adaptation, is needed to overcome this problem. In the process of striking the agility-robustness tradeoff, we identify a novel and fundamental tradeoff between directionality and base station diversity for uplink transmissions. We also demonstrate how the physical location of a mobile device can be used to aid resource parameter convergence at vehicular speeds.

While adaptation of bit-rate, transmit power, and directionality have all received significant attention [1, 33, 43–91], our systematic and exhaustive analysis of related work (Chapter 2) reveals that no work is able to strike the *agility-robustness* tradeoff effectively. In what follows, we look at our specific contributions in more detail.

1.3.3 Contributions

To the best of our knowledge, ours is the first effort to systematically address the key problem of resource management via agile and robust resource parameter adaptation in CSMA wireless networks. In this regard, *Symphony* is the first framework to simultaneously address the joint optimization of bit-rate and transmit power for indoor CSMA wireless networks. Likewise, *Sonata* is the first framework to jointly balance the benefits from directionality, base station diversity, and bit-rate for highly-mobile, outdoor CSMA wireless networks. Both represent

efforts that, unlike a majority of past solutions, are not limited to theory and simulations. More specifically, this dissertation introduces the following innovations for indoor CSMA wireless networks:

- We design a synchronous, two-phase, bit-rate and transmit power adaptation algorithm for per-link resource management. Central to this algorithm's design is its *agility*, which allows coping with the time-varying wireless channel and user mobility. Our experiments demonstrate that this algorithm addresses the key challenges in a comprehensive and easily realizable manner.

- Through controlled, large-scale testbed experiments, we expose the inability of existing state-of-the-art bit-rate adaptation algorithms in coping with scenarios involving high node densities. In particular, we identify the fundamental limitation of these algorithms in differentiating between the diverse causes for poor performance. To remedy this situation, we propose a novel mechanism, called *Utility-RTS (URTS)*, which can be used to reliably detect high levels of MAC collision either due to high node density or due to hidden terminals. This mechanism does not require extra hardware or protocol changes and can be incrementally deployed on existing wireless cards. An added incentive of this mechanism is that it is robust to user mobility. Our experiments demonstrate that URTS supports dynamic detection of hidden-node interference even when it exists for small time durations ($>=$ 1 second).

- To detect and prevent channel access asymmetry, which is a fundamental and unique challenge to adaptive transmit power control in CSMA wireless networks, we design and evaluate an innovative, distributed mechanism using the *expected transmission time* of each frame. To our knowledge, we are also the first to provide an implementation for measuring channel access time on off-the-shelf 802.11 wireless cards without explicit hardware support. We expect this implementation to be useful for a variety of other algorithms at other layers of the protocol stack [92].

- For improved stability and convergence of the proposed dynamic rate and power control algorithms, we implement and evaluate *learning* mechanisms and *stochastic* control. This enables using past history to reward sound decisions and penalize poor ones. Our

experiments show that the resulting algorithms significantly reduce losses and retransmissions in the network and make a link more robust to user mobility.

For outdoor CSMA wireless networks, our contributions include:

- We introduce a novel and fundamental tradeoff between directionality and base station diversity, while identifying several parameters of importance—such as transmission bit rates available, link SNR, packet error rate, beam width, number of receivers covered for a given beamwidth, and number of antenna elements on the client—that guide the design of solutions.

- We design the *Sonata* framework that intelligently combines directionality, base station diversity and bit-rate adaptation to maximize the uplink throughput of a mobile client. As part of this framework, we design and implement *Robust Rate with Directionality and Diversity* or R2D2, a location-based adaptation algorithm that strikes the right tradeoff between directionality and diversity, and is robust to fluctuations in link quality. We also address the bit-rate adaptation issue at high-speed mobility, that is often left unexplored in past works [93].

- For improved convergence at vehicular speeds, we demonstrate the importance of keeping track of resource parameter adaptation choices at each location.

1.4 Outline

The rest of the dissertation is organized as follows. In Chapter 2, we provide the necessary background and qualitatively discuss related work. This is followed by a quantitative analysis of why parameter adaptation is challenging in Chapter 3. In Chapters 4 and 5, we describe the overall architecture of the Symphony and Sonata frameworks for indoor and outdoor CSMA wireless networks, respectively. Finally, Chapter 6 summarizes our contributions and presents future directions.

Chapter 2

Background and Related Work

Resource parameter adaptation has been extensively studied over the past two decades [1, 33, 43–91]. In this chapter, we conduct a comprehensive survey of the proposed approaches and show that, despite significant research, no solution provides a complete and easily realizable approach to address all the principal issues in CSMA networks. We begin our survey with a discussion on bit-rate adaptation.

2.1 Bit-Rate Adaptation

Bit-rate (or rate) adaptation enables IEEE 802.11 radios to cope with time-varying channel environments. The IEEE 802.11 standard mandates twelve bitrates between 1 and 54Mbps. Generally, higher bitrates correspond to higher nominal throughput but require higher signal-to-noise ratios (SNR) for correct demodulation. In an SNR-limited environment, higher bitrates will suffer from frame errors, limiting the effective goodput. In such an environment lower bitrates may provide higher effective goodput than high rates. Rate adaptation aims to dynamically adjust the transmission rate to maximize goodput depending on channel conditions. Since the standard does not specify any particular rate adaptation mechanism, manufacturers use different proprietary implementations. In the following subsections, we classify existing literature primarily based on how they estimate channel quality. Published rate adaptation mechanisms can be classified into frame error-based, throughput-based and SNR-based adaptation.

2.1.1 Frame-error-based adaptation

Auto Rate Fallback (ARF) [43], developed for WaveLAN-II 802.11 cards, and *Adaptive Auto Rate Fallback (AARF)* [44] use fixed and dynamic frame error thresholds to increase/decrease the bit-rate. *ONOE* [45], a frame-error based algorithm used in the MADWIFI driver for Atheros-based wireless NICs aims at selecting the highest bit-rate with less than 50% frame

loss rate. Periodically, for each destination station, the algorithm maintains a credit score that it increments if less than 10% of packets required a retransmission and no packets were dropped in the last time period. If the credit score surpasses a threshold (default 10), the bit-rate is raised. If each data packet required at least one retransmission, the bit-rate is lowered and the credit score is reset to zero. The current implementation also uses Atheros' multi-rate retry feature, which allows algorithms to select different rates for retransmissions of frames. [1]*Adaptive Multi-Rate Retry (AMRR)* [44], a modification of ONOE, adaptively raises the threshold for rate increases to prevent frequent attempts at bit rates higher than the optimal one in an SNR-limited channel. The recently proposed *Robust Rate Adaptation Algorithm* [47] builds on top of ARF by using a combination of short-term loss estimation and selective use of RTS/CTS. However, the authors themselves acknowledge the potential degradation in performance when the number of stations in the network increases (beyond 8), due to a lack of samples used to infer the channel quality.

2.1.2 Throughput-based adaptation

The *SampleRate* [46] algorithm selects the rate that minimizes mean packet transmission time. Initially, the lossless packet transmission times are calculated for each bit rate and an initial rate is chosen (36Mbps). Hereafter, for each successfully sent packet, the transmission time is updated (using an exponentially weighted moving average (EWMA)) based on the number of retransmissions, packet length and protocol timing overheads. The algorithm also periodically attempts transmission at bitrates whose lossless transmission time is lower than the measured time on the current rate. If these sample transmissions indeed show lower mean transmission time, the algorithm switches the rate.

2.1.3 SNR-based adaptation

Since the frame-error rate on a collision-free channel is determined by the receiver's SNR, these algorithms measure channel SNR and select the appropriate rate based on a precomputed table.

[1]Since rate selection is implemented in the device driver and retransmissions are handled on the microcontroller, the rate selection algorithm can at most be executed once for each packet inserted in the hardware transmission queue. To change the bit-rate after a certain number of unsuccessful transmissions, the hardware provides a multi-rate retry table that the algorithm can fill. This table specifies the rate to use dependent on the retransmission count for the packet. ONOE fills this table with default parameters (4 at the chosen rate, 2 at the next lower rate, 2 at next lower rate and 2 at the lowest bit-rate) and only varies the rate for the first transmission attempt.

In general, 802.11 implementations usually only provide the received signal strength indicator (RSSI). This indicator reflects the amount of energy measured on the channel during the reception of the PLCP header.[2] Receiver Based Auto Rate (RBAR) [49] defines a closed loop rate adaptation mechanism through which the receiver can inform the sender of the most suitable rate choice. Specifically, the receiver selects the bit rate based on the RSSI of RTS frames and piggybacks this information on the CTS frame. Pavon and Choi [50] propose a hybrid approach that utilizes the RSSI of acknowledgment frames to choose the bit rate. This algorithm attempts to address asymmetric channels through re-calibration of the SNR thresholds for rate choices based on the frame error rate. Another hybrid algorithm proposed in [51] utilizes RSSI to clamp frame-error based bit-rate changes.

The Opportunistic Auto Rate (OAR) protocol [52], which can be layered on top of any of the above rate adaptation mechanisms can optimize individual, as well as network throughput, by sending multiple back-to-back frames under favorable channel conditions.

As we will show later, most of these techniques are not able to scale to high node densities since they are unable to accurately distinguish between the different causes for poor performance (at a bit-rate).

2.2 Transmit Power Control

Transmission power control (TPC) for wireless networks has been a subject of extensive study. In this section, we review this large body of work, starting with their application in CDMA cellular networks. Note that, in a majority of existing multiple access technologies, the functionality of TPC is the same — regulating interference in order to control quality-of-service (QoS) metrics such as per-link throughput, fairness, and delay, as well as minimizing power consumption for mobile devices.

2.2.1 CDMA Cellular Networks

One of the central requirements in CDMA networks is that of uplink power control to solve the *near-far problem*. Due to the intrinsic nature of these networks where all active users transmit simultaneously, albeit using a different *code*, it is possible for the transmissions of a closer

[2] According to the standard, it is measured between the beginning of the start frame delimiter (SFD) and the end of the PLCP header error check (HEC)

active user to unintentionally ”jam” the transmissions of users that are relatively farther away. To avoid this situation, TPC mechanisms are used to equalize the received powers from all transmitters in the cell. Seminal work in this area involved iterative algorithms that ensured that each mobile user attains the target Signal-to-Interference Ratio (SIR) [53], followed by Foschini and Miljanic’s distributed algorithm to achieve the feasible set of target SIRs [54]. This distributed algorithm, with wide applicability in cellular networks, was followed by a large body of publications, important among which was the proof of convergence of a more general class of distributed TPC algorithms in [55]. This was followed by works [56, 57] that aimed to provide limited QoS in distributed manner. Since then, researchers have focussed on the problem of jointly optimizing SIR assignment and transmit power control over the feasibility region [58–62]. More recently, [63, 94] provide the first distributed and optimal algorithm for the joint optimization of SIRs and transmit power, according to criteria defined by the network operator, as well as a distributed mechanism to check the feasibility of SIRs.

TPC mechanisms in these networks essentially consists of continuous tracking of the SIR of each mobile at the base station receiver and corresponding control of TPC via the downlink channel. The goals are two-fold: (a) to ensure that the transmission power is high enough to attain the target SIR, and (b) to monitor link quality and correspondingly set the target SIR. This is achieved primarily through the use of closed loop power control [64], which consists of an inner loop (also called fast loop) and an outer loop. The outer loop is responsible for setting the target SIR based on estimated link quality while the inner loop is responsible for TPC adjustment to meet the target SIR. The TPC commands are affected using a single bit, which tells the mobile to either increase or decrease transmission power in a step-wise manner. To combat fast fading, the rate of change of power in the inner loop power could be anywhere between once every 10ms (100Hz) to once every 1.25ms (800Hz). In addition to closed loop power control, open loop power control is used by the mobile to set its initial value based on a measure representing the path loss, interference and target SIR at the base station receiver, which are broadcasted on a seperate control channel.

The main issue in adapting these mechanisms for CSMA wireless networks is that they do not address sender-side channel access asymmetry, which is a problem unique to CSMA networks. Another practical issue is that all these approaches assume accurate SINR estimation and as pointed out earlier (in section 1.3), accurately measuring interference in CSMA devices is infeasible. Majority of existing CSMA-based network cards provide an estimate of SINR,

called Received Signal Strength Indicator (RSSI), which is only a measure of the cumulative energy for a small duration (below 10μs) of the frame header. Thus, if there is more than one simultaneous transmission at the receiver, RSSI is Signal(S) + Interference (I) rather than S/I. This problem does not arise in CDMA cellular networks since each transmitter uses a distinguishable code. In general, this limits the applicability of SINR-based mechanisms.

2.2.2 Mobile Ad-hoc Networks (MANETs)

TPC schemes for 802.11-based MANETs have primarily focussed on three main issues: (a) network topology control, (b) energy consumption savings, and (c) spatial reuse.

In the first category [65–70], network topology is controlled through TPC. The main idea is to ensure certain topological properties while either minimizing the energy consumption for each node [65–68] or reducing interference [69, 70]. In [65], two centralized algorithms, CONNECT and BICONNECT-AUGMENT are proposed for static wireless networks with the objective of maintaining network connectivity at the lowest possible transmit power levels. In addition to proving the optimality of these approaches, the authors also propose two heuristic-based distributed algorithms, LINT and LILT, that use neighbor information collected by routing protocols so that their degree (number of one-hop neighbors) is bounded. In [66], a position-based, distributed network-layer protocol is proposed with the objective of minimizing energy while maintaining connectivity. In [67], the authors propose using directional information to maintain network connectivity, as opposed to location. In [68], the effect of heterogeneous transmit powers on energy consumption and end-to-end throughput is investigated. In [69], a greedy algorithm is proposed and analyzed that aims to reduce interference while maintaining connectivity. In [70], the authors first disprove the assumption that interference is reduced as a result of topology sparseness and then propose centralized (LIFE, LISE) and distributed algorithms (LLISE) that are proven to minimize interference and preserve connectivity. The main drawbacks of these algorithms are that they do not deal with interference issues at the MAC layer ([65] acknowledges that hidden node problems are not handled, [67] assumes that these issues are handled and [68] uses a contention-free MAC and a slotted Aloha MAC), thus not addressing the hidden node and channel access asymmetry issues. Moreover, they also do not address the issue of multiple bit-rates. Finally, the deployability of approaches that depend on GPS [66] or Angle-of-arrival (AOA) [67] or a seperate reverse channel [68] is also limited.

Kawadia and Kumar argue that power control should be a network layer function and develop the COMPOW protocol [71] in which routing layer agents are used to converge to a *common* power level for all nodes. Unfortunately, this approach will not be able to deal with the inherent sender-side asymmetry that is present in indoor environments even when transmitters use the same power level. The authors also realize that restricting all transmitters to a common power level is a conservative approach, especially when nodes are clustered [72, 73]. However, again the main drawbacks of these schemes are that they do not address both sender and receiver-side asymmetry — the LOADPOW protocol is the only one which may be able to deal with receiver-side asymmetry through its use of RTS/CTS frames but this is also is the one algorithm the authors do not to implement in a real system.

In the second category, TPC is applied on a per-packet basis to reduce energy consumption [74–77]. RTS/CTS frames are exchanged at maximum power while DATA and ACK frames are sent at lower power levels. While these schemes achieve reduction in energy consumption (relative to the 802.11 max. power approach), their best-case throughput is comparable to that of 802.11. Moreover, they do not address sender-side asymmetry.

In the third category, TPC is applied on a per-packet basis to increase spatial reuse [78–81]. The schemes in [78–80] use a seperate control channel to broadcast collision-avoidance information. As pointed out in [81], the practicality of using a seperate control channel is questionable especially since (a) it is not backwards compatible, and (b) it requires nodes to be equipped with two wireless cards. POWMAC [81] proposes the exchange of MAC control frames, on a per-packet basis, to determine the power level for DATA transmissions. Receivers measure the average interference in their vicinity and accordingly decide to either allow or deny a transmission request from the sender. All nodes periodically broadcast the maximum interference they can tolerate so that senders in their vicinity can bound their transmit powers. Another interesting feature is the use of an adjustable time window between control and DATA frames to allow for the scheduling of multiple concurrent transmissions. The main limitations of this approach are that

- it assumes that wireless cards can measure the average interference power over timescales of a few packet transmissions ($\tilde{1}$-10ms). To our knowledge, none of the cards available in today's market provide such a measure.

- it assumes that the receiver can accurately estimate the channel gain from each sender,

for which the path-loss propagation model is used. As we will show later, the pathloss propagation model does not apply in indoor environments, which makes estimation of channel gain non-trivial.

- it requires protocol modifications, which are not backwards compatible with existing wireless cards.

Since spatial reuse in CSMA networks can also be obtained using carrier-sense (CS) threshold adjustment [82], Kim et. al. [1] address the relationship between transmit power control, CS threshold adjustment and spatial reuse. Specifically, they show that, in the case that achievable channel rate follows Shannon capacity, spatial reuse depends only on the ratio of transmit power and CS threshold. Additionally, they argue using examples that tuning transmit power offers more control of SINR at the receiver than tuning CS threshold. They also propose a distributed power and rate control algorithm that requires SINR feedback from the receiver. The main limitations of this algorithm are that

- it approximates interference at receiver through interference measurement at the transmitter. As pointed out earlier, measuring interference is non-trivial in today's real-world WLANs.
- it assumes path-loss propagation, which does not apply in indoor environments.
- it requires per-packet SINR feedback, which will not work with legacy devices.
- it does not address sender-side asymmetry.

Having reviewed power control approaches in all the relevant categories in ad-hoc networks, we now proceed to look at proposals in the WLAN domain.

2.2.3 IEEE 802.11 Infrastructure Networks (WLANs)

TPC schemes for WLANs have generally focussed on improving spatial reuse and/or reducing energy consumption for hand-held devices.

In [83], a joint rate and power control algorithm is proposed with the objective of increasing the battery-life of mobile terminals. Each transmission is preceded by an RTS/CTS exchange to mitigate hidden terminal issues while the optimum rate and power is a function of the frame size, the path loss and the frame retry counts. The main drawbacks of this algorithm are that

(a) since CTS is sent at max. power, spatial reuse is only going to be as good as that of vanilla 802.11, and (b) it does not address sender-side asymmetry.

In [33], the authors propose an algorithm to jointly tune transmit power and CS threshold to prevent starvation due to sender-side asymmetry. They propose maintaining fairness by ensuring that the product of transmit power and CS threshold is constant — this implies that transmitters using high transmit powers will also be the most sensitive. Power control is performed *per cell* by the AP depending on (a) the number of clients per cell, (b) channel gain of the client with the worst channel conditions and (c) inter-AP interference. The main limitations of this algorithm are that

- it approximates inter-AP interference using RSSI of frames heard from neighboring APs. As we stated earlier, this is a crude measure of interference since RSSI is more a measure of cumulative energy (S+I) rather than SINR. Moreover, this interference measurement does not account for interference from transmitters whose frames are above the CS threshold but below the receive threshold.

- it assumes a very static setting with no mobility. The convergence time of the algorithm is reported to be 30 seconds, which clearly is too slow even with mobility at walking speeds.

- it is being conservative by setting the transmit power for a cell based on worst-case channel conditions.

- it does not address receiver-side asymmetry.

Recently, Broustis et. al. [84] use experiments on an indoor 802.11a mesh network to confirm that power control can significantly improve throughput and fairness. They identify three interference scenarios corresponding to overlapping links, disjoint links, and links that are hidden from one another. Based on an exhaustive search of power levels layered on top of the SampleRate algorithm [46], they find that keeping RTS/CTS turned ON all the time, in conjunction with power control, is detrimental to performance. We leverage this result by using an adaptive RTS/CTS mechanism for hidden node detection.

2.2.4 CSMA Sensor Networks

Reducing energy consumption and increasing network lifetime is the primary objective of power control in sensor networks [85–91] although topology control is also important.

In [87], an energy-efficient surveillance system consisting of a group of cooperative sensors to track moving vehicles, is described. As part of this system, although motes use a lower transmission power to reduce the effect of asymmetric channels, the primary focus is on power management where the motes intelligently cycle between *sleep* and *wakeup* states. Moreover, lower transmit power is only used for specific synchronization messages while a *common* transmit power is used by all the motes in the network for data messages. This is the case for the sensor system described in [89] as well.

The need for per-link power control is recognized by the authors of [86] via experiments on an indoor sensor testbed. They propose *power control with blacklisting (PCBL)*, in which, each node measures link quality, in terms of packet reception ratio, to each of its neighbors at max. transmit power. All links with quality below the *link quality control threshold* are blacklisted. The significant effect of environment and the time-varying nature of link quality motivate [91] to propose and evaluate the *adaptive transmission power control (ATPC)* algorithm, the main components of which include, an initial modelling phase and a feedback-based runtime phase. In the modelling phase, nodes exchange beacons at different transmit power levels and build a linear predictive model based on RSSI feedback from their neighbors. In the runtime phase, based on link quality feedback from the receiver, the transmit power is tuned to adapt to time-varying channel conditions. Another empirical study of TPC [88, 90] proposes a node-level TPC algorithm that aims to keep the degree (number of neighbors) of each node bounded. This algorithm is then compared, for different traffic patterns, with a fixed TPC approach. Similar node-level approaches that aim to maintain a bounded degree were proposed in [85] and evaluated using simulations. Finally, a survey of link and network-level approaches to reduce energy consumption in sensor networks is carried out in [103, 104].

The main issue with all these algorithms is that they do not address the sender-side channel access asymmetry issue. In some of these proposed approaches [86], effects of hidden terminals are mitigated using per-packet RTS/CTS, which is part of the S-MAC protocol [105]. In others, such as ATPC [91], a TDMA MAC protocol is used for evaluation and hence performance in the presence of hidden terminals is unclear at present.

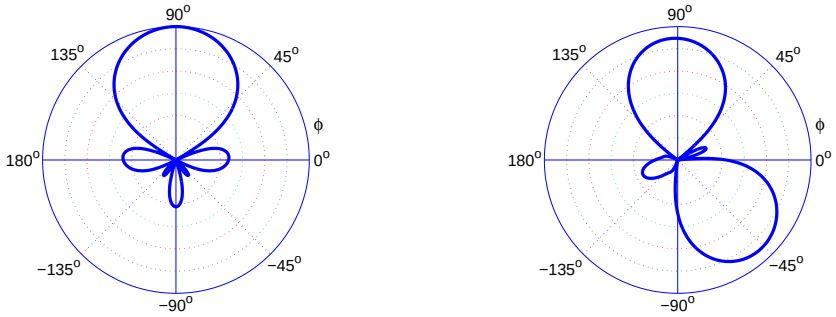

(a) 1 lobe at 90^o (b) 2 lobes at 90^o,-45^o

Figure 2.1: Example beams with different number of main lobes at different angles.

In summary, solutions in other domains were not designed to deal with the unique problems associated with TPC in CSMA wireless networks, while existing solutions for these networks do not address all the problems simultaneously or make assumptions that do not hold in real deployments. Table 2.1 provides a taxonomy of the relevant related work and identifies these drawbacks. Solutions having deployability constraints due to significant protocol modifications or impractical assumptions such as the requirement of precise interference measurements, are marked by $\times$ under 'Deployability'. Solutions not realized in practice and hence not addressing system-level challenges are marked by $\times$ under 'Realization'. Remaining columns provide information on the granularity of each solution, its objective (reduced energy consumption or increased capacity or both), and whether it addresses receiver-side interference, asymmetric channel access, and jointly adapting bit-rate.

2.3 Directionality

The notion of *directionality* corresponds to the ability of antennas to direct (beamform) energy in a desired direction, while suppressing the energy in all other unwanted directions. The footprint of the beam in the direction of maximum energy is often termed mainlobe. Increased directionality results in improved average link SNR in the desired direction, which is referred to as the *beamforming* gain.

One way of achieving directionality is to use arrays of antenna elements (referred to as smart antennas) placed in circular, linear, rectangular or other geometries. The signal sent to each of the elements is weighted in both magnitude and phase. The specific set of weights

applied to the antenna elements is responsible for the antenna radiation pattern that is created. The antenna radiation pattern for an N-element array is given by,

$$A(k) = a_0 \exp^{jkd_0} + a_1 \exp^{jkd_1} \ldots + a_{n-1} \exp^{jkd_{n-1}}$$

where a_n, d_n correspond to the magnitude and phase of the weight applied to the n^{th} antenna element respectively. An illustration of several beam patterns that we generated, along with the corresponding antenna element weights are provided later in Figure 2.1. If ρ is the average received SNR at a client due to an omni-directional transmission, then a beamformed transmission from an N element array will result in a received SNR of at most $n\rho$, i.e. the gain increases by a factor n. However, there exists a tradeoff between the beamforming gain and the mainlobe width. With increasing elements, the array gain increases by a factor proportional to the number of elements, n. However, this is achieved by focusing energy in a thin lobe of width $\frac{2\pi}{n}$, thereby decreasing the width with increasing elements. Further, since practical beamforming antennas cannot completely eliminate the energy radiated in undesired directions, they do result in some spill-over of energy in the unwanted directions, which are referred to as side-lobes. These side-lobes also increase with thinner main lobes.

In terms of related work, smart antennas are an integral part of most future wireless standards (WiMAX [106], LTE [107], 802.15.3c [108], etc.). Beamforming (directionality) is one of the core features adopted by operators to meet the high spectral efficiency requirement of future mobile applications. While several of the future wireless standards advocate the concept of directionality, they deal predominantly with protocol issues and do not consider algorithms or systems that instantiate the core mechanisms. The algorithms are left open for implementation and innovation by individual vendors. Further, the bulk of the mechanisms and sophistication in today's cellular networks is downlink-oriented. It must be noted that the directionality-diversity tradeoff explored in this work is uplink-specific, where the receivers (BS or AP) can collaborate unlike the case of mobile clients on the downlink. With the growing demand for uplink bandwidth, we believe that the identified tradeoff and the proposed solutions will be equally applicable to other wireless broadband technologies as well.

Apart from the standards, the recent work most relevant to our study is Mobisteer [109], which looks at directionality in isolation. Mobisteer attempts to improve the uplink performance by forming a beam directed at a single receiver. Several medium access control solutions also have been designed using directional antennas in multi-hop wireless networks [110, 111].

However, all these works either focus on just one facet of the agility-robustness tradeoff or address complementary issues related to directionality.

2.4 Diversity

While directionality helps improve the average link SNR, it does not alleviate the pitfalls of variance in SNR, or deep fades, resulting in channel outages and consequent packet errors. Diversity is a mechanism that tackles such deep fades. Diversity constitutes the idea of leveraging the broadcast nature of the wireless medium to receive the transmitted signal at multiple receivers and exploit the statistical independence between the channel paths to the different receivers to successfully decode the packet. Essentially, with multiple observations of the same signal, the probability that all of the independent paths fail simultaneously reduces significantly, thereby alleviating channel outages and packet errors.

While diversity-combining can be considered at multiple layers - bit, symbol, packet, etc., we focus on packet-level diversity that is amenable to implementation using off-the-shelf equipment, and is also shown to provide a large fraction of the benefits of diversity in CSMA wireless networks [112]. If p_i is the packet error rate (PER) at a receiver i, then the PER after diversity combining reduces to $\prod_{i \in L} p_i$, where L is the set of receivers involved in diversity combining. Thus, it can be seen that the resulting diversity gain is dependent on the number of receivers involved, which in turn depends on the broadcast nature of transmissions.

When the average SNR on a link is improved through beamforming, it reduces the broadcast nature of transmissions, thereby limiting its ability to leverage diversity combining to reduce PER, and vice versa. Consequently, there exists a fundamental tradeoff between using the available elements at a transmitter for directionality and diversity. This tradeoff maps to the more general agility-robustness tradeoff — using an omni-directional transmitter and leveraging full diversity across base stations is the most robust option; however, more agile solutions that leverage the available beamforming gains have a higher likelihood of maximizing network capacity. Thus, solutions that can strike directionality-diversity tradeoff effectively are of value. We explore this tradeoff analytically in Chapter 3.7, and describe a framework designed around effectively striking this tradeoff in Chapter 5.

With respect to prior work, the concept of diversity is not new — it is already being used in CDMA cellular networks (and is part of future standards as well), following the work of Viterbi

et al in [113], where the performance of the cell-edge users is improved by allowing them to receive simultaneously from adjoining base stations.

Solutions in the WiFi domain [93, 114–116] utilize *opportunistic reception* of packets due to omni-directional transmissions in order to mask-off packet loss at any individual receiver due to SNR fluctuations. However, none of the above works explore the issues involved in combining diversity and directionality. Tao et al. [117] address this issue to some extent in an ad hoc network setting, and conclude that diversity and directionality have conflicting parameter settings and further exploration is required to incorporate both approaches. They also focus on a complementary issue of developing a MAC protocol to enable nodes in an ad hoc network leverage directionality and diversity together. In contrast, this dissertation addresses the tradeoff by devising the system support and run-time adaptation required for choosing the beams, receivers and bit-rate jointly in a highly mobile environment.

2.5 Mitigating the Effects of Vehicular Mobility

Recently, a number of research efforts have focused on different aspects of improving connectivity to moving vehicles. Bychkovsky et al. [118] study the possibility of using organic WiFi deployments for providing network connectivity to moving vehicles. They investigate the effectiveness of a caching technique to reduce the overhead of IP address acquisition. The work also focuses on uploads rather than downloads to cars. Ott et al. [119] discuss an architecture and protocol to make applications disconnection-tolerant by maintaining application sessions despite connectivity interruptions. System support for fast association to APs and optimizations at the TCP level to improve throughput for moving vehicles is discussed by Eriksson et al in [120]. Rodriguez et al [121] introduce a wireless multi-homed device (MAR) for moving vehicles that dynamically aggregates channels (and hence bandwidth) across several technologies to meet the bandwidth requirements of moving users. Detailed studies on the factors affecting connectivity to moving vehicles is performed by Hadaller et al in [122], where they conclude that lack of environmental awareness is the fundamental underlying cause of several problems. Our exploration in this work is complementary to the above approaches and could hence be integrated. Further, our location-based beam and bit-rate selection algorithm instantiates environmental (location) awareness into the adaptation process, the benefits of which have already been demonstrated in several other works [109, 123–125].

Domain	Solution	Granularity	Realization	Deployability	Objective		Rate Adaptation	Channel Access Asymm.	Hidden Nodes
					Energy	Capacity			
WLANs	[Sheth02] [95, 96]	Per-link	√	√	√	√	√[a]	×	×
	MiSer [97]	Per-link	×	√	√	×	√	×	√
	[Qiao03] [83]	Per-link	×	√	√	×	√	×	√
	PARF, PERF [30]	Per-link	√	√	×	√	√	×	×
	[Chevillat05] [98]	Per-link	×	√	√	√	√	×	×
	Contour [99]	Per-link	√	×[b]	√	√	×	√	×
	[Mhatre07] [33]	Per-cell	√	√[c]	×	√	×	√	×
Wireless Ad-hoc Networks	LINT, LILT [65]	Per-node	√	√	√	√	×	×	×
	ConeBased [67], R&M [66]	Per-node	×	×[d]	√	×	×	×	×
	LLISE [70]	Per-node	×	×	√	√	×	×	×
	PCMA [79]	Per-link	×	×	×	√	×	×	√
	BASIC, PCM [76]	Per-link	×	√	√	×	×	×	√
	PCDC [80]	Per-link	×	×	√	√	×	×	√
	POWMAC [81]	Per-link	×	×[e]	√	√	×	×	√
	SHUSH [100]	Per-link	×	×[f]	√	√	×	√	√
	PRC [1]	Per-link	×	×	√	√	√	×	×
	TACP [101]	Per-link	×	×	√	√	×	√	√
	CONNECT, [65]	Per-cell	√	×[g]	√	√	×	×	×
	BICONN-AUGMENT								
	COMPOW [71]	Per-cell	√	√	×	√	×	√	×
	LIFE, LISE [70]	Per-cell	×	×[h]	√	√	×	×	×
Wireless Sensor Networks	LMA, LMN [85]	Per-node	×	√	√	×	×	×	×
	[Son04] [86]	Per-link	√	√	×	×	×	×	√
	[Jeong05] [88, 90]	Per-node	√	√	√	×	×	×	×
	ATPC [91]	Per-link	√	√	√	×	√	×	×
CDMA Cellular Networks	[Foschini93] [54]	Per-link	×	√[i]	×	√	√	×	√
	[Saraydar01] [58]	Per-link	×	√	√	√	√	×	√
	[Hande08] [94]	Per-link	×	√	×	√	√	×	√
	[Zander92] [53, 102]	Per-cell	×	×[j]	×	√	√	×	√
	UBPC [62]	Per-cell	×	×[k]	×	√	√	×	√
	[Chiang04] [60]	Per-cell	×	√	√	√	√	×	√

[a] Rcvr. tracks avg. RSSI

[b] Requires tight time sync.

[c] Frame format change

[d] Require GPS or precise Angle-of-arrival measurements

[e] Protocol change

[f] Protocol change

[g] Centralized solutions

[h] Centralized solutions

[i] Requires accurate interference measurement

[j] Unclear whether proposed solutions will work in indoor environments

[k] Centralized solution

Table 2.1: Taxonomy of existing transmit power control algorithms in WLANs, ad hoc networks, sensor networks and CDMA cellular networks.

Chapter 3

The Resource Parameter Adaptation Challenge

In this chapter, using an extensive set of experiments, we show that existing approaches for adapting bit-rate, transmit power, and directionality, in CSMA wireless networks, fail to strike the agility-robustness tradeoff effectively.

With bit-rate, the most robust choice would be the lowest bit-rate as it has the maximum FEC. However, this choice would also result in the lowest overall throughput. For bit-rate adaptation, the agility-robustness tradeoff maps to using the highest bit-rate possible without compromising on robustness. Similarly, with transmit power, the maximum value permitted by the regulatory bodies (e.g. FCC), which is frequency-dependent, would be the most robust choice for any single link. However, this would also reduce opportunities for spatial reuse and increase energy consumption. Striking the agility-robustness tradeoff for dynamic transmit power control maps to using the lowest transmit power level possible without reducing the link throughput. Finally, to counter fast fading and shadowing at vehicular speeds, with regards to directionality, the most robust choice would be to use an omni-directional transmitter and leverage base station diversity [93]. However, this increased link robustness would come at the cost of reduced reduced link and network throughput. By increasing the average link SNR, a directional transmitter would allow the usage of higher bit-rates, and hence achieve higher link throughput. At the same time, by reducing the spatial footprint of transmissions, it would allow for increased spatial reuse, thereby increasing network throughput. For directionality, the agility-robustness tradeoff maps directly to the directionality-diversity tradeoff that we explore in this work.

In this chapter, we primarily focus on the need for resource parameter adaptation to be agile and robust to interference and mobility. With bit-rate, we first demonstrate how existing bit-rate adaptation approaches fail to differentiate between packet losses caused by interference, and those caused by channel degradation. We then focus on how approaches are either robust to mobility or agile but not both. For adaptive transmit power control, we identify the

different cases of link asymmetry that can be introduced without interference-awareness, discuss its interaction with bit-rate adaptation, and address agility-related issues to mobility. For directionality, we analytically explore the tradeoff between directionality and diversity, while identifying several parameters of importance—such as bit-rate, link SNR, packet error rate, beam width, number of receivers covered for a given beamwidth, and number of antenna elements on the client—that guide the design of our solution in Chapter 5. We begin by stating our assumptions.

3.1 Assumptions

We make the following assumptions in our work:

- The distributed network architecture of today's CSMA wireless networks is here to stay. This distributed architecture is an important factor that results in the requirement for parameter adaptation techniques to be distributed. As this architecture has seen little change even in next generation WiFi standards (such IEEE 802.11n), we expect this assumption to hold at least in the near future.
- Throughput or network capacity maximization is important for a majority of applications.
- With regards to the effect of dynamic transmit power control on energy consumption, future hardware and protocol improvements will reduce the idle-time power consumption of wireless network interfaces. We also expect technology improvements to reduce the power consumed by the other components of a mobile device (such as the processor, graphics display, etc.) [126].

3.2 Bit-Rate: Agility and Robustness to Interference

Our analysis begins by comparing existing state-of-the-art rate adaptation algorithms by how well they maximize performance under increased co-channel interference. This is followed by an analysis of how they adapt to user mobility.

3.2.1 Effect of Node Density: MAC collisions

A frame collision occurs if two simultaneously transmitted frames interfere at the receiver, so that frames are lost. To avoid collisions, the IEEE 802.11 distributed coordination function

(DCF) employs a CSMA/CA mechanism [15]. Carrier sensing prevents transmissions that start while another transmission is in progress. A random backoff mechanism is used to reduce the probability of two stations simultaneously starting a transmission. Specifically, on detecting the wireless medium to be idle for a DCF interframe space (DIFS) duration, each station initializes a counter to a random number selected uniformly from the interval [0, CW(*retransmission*)-1] and starts counting down.

Time is slotted and the countdown halts when the medium becomes busy, resuming only after the medium is idle again for a period DIFS. Given its half duplex nature, 802.11 transmitters require the receiver to send an acknowledgment (ACK) after a short interframe space (SIFS) duration. The absence of an ACK is interpreted as a collision, following which, CW is doubled (until a maximum value of *CWmax*) and the process repeated. CW is reset to its minimum value, *CWmin* (16 for 802.11a and 32 for 802.11b) after successful transmissions, as well as when the maximum retry limit is reached. Note that, for a CSMA/CA MAC, simultaneous transmission of frames can occur either because the two senders (a) happen to select the same time slot for transmission or (b) cannot hear each other's transmissions (hidden terminals).

3.2.2 Discussion: expected performance

In the absence of hidden terminals, the probability that two nodes select the same time slot increases with the number of stations and the load on the network [127], since nodes reset their contention window to *CWmin* after every successful transmission. This increase in collision-based packet errors leads the auto rate fallback (ARF) algorithm to unnecessarily decrease bitrates as observed in [128], even though the interference from collisions is usually strong enough to prevent decoding even at the lowest rate. Moreover, transmissions at lower rate consume more time, decreasing the overall network throughput [129]. This anomaly occurs due to bit-rate diversity — hosts using lower bit-rates limit the throughput of hosts using higher bit-rates. These shortcomings are addressed in the design of the Collision-Aware Rate Adaptation Algorithm [130] based on ARF, which uses RTS packets to probe the state of the channel—the loss of an RTS frame is interpreted as a collision loss rather than being due to low SNR. To avoid the overhead, the sender invokes the RTS/CTS exchange only after a DATA frame transmission failure at the current bit-rate. If an ACK is not received (after an RTS/CTS exchange), the

Figure 3.1: ORBIT testbed setup

algorithm interprets this event as being due to poor channel conditions and drops the bit rate.[1]

For throughput-based algorithms, one might expect that the collision probability remains independent of the rate choice and that the collisions should cancel each other out when comparing different rates. Based on this assumption, these algorithms should be resilient to congestion. SNR-based algorithms are expected to perform optimally in congested environments. However, it is unclear whether the RSSI provided by a majority of 802.11 implementations reflects the SNR or the Signal-to-Interference-and-Noise ratio (SINR).

Overall, the collision resiliency of many of the above-mentioned rate adaptation techniques (apart from ARF) remains, to our knowledge, experimentally unexplored. This motivates our experimental study of these algorithms in a controlled, high-density setting.

3.2.3 Experiment Design

The experimental setup primarily consisted of infrastructure WLANs emulated on a large-scale indoor testbed. This section describes the testbed setup and our methodology in detail.

Controlled, large-scale testbed setup: Our study is based on systematic experiments on a preliminary version of the ORBIT indoor testbed [131]. This testbed consists of 64 nodes (standard Linux PCs), each of which is equipped with two wireless 802.11a/b/g interfaces. Half of these nodes use the Atheros 5212 chipset-based wireless NICs and the remaining use Intel 2915 chipset-based wireless NICs. The nodes are placed in a two-dimensional rectangular grid separated by 1-meter distance (see Figure 3.1) and the antennas are mounted on the sides in 135 and 215 degree positions (viewed from the top). The testbed nodes run Linux and we

[1]The authors also present an enhancement using the Clear Channel Assessment (CCA) feature of 802.11 – if the channel is not idle immediately after the reception of a DATA frame for SIFS period, it is interpreted as a frame collision. However, this functionality will be difficult to implement without requiring firmware changes for current wireless NICs.

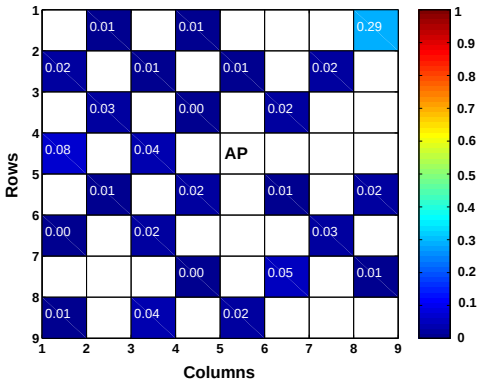

Figure 3.2: Single link (no collision) mean packet error rate from each sender node to the AP at 54 Mbps (across five six-minute experimental runs.

utilize only those nodes with Atheros chipset-based wireless NICs. Several features of this platform facilitate our research on rate adaptation. First, since the rate adaptation mechanisms are implemented in the open-source MadWiFi driver [45] (rather than in the firmware), we can develop new algorithms and modify existing ones. Second, by offloading most of the MAC protocol processing to the node's CPU, these cards are more open to protocol modifications. MadWiFi allows for the configuration of a number of MAC parameters, including the transmission rate, on a per-frame basis. Third, the platform provides a controlled and repeatable environment, where surrounding objects are stationary. Shielding in the walls of the room, housing this testbed, limit the effect that outside interference could have on experimental results. Also, we are not aware of any other adjacent 802.11 networks operating in the 5GHz band (confirmed using an additional sniffer to ensure that no background traffic exists on the channel in question).

In addition, we have instrumented the MadWiFi driver to report both successful and failed transmissions at the sender, as well as successful frame reception at the receiver (every 100ms). Given a constant packet size this allows for goodput calculations. The driver was also modified to report the source MAC address, RSSI, bit-rate and hardware timestamp (microsecond resolution) for each successfully received frame.

Target environment and node calibration: We focus on an infrastructure-based 802.11a system in which, all nodes are within communication range of each other emulating future very high density deployments. We vary the number of clients from 2 up to 20 (we could not use the remaining 11 Atheros-based nodes due to hardware issues). To characterize the radio links, we rely on single-link received signal strength indicator (RSSI) values and packet

error rate (PER). RSSI is an estimate of the signal energy at the receiver and is reported by all commodity wireless NICs on proprietary scales (Atheros cards report RSSI in dB relative to the noise floor). In our setup, RSSI measurements serve to approximate "true" SNR values, which would require a calibrated comparison with an accurate RF measurement device. In essence, they represent SNR as measured by actual 802.11 radios and although our results may not apply to future radios with improved measurement accuracy, we believe that our findings have significant implications with regard to practical mechanisms which must depend on similar measurements in real deployments. From our single-link experiments, we observe that link RSSI values (not shown here due to space constraints) range between approximately 30 and 60, which translates to an SNR of -65dBm to -35dBm (assuming constant noise floor of -95dBm), indicating good to excellent connectivity [132]. Figure 3.2 confirms that all links support the highest bitrate (54Mbps) with near-zero packet error rate (in the absence of contention).[2]

Note that, although our experiments use a single AP, we believe that our results serve to highlight the significant issues, related to rate adaptation, in congested environments. We believe that the same issues will assume significance in networks consisting of multiple APs on the same channel, albeit with fewer clients per AP (for e.g. home wireless networks with 3-4 clients per AP).

Network traffic and candidate rate adaptation techniques: In our experiments, clients generate constant bit-rate UDP traffic (using the ORBIT traffic generator [133]) to emulate streaming media applications. Further, recent IETF measurement studies [134], which show that highly congested environments represent realistic scenarios, motivate our study of these algorithms under network saturation. Other advantages of this approach are that it enables comparisons with prior theoretical work [127] and provides an estimation of the worst-case performance.

We carry out multiple runs of each experiment, and the results presented are the average

[2]Note that we do not use the nodes with poor PER (due to defective NICs).

Parameter	Default Setting
Mode	802.11a
Channel	36
Transmit Power	18 dbm
Packet Size	1350 bytes

Table 3.1: Default configuration parameters

over all runs. We empirically choose the experiment durations so as to provide low variance in results. We also vary the packet size in our experiments. Due to space constraints, we only present the results with 1350-byte packets in this paper. Unless otherwise specified, for the results presented in this paper, the default configuration parameters are specified in table 3.1.

We evaluate and compare ONOE [45], SampleRate [46], and Pavon and Choi's algorithm [50] as representatives of the packet-error-, throughput-, and SNR-based categories. We attempt to measure the performance of CARA [130] by approximating its behavior using a combination of ONOE, which is similar to ARF [43], with RTS/CTS enabled for all frames. We also report results for SampleRate with RTS/CTS enabled, since this configuration can serve as an indicator of how throughput-based approaches, in conjunction with RTS/CTS, will perform. Note that we use vanilla versions of SampleRate and ONOE and implement Pavon and Choi's RSSI-based algorithm. We believe that these selected algorithms provide a good sample of representative designs in literature.

3.2.4 Experiment Results and Analysis

We begin by comparing the rate adaptation algorithms by how well they maximize cumulative throughput under congestion.

Cumulative Throughput: Figure 3.3(b) shows an experimental comparison of the rate adaptation algorithms in the access point scenario, with a fixed rate of 54 Mbps (i.e., deactivated rate adaptation). For reference, Figure 3.3(a) also depicts analytical saturation throughput curves for the same scenario, obtained using Bianchi's model [135] with 802.11a parameters. These curves assume fixed (no adaptation) PHY rates and predict a graceful degradation in cumulative throughput. The analytical results show about 14% reduction in throughput when the number of transmitting nodes increases from 2 to 20. The experimental results for fixed rate (deactivated rate adaptation) closely track this performance. While deactivated rate adaptation cannot represent a useful approach in general, it illustrates that the basic MAC protocol scales as expected.

With both ONOE and SampleRate, the cumulative throughput drops with an increase in the number of transmitters. As the number of transmitting nodes increases from 2 to 20 nodes, the throughput falls by more than half, compared to a drop of less than 15% corresponding to the single bitrate analytical results. SampleRate performs slightly better than ONOE when

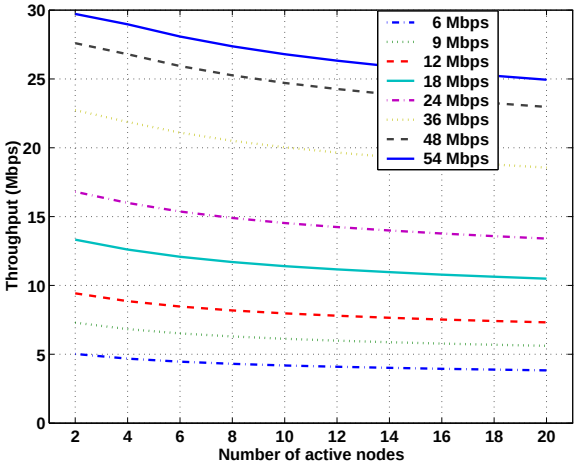

(a) Analytical saturation throughput for fixed bit-rate as a function of the number of active nodes.

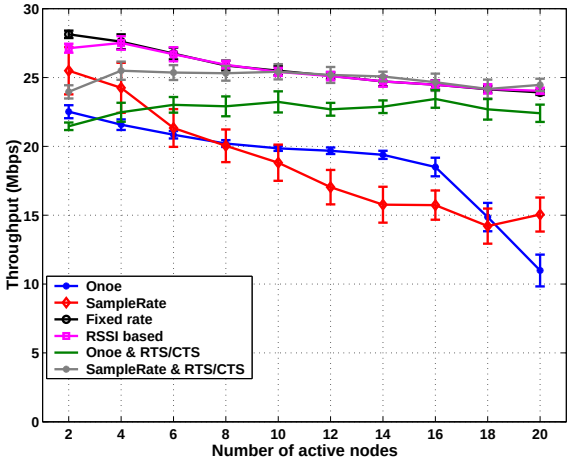

(b) Experimental saturation throughput for the different rate adaptation algorithms vs a fixed 54Mbps rate approach.

Figure 3.3: Comparison of theoretical and empirical throughput.

the network size approaches 20 nodes, whereas, ONOE maintains a higher throughput when the number of users is between 6 to 16 nodes. Given that the average frame transmission time would be minimal at the highest bit-rate, even in congested environments, we would expect SampleRate to perform much better.

RSSI-based rate adaptation appears resistant to collisions and shows excellent performance, in terms of cumulative throughput.

Note that ONOE shows significant throughput improvement when RTS/CTS is enabled. Additionally, RTS/CTS benefits SampleRate as well, with cumulative throughput approaching

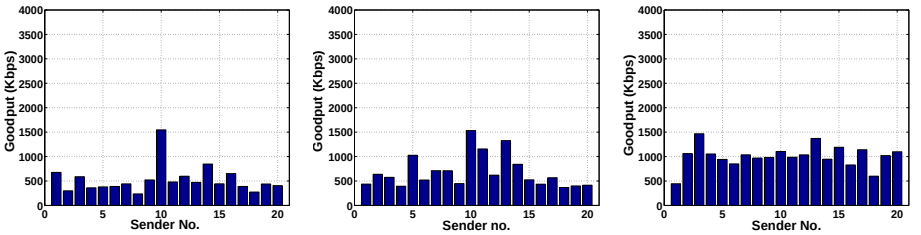

(a) Individual flow throughputs for ONOE

(b) Individual flow throughputs for SampleRate

(c) Individual flow throughputs for fixed rate (54Mbps, no adaptation)

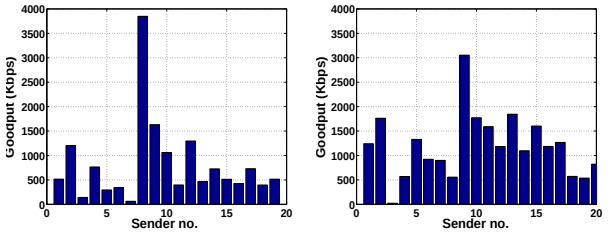

(d) Individual flow throughputs for ONOE with RTS/CTS

(e) Individual flow throughputs for SampleRate with RTS/CTS

Figure 3.4: Throughput fairness characteristics of rate adaptation algorithms from an experimental run. We observed similar trends across multiple runs of this experiment.

the performance of fixed rate. For both algorithms, the throughput improvements can be explained, in part, by the smaller time spent in collisions through the use of RTS (which is smaller than the frame header for 802.11 DATA frames), even though it is sent at the lowest bit-rate. More significantly, this implies that the reduced time spent in collisions outweighs the overhead of using the RTS/CTS exchange in such environments.

In summary, we observed improved throughput for RSSI-based adaptation and through the use of RTS/CTS. We also notice lower than expected throughput for SampleRate. Since throughput gains can be easily achieved at the expense of fairness, let us now look at the throughput fairness characteristics of these algorithms.

Table 3.2: Fairness comparison for the 20-sender case. Mean and Std. Dev. in JFI across 5 runs is reported.

Rate adaptation scheme	Avg. JFI	Std. Dev. in JFI
ONOE	0.822	0.032
SampleRate	0.819	0.024
Fixed Rate (54Mbps)	0.917	0.027
ONOE w/ RTS/CTS	0.491	0.020
SampleRate w/ RTS/CTS	0.709	0.034

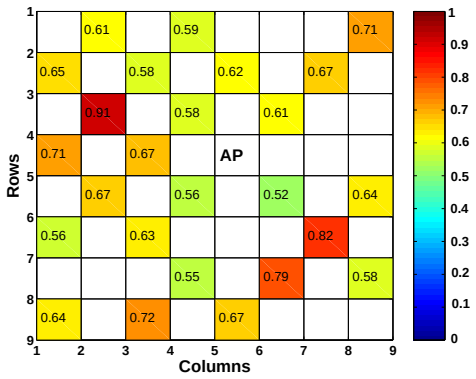

Figure 3.5: Mean packet error rate (PER) from each sending node to the AP when the network is saturated (across five six-minute experimental runs).

Fairness: Table 3.2 reports the mean and std. dev. in Jain's fairness index (JFI) [136][3] for the RTS/CTS-based approaches with that for ONOE, SampleRate, and fixed rate for the standard 20 sender experiment. ONOE with RTS/CTS adaptation stands out with the lowest fairness index. To analyze fairness in more detail, figures 3.4(a) through 3.4(e) compare the throughput distribution across senders for each of the algorithms. Fixed rate shows slight imbalances that, as we will see, is due to the physical layer capture (PLC) effect.

The presence of PLC is illustrated in the PER imbalances observed in a 26 node setup in Figure 3.5 (we also confirmed PLC by looking at packet traces from multiple sniffers). PER for each link in saturation ranges from approximately 50% to 90%. Since all other parameters in this experiment were the same as that for Figure 3.2, we can attribute these PERs solely to collisions. Note that the PER of nodes with lower RSSI (relative to the stronger sender) at the access point tends to be higher, a typical result under PLC.

SampleRate and ONOE both show more pronounced throughput variations, most likely because rate diversity increases the probability of capture as shown in the previous subsection. SampleRate with RTS/CTS shows significant throughput imbalances also reflected by its relatively low Jain fairness index. ONOE with RTS/CTS clearly shows the largest throughput imbalances. A closer inspection of the packet error traces in the experiments involving ONOE with RTS/CTS reveals that for a majority of senders, PER was higher than 10% (ONOE's

[3]The index, F, is calculated as $F = \frac{(\sum_i x_i)^2}{n \times \sum_i {x_i}^2}$ where x_i is the individual flow throughput and n is the total number of flows. An index value equal to one is considered to be perfectly fair.

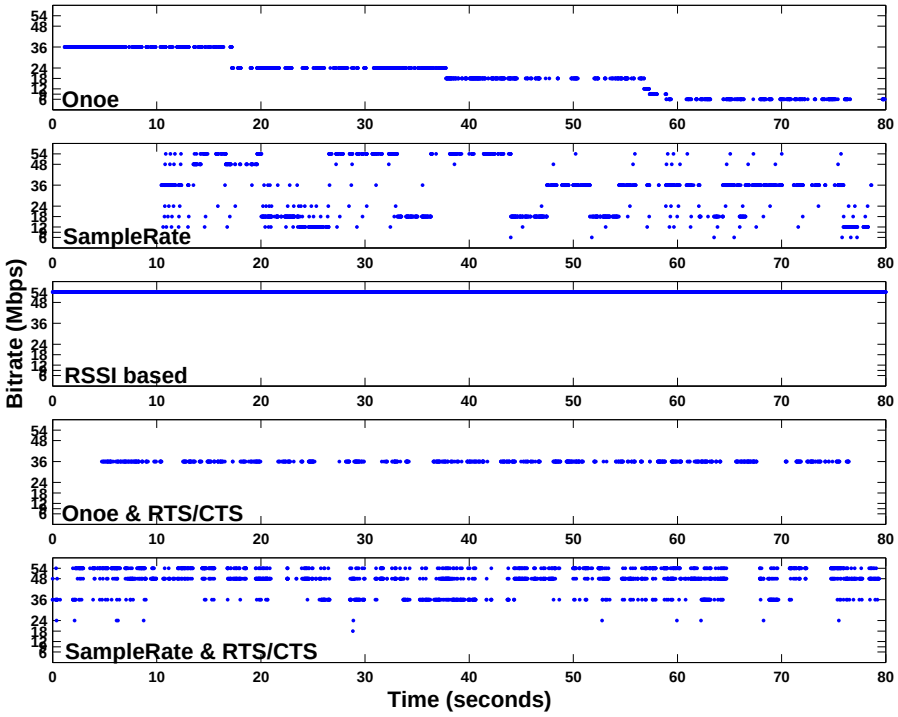

Figure 3.6: Rate vs time for rate adaptation algorithms from an experimental run. Transmission rate for both SampleRate and ONOE decreases with the addition of nodes with SampleRate showing higher variance.

threshold for rate increase) and below 50% (threshold for rate decrease). This indicates that significant DATA frame losses occur even though the channel is reserved and links are not SNR-limited. To investigate this anomaly, we carried out single-link measurements, with and without RTS, and observed that the PER, unexpectedly, shows a relative increase of approximately 4% when RTS/CTS is enabled. We speculate that this indicates an incorrect implementation of this mechanism on Atheros 5212 NICs (also discovered by [137] on different hardware) and we hypothesize that the throughput gains will be higher with more accurate implementations. In addition, the observed fairness reductions may be also caused by this issue. We plan to investigate this issue further as part of our future work.

The surprisingly low throughput of SampleRate motivates us to look more closely at the bit-rate choices of the individual algorithms.

Bit-rate choices: Figure 3.6 depicts transmission rate changes over time for one of the senders. Note that in this 80 second experiment, all 20 senders start simultaneously. The bit-rate (obtained from the PLCP header) for each received packet is logged at the access point—each

dot in the plot represents one packet. The results show that ONOE starts off at its default initial rate of 36Mbps and steadily decreases the bitrate until it reaches 6Mbps. SampleRate shows a similar trend even though it is not as pronounced because the algorithm tends to change rates more frequently. For example, around 30s into the experiment, transmissions occur at nearly all the rates between 12 and 54Mbps. This confirms, however, that poor bitrate decisions are also the cause for the reduced throughput obtained with SampleRate.

In comparison, RSSI-based adaptation shows nearly perfect rate choices, with all nodes choosing high bitrates. This indicates that RSSI-based adaptation is not affected by collisions. ONOE with RTS/CTS remains at 36Mbps for the entire experiment. This is contrary to what we expect – the use of higher bit-rates by transmitters, provided only DATA frame losses are taken into account (while calculating PER). However, as mentioned before, we do see DATA frame losses, even when RTS is turned ON. SampleRate with RTS/CTS shows bit-rate fluctuations, but critically, we can see that it is more prone to select the higher bit-rates (48Mbps and 54Mbps). We believe that more accurate bit-rate choices are the primary reason for the difference in cumulative throughput gains between ONOE and SampleRate, when RTS/CTS is enabled. The adaptation stability of both algorithms is significantly improved with RTS/CTS, nevertheless, SampleRate still lacks stability due to a reduction in the number of measurement samples. We highlight some of the more interesting performance details characterizing the selected rate adaptation algorithms in the sections that follow.

SampleRate with, and without, RTS/CTS: SampleRate's rate decisions compare the expected transmission time of different rates relative to each other. One might expect, based on arguments offered in section , that this algorithm is resilient in high collision environments. Surprisingly, SampleRate's performance degrades with increasing node density, similar to ONOE. We identified two reasons:

1. In highly congested environments, few samples (packets) per node are available to accurately estimate the transmission time.

2. Due to the PLC effect [138], some nodes can decrease their collision probability by decreasing their rate while maximizing their individual throughput.

The frequent rate changes observed in Fig. 3.6 support that the algorithm bases its rate choice on too few samples. To confirm that nodes can maximize throughput by lowering rates even on the high SNR channels in our setup, Figure 3.7(a) shows the goodput for different rate

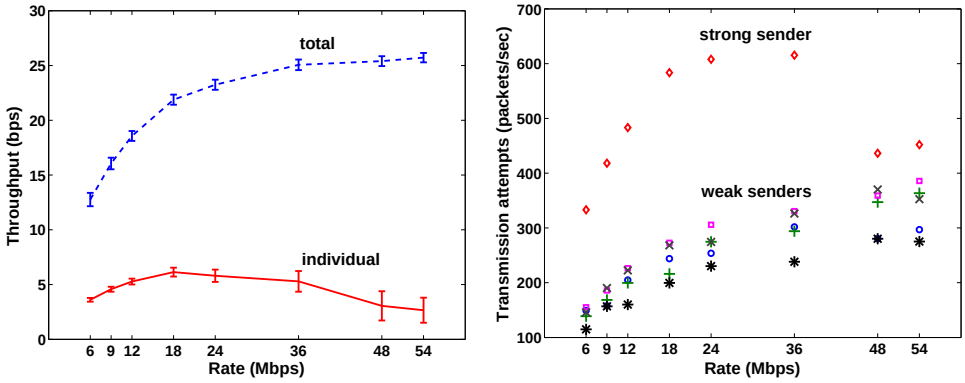

(a) Individual goodput vs network goodput for different rate choices of one node, under PLC.

(b) Number of transmission attempts for different nodes depending on the rate choice of one node.

Figure 3.7: Performance in the presence of Physical Layer Capture (PLC).

choices for a station that is closer to the access point as compared to the other 10 competing stations. The competing stations use a fixed rate of 54 Mbps and the average RSSI observed by the access point for the close and far stations are 60 and 47, respectively. Evidently, the closer station obtains maximum throughput while operating between 18Mbps and 36Mbps. We can explain this result with a capture probability that depends on the bit rate choice of the stronger sender—capture becomes more likely when the stronger sender reduces its rate. Each capturing transmission in turn causes the capturing node to reset its contention window to CW_{min} while the other colliding nodes double their current contention windows. This leads to a larger number of transmission opportunities for the capturing node, at the expense of reduced transmission opportunities for the other nodes, as shown in Figure 3.7(b). In effect, the station sacrifices total network throughput for a small gain in individual throughput. While Tan and Guttag [139] have noted the existence of such inefficient equilibria through simulations of mobile nodes under a Rayleigh fast fading model, these results show that these inefficiencies also exist in relatively stationary multiple-client single-access point scenarios.

To further validate this hypothesis, Figure 3.8(a) translates the observed packet error rates into the expected packet transmission time, on which SampleRate bases its rate decision[4]. Indeed, the minimal expected transmission time for the stronger sender occurs at 36Mbps. This

[4]The theoretical curve derivation is outlined in [48].

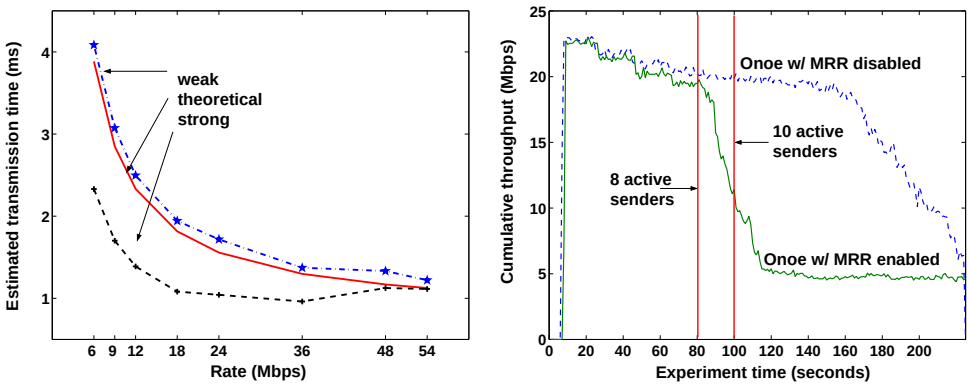

(a) Estimated average transmission time for two senders - the "strong" sender is able to capture transmissions from "weak" senders.

(b) Cumulative throughput with and without multi-rate retry (MRR) for ONOE. Behavior is similar for other algorithms.

explains why stations choose lower rates. Moreover, the differences in expected packet transmission times between the rates of 18Mbps to 54Mbps are small. This explains the oscillatory behavior in Figure 3.6. As mentioned earlier, SampleRate with RTS/CTS selects higher bit-rates with greater frequency and shows corresponding increase in cumulative throughput (relative to when RTS is disabled).

ONOE with, and without, RTS/CTS: Auto rate fallback is known to lead to degraded performance with less than 10 senders. Since vanilla ONOE also bases its rate decisions on packet errors, one might expect similar performance. Instead, the cumulative throughput with ONOE remains more stable until a significant reduction occurs with 18 active senders. The exact number of senders tolerated is, however, very sensitive to the detailed algorithm configuration. When the multi-rate retry (MRR) feature in Atheros cards is enabled, throughput collapse occurs with just 10 senders as depicted in Figure 3.8(b). Since this mechanism is configured to pick lower rates for retransmissions, we hypothesize that the pathological effects of collision on packet-error-based adaptation are amplified by MRR. All ONOE results in this paper were obtained with MRR disabled.

Results from ONOE with RTS/CTS indicate that the performance of packet-error-based adaptation can be stabilized through channel reservations, as proposed in CARA [130]. Here, ONOE is modified to only consider packets that did not contend with other stations (e.g., the data frame following a CTS), thus avoiding unnecessary rate decreases (due to RTS losses). However, as mentioned before, ONOE does not increase the bit-rate, as would be expected when RTS/CTS is used, in near-perfect channel conditions, likely due to implementation issues.

Now that we have looked at the performance of ONOE, SampleRate and RTS/CTS-based rate adaptation in detail, we proceed to highlight some practical implementation issues.

3.2.5 Implementation Experiences

In this section, we first discuss precision issues associated with the reporting of SNR in existing wireless NICs. This is followed by a discussion on how the use of RTS/CTS enables the accurate estimation of channel quality.

RSSI-based rate adaptation: From the comparative evaluation, RSSI-based algorithms proved to be more resistant to collisions in a congested scenario. However, given that the RSSI thresholds for all rates lie in a small interval of the total RSSI measurement range, there is a low margin of error w.r.t. comparison with thresholds to increase (or decrease) the bit-rate. Hence, we expect that these algorithms will fail to perform optimally in SNR-limited environments.

Table 3.3 lists the RSSI threshold values for which the frame error rate (FER) approaches 1.0 for any of the 802.11a rates. We measured these thresholds by placing an (additive white gaussian noise) AWGN source [131] near the receiver, fixing the sender's bit-rate and steadily increasing noise power until the receiver did not decode any frames. The RSSI values for the frames decoded last were noted as the approximate RSSI thresholds. These thresholds are specific to the Atheros 5212 card because the absolute interpretation of RSSI values is not defined in the standard. However, for convenience, many manufacturers use a similar scale where each step in RSSI signals an increase of approximately one dB in signal strength. According to simulations of the modulation schemes, they cover a range of 20 dB [140]. Thus, we would expect the thresholds to lie within an interval of 20 RSSI values in most implementations. With thresholds compressed into such small intervals, slight measurement errors might have a large effect on RSSI-based rate adaptation. To make RSSI useful as a primary indicator for rate selection, wireless NIC's should provide more fine-grained RSSI measurement differentiation in the range relevant to bitrate selection. In addition, protocols for exchanging receiver RSSI information must be defined to allow a pure RSSI-based approach, as opposed to Pavon's hybrid RSSI/frame-error-based algorithm.

Table 3.3: RSSI thresholds for 802.11a PHY bit-rates

Rate (Mbps)	6	9	12	18	24	36	48	54
RSSI Threshold	9	10	11	12	13	15	19	23

RTS/CTS-based rate adaptation: In congested environments, both SampleRate and ONOE make rate adaptation decisions based on "collision-tainted" measurements obtained during a fixed time interval (ONOE uses frame retry measurements and SampleRate measures average frame transmission time). Measurements are tainted since, an increase in congestion results in a reduction in the number of available samples and simultaneously, these samples are more likely to be "affected" by collisions. For ONOE, the number of frame retry and success samples that are available to infer PER will be lower, thereby affecting the accuracy of PER estimation. SampleRate will also suffer from the same issue — reduction in the number of samples, which in turn, will affect accuracy in estimating the average frame transmission time.

The use of RTS/CTS ensures that the rate adaptation decisions are made solely on measurements not tainted by collisions (provided RTS transmission errors are not taken into account). Both, packet error-based mechanisms (as proposed by CARA [130]) and throughput-based mechanisms infer channel quality on DATA frames. Thus, they exhibit a stable behavior with an increase in the number of nodes. However, in practice, we observed reduced gains due to the likely implementation issues with RTS/CTS, which results in DATA frame losses even when the channel is reserved and not SNR-limited.

An alternative to RTS/CTS-based collision detection includes the *passive estimation* of PER due to collisions [141]. However, although the proposed PER estimation technique appears promising, it is yet to be experimentally evaluated. Moreover, the proposed technique requires precise information regarding CSMA/CA slot usage at each IEEE 802.11 transmitter and this information, to our knowledge, is not exposed by the majority of existing open-source device drivers. Similarly, the dynamic tuning of MAC contention windows based on achieving a balance between the time spent in collisions and the time spent waiting in idle slots [142] requires a number of changes to the MAC, which may not be possible to implement on existing wireless NICs.

3.3 Bit-Rate: Agility and Robustness to Mobility

User mobility introduces additional requirements in the form of responsiveness and stability. In this section, we analyze the performance of SampleRate and *RRAA* [47] to demonstrate the need for enhancements to improve convergence and robustness characteristics.

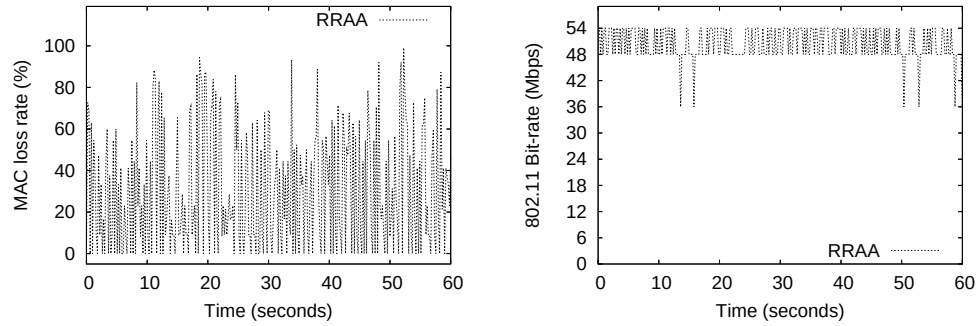

(a) RRAA MAC loss rate (b) RRAA bit-rate

Figure 3.8: Convergence characteristics of RRAA.

3.3.1 Discussion: expected performance

In wireless networks with user mobility, a rate adaptation algorithm should satisfy two characteristics: (1) it should react to changing channel conditions due to mobility, (2) it should be able to differentiate mobility-induced channel changes from other reasons for poor performance such as hidden terminals. To address the second characteristic, it is important for the rate adaptation algorithm to converge to a particular rate so as to reduce false positives.

With regards to the first characteristic, Wong et. al. [47] have shown through coarse-grained metrics (end-to-end throughput) that RRAA is more responsive than SampleRate. We expect that to be the case due to the faster rate of adaptation in RRAA as compared to SampleRate. With regards to convergence and stability, one may expect, based on results reported in [47] that RRAA is able to converge to the optimal rate. However, the lack of an explicit convergence mechanism in this algorithm does raise questions.

Algorithm 1 RRAA

1: **if** ($loss > HT_k$) **then**
2: k ← next_lower_rate
3: **else if** ($loss < LT_k$) **then**
4: k ← next_higher_rate
5: **end if**

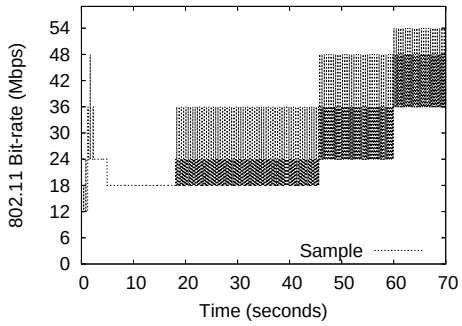

Figure 3.9: SampleRate's performance with user mobility

3.3.2 Experiment Design, Results and Analysis

To understand further the convergence characteristics of RRAA, we implemented it in the Linux MadWiFi driver. An implementation of SampleRate exists already in the MadWiFi package. We faced some non-trivial implementation challenges that we elaborate upon in chapter 3.4. We now perform two experiments to demonstrate the instability of RRAA and the slow adaptation of SampleRate. First, we consider one AP and one client, where a NLOS client is kept stationary at a distance of 15 meters, and a voice call between the client and the AP is emulated. We use the DITG traffic generator [143] to generate voice packets at 50 packets per second emulating a G.729.2 codec. We use RRAA as the underlying algorithm.

Second, in the above setup, we make the client mobile, and setup the voice call again. The client starts at a distance 40 meters, is LOS from the AP, and moves towards the AP to within a meter. The client is initially stationary and starts moving 20 seconds after the start of the experiment. We use SampleRate as the underlying algorithm.

Figure 3.8(a) and (b) show the frame lossrate and the bitrate chosen by RRAA when the client is stationary and NLOS. The graphs show that RRAA ends up incurring greater 802.11 frame loss and hence increased number of retries because of not *learning* the fact that 54Mbps is not appropriate rates for the setup. Increased frame loss leads to reduced overall network throughput. Figure 3.9 shows that SampleRate takes a long time to ramp-up to the chosen bitrate, when the client is mobile. Such conservative bitrate selection leads to inefficient channel usage and reduces overall network throughput.

In summary, a majority of existing bit-rate adaptation algorithms lack mechanisms to distinguish between interference and low SNR conditions. While more recent efforts have managed

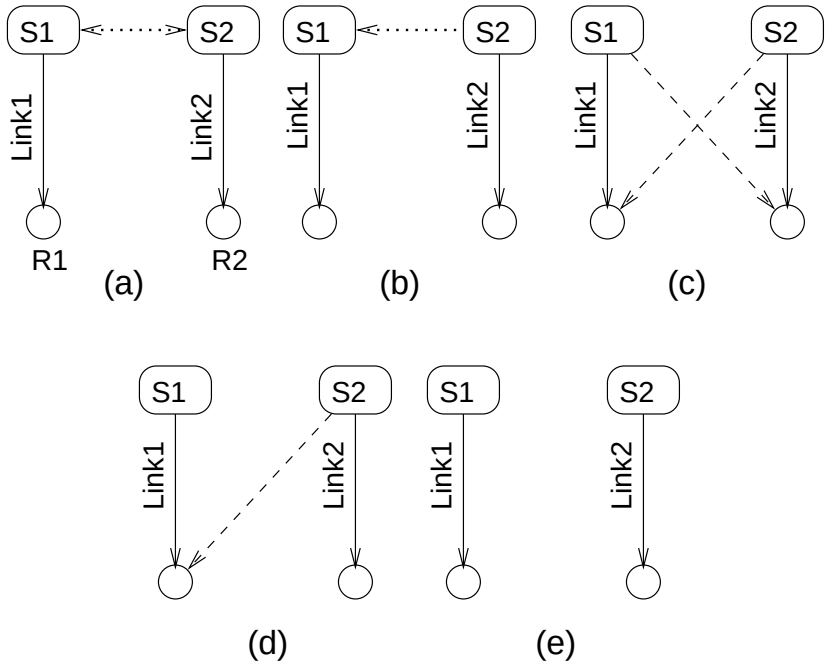

Figure 3.10: Problems introduced by power control: Scenarios of interaction between two links.

to fix this issue through the use of RTS/CTS frames, they are not robust to mobility and fail to converge in its presence. We next look at why agility and robustness is important for adaptive transmit power control in these networks.

3.4 Transmit Power: Agility and Robustness to Interference

As users increasingly make WLANs the first choice for last-mile network access, both spatial reuse and battery life are crucial metrics for better user experience. With emerging mobile applications leading to increased data transfer over WiFi interfaces, and hardware and protocol improvements reducing the idle-time power consumption of these interfaces, transmit power becomes the dominating factor influencing battery lifetime. Secondly, with increasingly dense deployments of WLANs for continuous coverage to users, mitigating interference to maximize spatial reuse is a crucial design goal. Adaptive transmit power control on a per-link basis promises to improve both the above metrics. While per-link power control is beneficial, performing it can be challenging due to several reasons. In what follows, we first look at TPC's interaction with interference and CSMA/CA scheduling followed by its interaction with rate adaptation, and the issues that need to be addressed in the presence of user mobility.

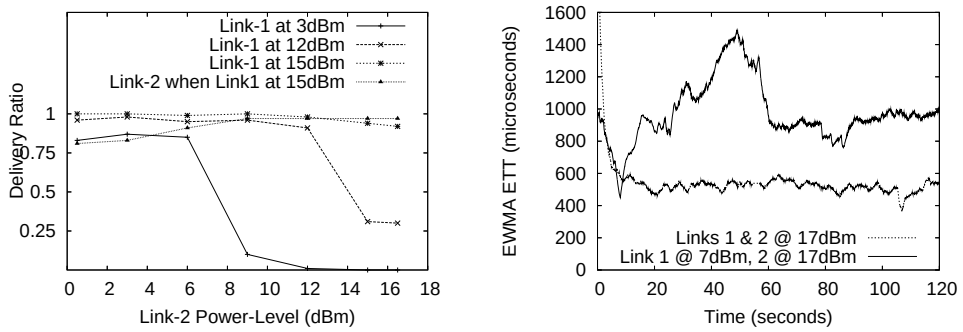

(a) (b)

Figure 3.11: Problems introduced by power control: (a) Receiver-side interference, and (b) Asymmetric channel access.

Transmit power control can introduce link asymmetry that leads to two problems: Receiver-side interference and Asymmetric channel access. Several previous works have already observed these two problems with power control [33, 73, 76, 83]. To provide a more quantitative characterization of (1) their *effect on performance in realistic settings* and (2) their *likelihood of occurrence*, consider a canonical network of four nodes—two senders and two corresponding receivers using the same 802.11 channel. Figure 3.10 identifies the different scenarios of interaction when the two links use different transmit powers. Dotted arrows indicate that the senders are in carrier-sense range; in (b) S1 can hear S2, but not vice versa. Dashed arrows indicate unintended interference at the receivers.

Scenario (a) represents fair channel sharing since S1 and S2 can carrier-sense each other. Scenario (b) represents the case of asymmetric channel access; whenever S2 has data to transmit, S1 does not get a fair chance to transmit. Scenarios (c) and (d) are two instances of receiver-side interference; while the senders are oblivious of each others' presence, packets sent by them collide at their receivers. Finally, scenario (e) is the ideal case of no interference and simultaneous transmissions on each link. Scenarios (b), (c) and (d) for any two links in the network can degrade the link and network throughput and fairness.

To demonstrate the effect of receiver-side interference, we set up four laptops—two senders and two corresponding receivers, with receiver 1 in between both senders. The laptops have Atheros PCMCIA wireless cards. We start backlogged UDP transfers at a fixed data rate of 54 Mbps from each sender to its receiver. We plot in Figure 3.11(a), the *delivery ratio* observed by the two links when the transmission power of each link is changed. Delivery ratio

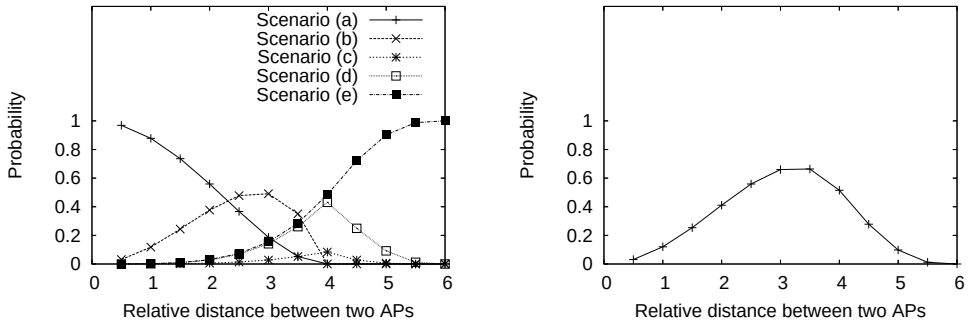

(i) All Scenarios (ii) Aggregate Bad Scenarios

Figure 3.12: (i) Probability of scenarios (a)-(e), (ii) Total probability of bad scenarios (scenarios (b), (c) and (d)).

is calculated as the ratio of packets successfully received at the client to those sent out by the AP. In Figure 3.11(a), the line "Link-1 at 3dBm" shows that Link-1 observes a good delivery ratio when Link-2 is also at a relatively low power level (Scenario (e)), but the delivery ratio decreases when Link-2's power increases due to receiver side interference (Scenario (d)). The lines "Link-1 at 15 dBm" and "Link-2 when Link-1 at 15dBm" show that when Link-2 is also sending at high power level ($>$ 12 dBm), both the links share the channel as in Scenario (a).

To demonstrate the effect of asymmetric channel access, we consider the same setup, but with receivers moved away from the other senders, so that they cannot be in scenarios (c) and (d). We now plot in Figure 3.11(b) the expected transmission time (ETT) [46, 92] of a packet on Link-1, with the minor modification that we only consider packets that succeed without any retries. Our intention is to capture the *channel access delay* for each frame through ETT measurement. When links are asymmetric (scenario (b)), Link-1 has higher channel access delay and hence higher ETT than when both links are symmetric (scenario (a)).

We now address the question, how frequently do the problematic scenarios ((b),(c) and (d)) occur? Since an experimental approach cannot sufficiently answer this question, we take an analytical approach. We derive the probability that each of the scenarios occur in a random geometrical graph with four nodes as above when the senders employ transmit power control. For brevity, we present the detailed mathematical formulation and analysis in a technical report [144], and just state the results here. Figure 3.12(i) shows the probability of occurrence of each scenario with distance between the senders, where the distance is shown as a factor of the communication range of the senders. Figure 3.12(ii) shows the sum of probabilities of all

problematic scenarios. The graph clearly shows that these scenarios can occur very frequently in a real deployment. Detecting and avoiding these problems in mobile environments is even more challenging since such scenarios can be dynamically introduced for short periods of time.

3.5 Transmit Power: Interaction with Bit-Rate

In WLANs based on 802.11a/b/g technology, senders use one of multiple transmission rates for sending packets. The choice of the rate is determined by an estimate of the channel conditions either through packet loss [43–45], delivery ratio [47], throughput [46], or Signal to Interference and Noise Ratio (SINR) estimates [1, 49]. Conceptually, a link is expected to perform well at a chosen rate if the SINR at the receiver is above a threshold (Table 3.4). Rate selection and transmit power control are tied together; power control without considering rate can reduce the SINR, leading to reduction in rate and hence the link and network throughput. In this paper, we take a system perspective and choose a (minimum) power level for a link that does not compromise the achievable rate. From the table, it can be seen that for supporting 54Mbps, the transmit power can be reduced such that the SINR is close to 24.56dB. Similarly, say, if 54Mbps and 48Mbps cannot be supported even at maximum allowed transmit power, then power can be reduced such that the SINR is close to 19dB to operate at 36Mbps.

While such rate and power selection is easily realizable with precise knowledge of SINR at receivers [1], reliable SINR measurements and reports in the presence of mobility cannot be achieved at fine timescales due to their overhead. Consequently, we rely on *estimating* the channel conditions based on the delivery ratio of a window of packets, similar to past works [47]. Such an approach, however, makes rate and power selection non-trivial. To illustrate, consider Figure 3.13. If the link is in a state of rate and power allocation (r_j, p_k) at a given instant, and the delivery ratio deteriorates (negative feedback), the reaction can either be to reduce rate or increase power. While increasing power appears to be a natural choice (as in PARF [30]), it is

SINR Range	Rate	SINR Range	Rate
$\geq$ 24.56	54	[10.79,17.04)	18
[24.05,24.56)	48	[9.03,10.79)	12
[18.8,24.05)	36	[7.78,9.03)	9
[17.04,18.8)	24	[6.02,7.78)	6

Table 3.4: SINR (in dB) vs. Rate (in Mbps) for BERs $\leq 10^{-5}$ in 802.11a [1].

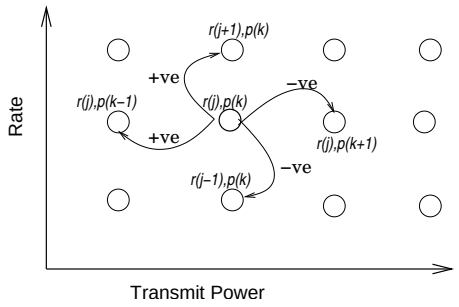

Figure 3.13: The dilemma of rate and power adaptation.

possible that even at the maximum power, the current rate cannot be supported, in which case reducing rate is the right choice. Lack of knowledge of whether a rate can be supported by increasing power to the maximum, can increase the convergence time, which is prohibitive in the presence of mobility. Similar dilemma exists on positive feedback.

3.6 Transmit Power: Agility and Robustness to Mobility

Adaptive transmit power control becomes more challenging in the presence of user mobility. Link conditions change frequently due to distance-based path loss, short-lived hidden terminals and occasionally severe destructive multipath interference at certain locations in a user's path. In such environments, both rate and power control algorithms need to address the following questions effectively: *How frequently* should the adaptation take place? and *At what granularity* should rate and power be adapted?

For the first question, solutions have to strike the right balance between reliability and responsiveness: waiting for enough samples avoids reacting to short-lived drops in link conditions, while waiting too long can also be detrimental to performance. For the second question, changing power at coarse granularity allows adapting less frequently, but compromises on battery life and spatial reuse; whereas fine granularity changes require frequent adaptation.

The exact answers to these two questions mainly depend on the speed of the user, which is 0.5-1m/s in typical WLAN environments. To understand the solution requirements in such scenarios, we study the effect of mobility on the received signal strength (RSS) of a client, understanding which will shed light on how to tune transmit power and rate. In this experiment, we use one AP that is stationary, and one client running a voice call (generating 50 packets

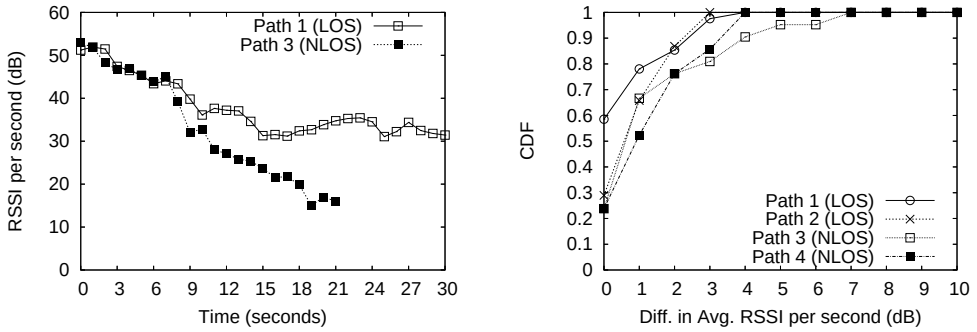

(a) (b)

Figure 3.14: Effect of mobility (at about 0.75 m/s) on RSSI.

per second) with the AP. The client moves away from the AP at approx. 0.75m/s along four different paths in our office building. The client is in line-of-sight (LOS) on two of the paths, and non-line-of-sight (NLOS) on the others. Figure 3.14 (a) shows the average RSSI per second with time for one LOS and one NLOS case. In Figure 3.14 (b), we show the CDF of the difference in avg. RSSI per second. We observe that 95% of the time, the avg. RSSI in one second will change by at most 5dB.

In summary, mechanisms are needed to detect channel access asymmetry and hidden nodes, and whether adaptive TPC caused these inefficient states. These mechanisms also need to keep in mind that, due to user mobility, such inefficiencies can occur at small time scales. Finally, the experimental results in this section can be leveraged to set the granularity and frequency with which TPC algorithms adapt.

We now turn our attention to leveraging directionality in outdoor wireless networks. In particular, we focus again on balancing agility and reliability in the presence of high-speed or vehicular mobility.

3.7 Directionality: Agility and Robustness to Mobility

As we stated earlier, we believe that directionality, base station diversity (or diversity) and bit-rate are critical resource parameters for future outdoor wireless networks. From Chapter 2, we know that directionality represents the idea of forming a beam towards a receiver (i.e. direct the transmitted energy in an intended direction) in order to increase the *mean* signal-to-noise-ratio (SNR) (pictorially represented by beam B1 in Figure 3.15). Meanwhile, diversity

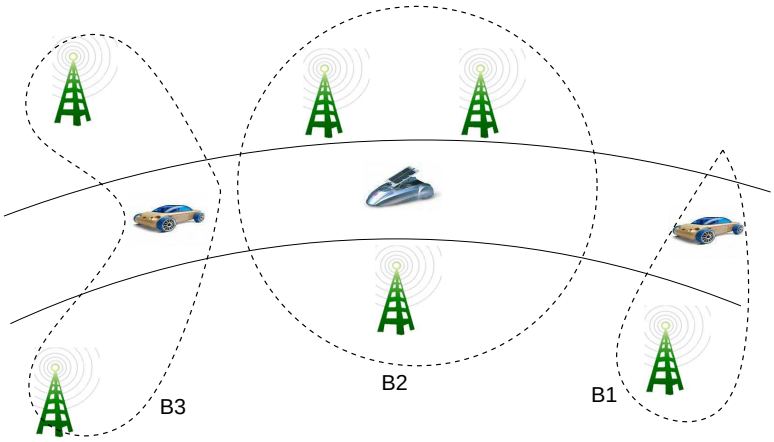

Figure 3.15: Different beam configurations for communication between clients and the receivers. B1 uses directionality with a single receiver, B2 uses diversity with all visible receivers, and B3 uses a combination of directionality and diversity to a subset of the visible receivers.

represents the idea of making omni-directional transmissions so as to combine packets received at multiple receivers in the vicinity of the transmitter, thereby masking off packet losses due to SNR *variance* or deep fades at any individual receiver. This method is pictorially represented by beam B2 in Figure 3.15. Both mechanisms have individually been included in several wireless standards [106–108] and research proposals [93, 109].

In this dissertation, we argue that in several locations that a mobile client traverses, a *combination* of directionality and diversity is more appropriate than using them in isolation; i.e. a beam that covers more than one but not all of the visible receivers is more appropriate for maximizing the uplink bandwidth (as depicted by B3 in Figure 3.15). To support the hypothesis, we perform an experiment (details explained in subsequent sections) with a real deployment of 4 road-side WiFi APs and a moving vehicle uploading data. We measure the effectiveness of (i) choosing one AP at a time as a receiver and forming a beam pointed to it (represented by C1), (ii) all subsets of two receivers and forming a multi-lobe beam to them (C2), (iii) all subsets of three receivers (C3), and (iv) making omni directional transmissions to emulate complete diversity (C4). Figure 3.16 shows the percentage of locations where a combination gives the best throughput. We observe that at significant number of locations, either C2 or C3 is a better choice than the extremes: C1 or C4. Further, bitrate has an influence on which choice is the best. Note that although we use WiFi devices in this paper, the observations themselves are technology-independent.

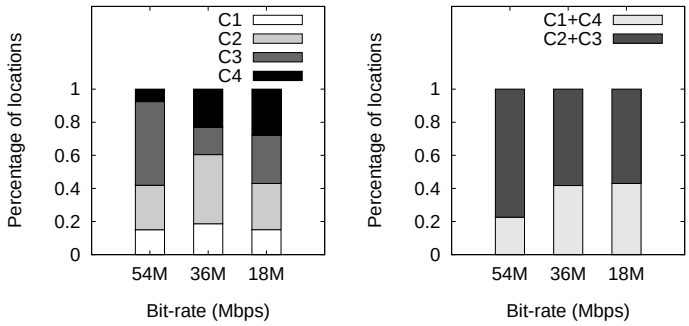

Figure 3.16: Tradeoff between directionality and diversity. A combination of directionality and diversity (C2 or C3) achieves lower PER at a majority of locations relative to vanilla directionality (C1) and diversity (C4).

Identifying the right operating point between directionality and diversity, however, is non-trivial due to conflicting parameter settings required to achieve the best of each mechanism: directionality stipulates using as narrow a beam as possible for higher gains, whereas diversity requires the wider beams to include as many receivers as possible. Further, high-speed mobility and real-world situations such as shadowing and lack of line-of-sight (LOS) that influence link quality require *agile* and *robust* solutions that choose between these choices in an informed manner at run-time.

In the following section, we first analytically explore the directionality-diversity tradeoff, and identify the impact of different parameters to guide the design of our solution. We also demonstrate that bit-rate adaptation is a crucial consideration that makes striking the tradeoff even more involved.

3.7.1 Directionality vs Diversity Tradeoff

We use outage probability to model the resulting throughput on a link. Outage probability refers to how often (probability) does the bit error rate (BER), or equivalently SINR, experienced falls below a certain threshold. It is both a popular and practical measure for robustness to fading, especially for block fading where it can directly be related to frame/packet error rate. It can be measured by determining the probability that the mutual information of communication (capacity) is less than the information rate. We consider an independent, quasi-static, frequency non-selective Rayleigh fading (complex channel coefficients being uncorrelated and circularly symmetric Gaussian random variables with zero mean and unit variance) along with free space

path loss for our channel model. Let a channel realization between an n-element transmitter and an omni receiver be denoted by h_{tr}. Now, the mutual information for a given channel realization in the case of beamforming is given by the asymptotic Shannon capacity formula,

$$C(h_{tr}) = \log_2(1 + n\rho|h_{tr}|^2))$$

where ρ is the average receiver SNR at Rx for an omni transmission. If $\tilde{R}$ is the information (data) rate (bits/s) applied to the system (normalized to bandwidth), then the outage probability can be given as,

$$P_o(\tilde{R}, n) = Pr[C(h_{tr}) < \tilde{R}] = Pr\left[|h_{tr}|^2 < \frac{2^{\tilde{R}} - 1}{n\rho}\right]$$

By way of definition of $|h_{tr}|$, $|h_{tr}|^2$ follows an exponential distribution. Hence, on averaging over all possible channel realizations, we have,

$$P_o(\tilde{R}, n) = 1 - e^{-\left(\frac{2^{\tilde{R}}-1}{n\rho}\right)} = 1 - e^{-\left(\frac{S}{n\rho}\right)}$$

where $S = 2^{\tilde{R}} - 1$. When a beam of width $\frac{360^o}{n}$ is formed with n elements, let the number of receivers falling in the reception zone of the beam be given by the function $\ell(n)$. Note that, $\ell(n)$ is a monotonically decreasing function, with number of accessible receivers decreasing with finer beamwidths. Now, the resulting probability of successfully receiving the packet/frame can be given as,

$$P_s(\tilde{R}, n) = 1 - \prod_{i \in \ell(n)} \left(1 - e^{-\left(\frac{S}{n\rho_i}\right)}\right)$$

where ρ_i is the average SNR at receiver i corresponding to an omni transmission. The resulting throughput can be given as,

$$T(\tilde{R}, n) = \tilde{R} \cdot \left(1 - \prod_{i \in \ell(n)} \left(1 - e^{-\left(\frac{S}{n\rho_i}\right)}\right)\right)$$

The above equation captures the tradeoff between diversity and directionality. With larger n, smaller beamwidths can be formed. This results in higher directionality or array gain that improves the individual link success probability. However, the decreased beamwdith reduces the set of accessible receivers $\ell(n)$ and hence the diversity combining ability, thereby reducing the collective success probability across multiple receivers.

To illustrate this tradeoff, we quantitatively evaluate throughput as a function of n, ρ and $\tilde{R}$ under simplifying assumptions. We assume $\rho_i = \rho,\ \forall i \in \ell(n)$, and $\ell(n) = \frac{8}{n}$ such that

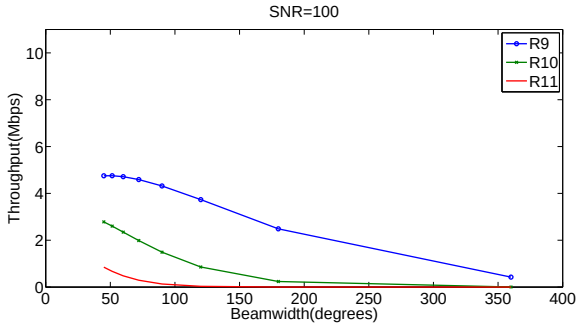

(a)

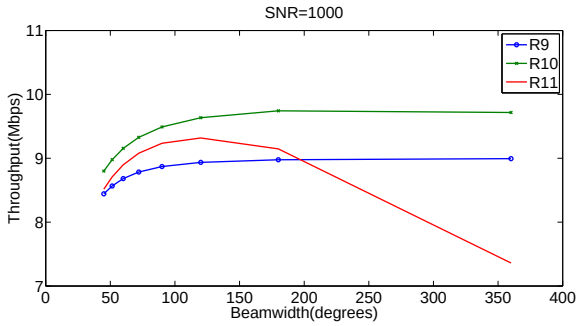

(b)

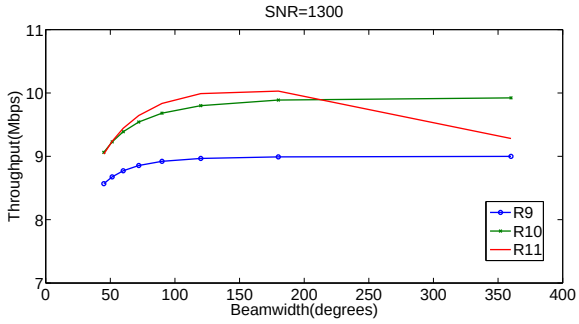

(c)

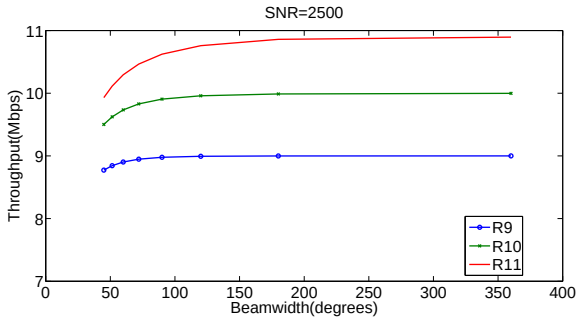

(d)

Figure 3.17: Tradeoff Illustration.

there are eight receivers reachable with omnidirectional transmission, and they decrease proportionally with increasing antenna elements (i.e. decreasing beamwidth). We allow fractional values of $\ell(n)$ for analytical simplicity. Figure 3.17 shows the results as an increasing function of average SNR across the graphs. Within each graph, we examine the behavior of different rates for a given average SNR with increasing beamwidths on x-axis. Increasing the number of antenna elements n increases directionality (beamwidth of $\frac{360^o}{n}$), with $n = 1$ representing pure diversity from an omni-directional transmission, and $n = 8$ representing pure directionality with an eight element array, with $1 < n < 8$ corresponding to combinations of diversity and directionality. The results reveal the complex interaction between diversity, directionality and rate. In particular, we make three key observations:

Observation 1: Figure 3.17(a) shows that at low average SNR, lower beamdwidths are needed (i.e. increased directionality) to increase throughput, since making the gain on at least one link is better. Whereas, Figure 3.17(d) shows that at high SNRs, increased directionality does not fetch as much benefit as diversity does; diversity alleviates packet drops due to deep fades that directionality can not handle.

When the average SNR is in the moderate region, as in Figures 3.17(b) and 3.17(c), there exists a clear tradeoff between directionality and diversity that results in a mixed directionality-diversity strategy being the optimal. This is demonstrated by the peak being somewhere in the middle for a given rate.

Observation 2: Now, let us examine the behavior across rates for a given average SNR. Note that for a given average SNR, with increasing transmission rates, the packet error rate (PER) increases. Thus, the different rate curves can be considered to correspond to different PERs. It can be seen in Figure 3.17(a) that when the average SNR is low, higher rates cause large PER. Hence it is important to keep the rate low and use maximum directionality to reduce PER. However, when the SNR is very high it is best to operate at the highest rate using maximum diversity as in (d). For moderate SNRs, the rate curves exhibit a tradeoff as identified before. However, picking the optimal rate now requires more care.

Observation 3: Finally, Figures 3.17(b) and 3.17(d) show that for the best rate in the setting, the curves flatten with increasing beamwidth showing that it is sufficient to choose a low enough beamwidth (i.e. operate at the knee of the curve) for maximum benefit and not operate at maximum diversity. Operating at lower beamwidths increases the opportunities for *spatial reuse*, and hence it is always desirable to operate at as low beamwidths as possible.

3.7.2 Design Considerations

In summary, the following issues have to be considered for developing a solution to jointly tune directionality, diversity and bit rate for maximizing the uplink throughput of a client. These options have conflicting parameter settings, hence making the problem non-trivial:

- For best diversity gains: the larger the number of receivers, the better are the benefits of diversity. Covering larger number of receivers naturally engenders using the widest possible beamwidth (omni-directional being the extreme) for transmission at the client. The larger the beamwidth, however, the lower is the gain in any direction, and hence the lower the SNR at a receiver.

- For best directionality gains: The thinner the beamwidth used by the client, the higher is the signal gain, and hence the higher is the benefit of directionality. Increased SNR on the link also allows using higher bit rates for transmissions. However, this also implies that using highly directional beams reduces the number of receivers available for obtaining diversity gains.

- For best bit rate gains: The higher the bit rate, the better are the throughput gains on a link . For achieving high bit rate, however, the average SNR on a link should be high enough to cross a certain threshold. Consequently, among a set of receivers that are used for diversity, the receiver with minimum SNR controls the bit rate at which packets should be sent to be successfully decodable at all the receivers.

Further, a complete solution has to address two deployment-specific challenges:

- Short timescales at locations: Due to high-speed mobility of user devices (that is the focus of this work), the amount of time spent at a location is relatively short, and the number of parameters that need to be explored for finding the best combination is relatively large. For instance, a car traveling at 50 MPH is in a 10 x 10 meter location for about 500 ms. Correspondingly, with an eight element antenna, there are tens of beam angles, several receivers depending on the location, and at least nine bit rates to choose from (assuming 802.11g). To adjust to the high mobility, a solution has to rapidly converge to the right parameter settings.

- As wireless technologies evolve, clients are bound to have different capabilities for forming beams, i.e. number of elements, use different transmit powers, etc. Hence, the solution should not be completely tied to a specific client configuration. For instance, which direction a beam may be formed towards a receiver can be decoupled from how exactly we form the beam, or how many elements are used to do so.

We now proceed to describe two resource management frameworks, *Symphony*, and *Sonata*, for indoor and outdoor environments respectively that leverage the insights gained so far.

Chapter 4

The Symphony Framework: Design and Implementation

This chapter describes the design and implementation of Symphony, a synchronous two-phase rate and power control framework for indoor CSMA wireless networks. Symphony aims to increase the battery life of mobile devices and improve spatial reuse, while addressing the challenges of undesirable rate adaptation, receiver side interference, asymmetric channel access and the effects of user mobility.

4.1 Overview

Due to the possible adverse effects of power control, the goal of Symphony is to tune the transmit power and rate of each link in a WLAN such that *the link's performance is at least as good as in the baseline maximum-power network.* At the core of Symphony is a synchronous two-phase execution (Figure 4.1) strategy, in which all nodes (APs and clients) in the WLAN cycle through two phases in synchrony—the REFERENCE (REF) phase and the OPERATIONAL (OPT) phase. In the REF phase, Symphony estimates for each link the best achievable performance, and in the OPT phase, it tunes the link to the lowest transmit power to achieve the same performance as in the REF phase. Due to mobility (of users or the environment) the best attainable performance may continuously change, and a power/rate setting may suddenly be affected by asymmetry. The reference phase provides a convenient solution to periodically verify that power and rate control have not unnecessarily degraded system performance.

Similar to most rate control algorithms, Symphony executes at the sender side on each unidirectional $(sender, receiver)$ link. The sender can be either an AP or a client, and the receiver a client or an AP. In the REF phase, each sender performs rate adaptation for each link at the maximum power to choose the best data rate for the current channel conditions. In the OPT phase, the sender performs both rate and power adaptation. The rate and power adaptation algorithms maintain two contexts—ref_ctxt and opt_ctxt, one for each phase for each link.

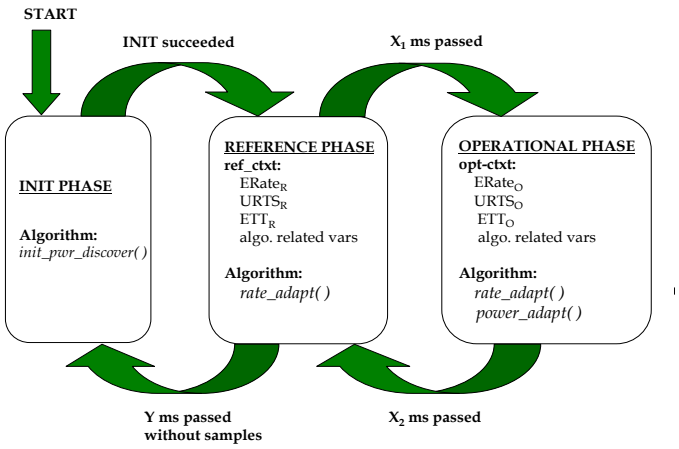

Figure 4.1: Symphony's two-phase synchronous strategy.

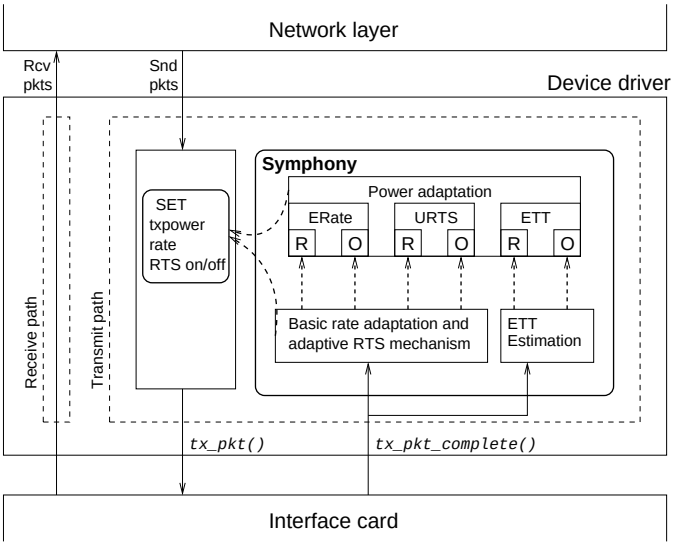

Figure 4.2: Architecture of Symphony. The blocks R and O represent REF and OPT contexts.

Each context contains several performance metrics and other variables needed for executing the rate and power adaptation algorithms. We choose three metrics—EWMA (Exponential Weighted Moving Average) of data rate, utility of RTS and EWMA of ETT—to help detect and avoid the problems outlined in Section 3.4. The performance metrics in the `ref_ctxt` serve as reference values for the OPT phase. In the OPT phase, each link is tuned to the lowest power such that each performance metric in the `opt_ctxt` is *no worse* than the corresponding metric in the REF phase by more than a threshold.

Before entering the REF-OPT cycle, each link goes through an INIT phase, in which the sender starts from the minimum power level and rapidly discovers the initial power level necessary for communication with the receiver. To avoid impacting applications, we use probe packets at the highest rate and additively increase power for each packet till a probe packet succeeds in reaching the receiver. If we reach the maximum power and still do not succeed, we lower the rate and start over at the lowest power. After succeeding, the sender initializes the OPT phase with the successful power level, and enters the REF phase with the successful rate at the appropriate synchronized time. If the sender is idle for a threshold number of seconds (=2 in our prototype), and then a packet from the network layer arrives that does not succeed in reaching the receiver, we determine that the rate and power information is stale (e.g., due to mobility) and reset the sender to the INIT phase to repeat the rapid discovery process.

We achieve synchronized phase execution on all APs and clients in two steps. First, the APs are synchronized to a global real-time clock by a central controller[1]. The controller configures the lengths of the two phases on each AP, and specifies at what real-time the phases should start executing. Second, for each phase change, each AP broadcasts a message (at maximum power) informing the change to the clients, and the clients switch phase. These broadcast messages are sent at high priority to ensure minimum skew across nodes; in our prototype, we use the hardware queue reserved for voice traffic in the Atheros cards.

We implement Symphony in the Linux MadWifi driver 0.9.3.1 [45]. Figure 4.2 shows the architecture of Symphony. As shown, Symphony executes in the transmit path. We represent rate adaptation as a separate block to make Symphony extensible. Any rate adaptation algorithm can fit into Symphony as long as it executes in the two contexts and provides rate information for power adaptation. Similarly, different mechanisms can be implemented for ETT estimation and determination of RTS utility. In what follows, we describe the important components of Symphony.

4.2 Bit-Rate adaptation

Based on our experiments in Section 3.2, we propose the following enhancements to the state-of-the-art RRAA algorithm before incorporating it as part of our framework. We call the new algorithm ***RRAA+***.

[1]The controller and thin-AP architecture is the most common way WLANs are built today.

Algorithm 2 RRAA+

```
1:  if (loss > HT_k) then
2:      p[k] /= α_1
3:      k ← next_lower_rate
4:  else if (loss < LT_k) then
5:      for (all rates j ≤ k) do
6:          p[j] *= α_2
7:      end for
8:      if (rand() < p[nexthighk]) then
9:          k ← next_higher_rate
10:     end if
11: end if
```

4.2.1 The Probabilistic Rate Increase (PRI) mechanism

To make the RRAA algorithm converge, we propose adding a probabilistic rate increase (PRI) mechanism to the algorithm as shown in Algorithm 2. In brief, RRAA+ maintains for each bitrate, the probability that it transitions to this bitrate from the next lower rate. Every time the loss at a bitrate exceeds the high threshold, the probability of returning to this bitrate is reduced before transitioning to the next lower rate. The MIMD parameters $\alpha 1$ and $\alpha 2$ are chosen such that the algorithm becomes stable. In our prototype, $\alpha 1 = 2$ and $\alpha 2 = 1.0905$; it takes 8 increments to match one decrement.

4.2.2 Use of both time and number of samples

Another drawback of RRAA that we rectify is that it proposes the use of a window of frames, before each rate control decision. This approach is dependent on application traffic characteristics. Thus, for VOIP packets, we observe that RRAA will wait for up to 800ms (40 frame window at 54Mbps) before deciding to switch the rate. A more elegant approach to this problem is to wait for a certain number of packets, upper bounded by a duration of time, which is what we adopt in our implementation. This duration of time has to be chosen well since the metric calculated over a particular window should have enough samples to be reasonably accurate.

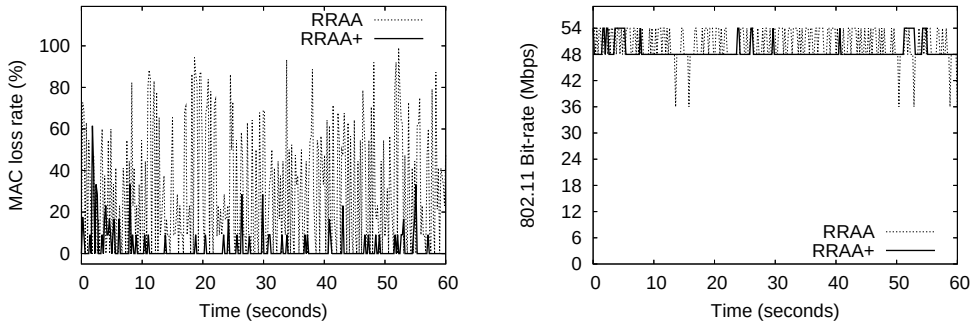

(a) RRAA vs RRAA+ (MAC loss rate) (b) RRAA vs RRAA+ (MAC bit-rate)

Figure 4.3: Comparing convergence characteristics of RRAA+ with RRAA.

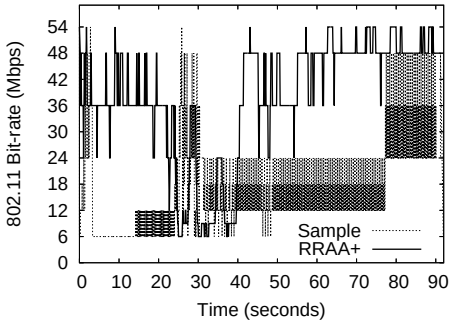

Figure 4.4: Performance with user mobility: SampleRate vs RRAA+

4.2.3 Implementation challenges

Before we could proceed to implement these enhancements, we had to first address non-trivial implementation challenges related to RRAA's assumption that the underlying platform provides *per-frame* control and feedback. The original RRAA is implemented in the firmware of a programmable AP platform, which provides it with the ability to immediately control, and get feedback on, each frame transmission. However, a number of existing hardware implementations [145–147], queue a group of *packets* for transmission and have the option of getting feedback on a *per-packet* basis rather than a *per-frame* basis. We overcame this challenge by dealing in units of packets rather than frames. In our implementation of RRAA and RRAA+ in the MadWifi driver, we utilize the window sizes specified by [47], upper bounded by a duration of 200ms, which, even with VOIP, should provide us with up to 110 samples (10 packets with 11 retries each).

We repeat the experiments outlined in section 3.3.2 with RRAA+. Figure 4.3(a) and (b) show the frame lossrate and the bitrate chosen by RRAA and RRAA+ in the static, NLOS experiment. Unlike RRAA, RRAA+ recognizes that 54Mbps is unstable and settles at 48Mbps. This is also reflected in the reduced MAC loss rate, which is near-zero for a major duration of time. Figure 4.4 shows that, in the mobility experiment, RRAA+ is able to ramp-up its bit-rate much faster than SampleRate.

In summary, we choose RRAA+ for rate adaptation in Symphony due to its three features—agility, convergence to appropriate rate and avoidance of rate adaptation because of collision-induced packet losses.

4.3 Transmit Power control

Our goal for power adaptation is to tune each (sender, receiver) link in a WLAN to the lowest transmit power such that the performance metrics in the OPT phase are no worse than the corresponding metrics in the REF phase. Algorithm 3 shows the basic behavior of power adaptation in Symphony. The three conditions in line 1 detect undesirable rate adaptation, receiver side interference and asymmetric channel access introduced by power control. Similar to RRAA+, the power control algorithm learns the lowest appropriate power level by maintaining the probability with which it should transition to a particular level. The algorithm executes once for every two intervals of the rate adaptation algorithm to adapt to user mobility. Further, several rate adaptation intervals can occur in each of the REF and OPT phases, depending on traffic.

4.3.1 Preventing undesirable rate adaptation

For detecting and preventing undesirable rate adaptation due to power control, for each link, the two contexts maintain an EWMA of the rate chosen by the rate control algorithm in response to the measured packet loss: for each $rate_i$ chosen in interval i, we set $ERate = ERate * \psi + rate_i * (1 - \psi)$ at the end of interval i. Every time the power control algorithm is triggered, if the EWMA of rate in OPT phase ($ERate_O$) is lower than that in the REF phase ($ERate_R$) by a threshold τ_1, transmit power is increased. In our implementation, the EWMA parameter ψ is chosen as 0.8. We choose τ_1 to be 3 Mbps if $ERate_R$ is above 48Mbps or below 24 Mbps, and 6 Mbps otherwise. We make this choice because of the non-uniformity in 802.11 a/g bit rate granularity. Our idea is to place the threshold between the two consecutive rates.

Algorithm 3 power_control

1: **if** (($ERate_R - ERate_O > \tau_1$) OR
($URTS_O > URTS_R$) OR
($ETT_O - ETT_R > \tau_2$)) **then**
2: p[curpwr] /= β_1
3: curpwr ← next_higher_pwr
4: **else**
5: **for** ($all\ pwr\ levels\ i \geq curpwr$) **do**
6: p[i] *= β_2
7: **end for**
8: **if** ($rand() < p[next\_lower\_pwr]$) **then**
9: curpwr ← next_lower_pwr
10: **end if**
11: **end if**

4.3.2 Detecting hidden terminals

To detect that power control introduces receiver side interference we utilize the adaptive RTS/CTS mechanism. Similar to RRAA [47], RRAA+ includes a mechanism to detect if packet losses are happening due to collisions as opposed to degraded channel conditions. Our implementation of the adaptive RTS/CTS mechanism, however, differs significantly from RRAA; while RRAA was implemented on a per-frame basis (because of the availability of card firmware), we implement it on the basis of a window of packets, both for more reliable estimation of receiver side interference, and for obviating the need of modifying the firmware. In the interest of brevity, we elaborate the implementation details of this mechanism for Atheros cards in a technical report [144].

Using this mechanism, we maintain a performance metric—the utility of RTS (URTS)—that is set to one if the loss rate with RTS/CTS is less than the total loss rate in at least 2 out of 4 last rate adaptation intervals, i.e. enabling RTS/CTS is helpful to reduce losses. Otherwise, URTS is set to zero. The rationale for waiting for 4 intervals is to determine the utility with greater reliability, while trading off responsiveness; further, unless the receiver interference problem is sustained, we do not increase transmit power and let adaptive RTS address the problem. Now, if URTS is 0 in the REF phase and 1 in the OPT phase, it indicates that power control introduced the receiver side interference that didn't exist in the REF phase. Line 1 in Algorithm 3 captures this condition and triggers a power increase.

Algorithm 4 ETT Estimation

```
1:  VARIABLES mark_seqno, itime, etime
2:
3:  FUNCTION driver_send_pkt (seqno)
4:  if (mark_seqno not set) then
5:      mark_seqno = seqno
6:      etime = curtime
7:  end if
8:
9:  FUNCTION card_sent_pkt (seqno)
10: if (mark_seqno = seqno) then
11:     ETT = curtime - MAX (itime, etime)
12:     itime = curtime
13:     unset mark_seqno
14: else
15:     ETT = curtime - itime
16:     itime = curtime
17: end if
```

4.3.3 Preventing channel access asymmetry

To detect that power control introduces asymmetric channel access, we measure the EWMA of the expected transmission time (ETT) of each packet. The key idea here is that if a sender does not get a chance to transmit as frequently due to asymmetry in the OPT phase, the ETT in the OPT phase increases compared to the REF phase. If the ETT increases by more than a threshold τ_2, we trigger power increase (as in Algorithm 3). In our implementation, τ_2 = 100μs.

If the wireless card provides the device driver with the transmission time for each packet, EWMA of ETT can be easily calculated. However, a majority of today's WiFi cards (including Atheros cards) currently do not provide this information. Further, multiple packets can be queued by the driver in the buffer of the card for efficiency, which makes ETT estimation non-trivial. In our implementation, we overcome the above problem with the two functions in Algorithm 4 that exploit the interaction between the device driver and the interface card. We use the unique sequence numbers in packets that are sent to the card, and keep one outstanding marked packet. The variable $etime$ represents the time a marked packet was sent to the card, and $itime$ represents implicit time when a packet's transmission actually started on the card. This method ensures that ETT can be estimated even on packets that are buffered back-to-back and hence are not explicitly marked. Further, we only consider packets that do not incur any retransmissions to reliably estimate the channel access delay, and consider packets sent at the

same rate as in REF phase to avoid false positives due to rate-based ETT changes.

Finally, since packet sizes affect the ETT because of the varied transmission time on the air, we maintain ETT in terms of 1500 byte packets. For smaller packets, the ETT is scaled to 1500 bytes before using it to calculate EWMA. To do so, we set for each packet $fixedETT = ETT + (1500 - pkt\_size)/rate$, where rate is the data rate used to transmit the packet. We validate experimentally that this approach of scaling ETT is reasonably accurate, and present the result in the technical report [144].

4.3.4 Implementation challenges

Granularity of power control: Learning from our observations in Figure 3.14, and given that power adaptation gets triggered at least twice in a second, in our implementation, Symphony increases and decreases power at a granularity of 3dB, between [MIN_PLEVEL, MAX_PLEVEL]. This ensures agility to typical user mobility in WLANs. Further, [148] observes that transmit power control at a finer-granularity than 3dB may not always be useful in indoor environments. The minimum and maximum transmit power values can be different on different 802.11 cards, vary with the frequencies used (such as in the 5GHz band) and also vary based on the gains of the external antennas connected to the cards. However, we assume that the levels can be discretized at the granularity of 3dB. In our prototype with the Atheros cards, we vary the power levels between 0 and 18 dBm.

Convergence: The process of increasing and decreasing power is similar to rate adaptation in RRAA+. Symphony maintains for each power level, the probability that it transitions to this level from the next higher level. Every time at least one of the conditions on the performance metrics satisfies, the probability of returning to this power level is reduced before transitioning to the next higher power. The MIMD parameters $\beta 1$ and $\beta 2$ are chosen to make the algorithm stable. In our prototype, $\beta_1 = 3$ and $\beta_2 = 1.14$; it takes 8 increments to match one decrement. Again, the choice of β_1 and β_2 strikes a tradeoff between the benefits of power control and stability of the algorithm; we arrive at the above values after experimenting with several scenarios.

With Symphony's approach of maintaining probability per power level, the transmit power of each sender will eventually converge to a point where the performance of each link is at least as good as in the REF phase. If the performance of a link at a given power level is similar

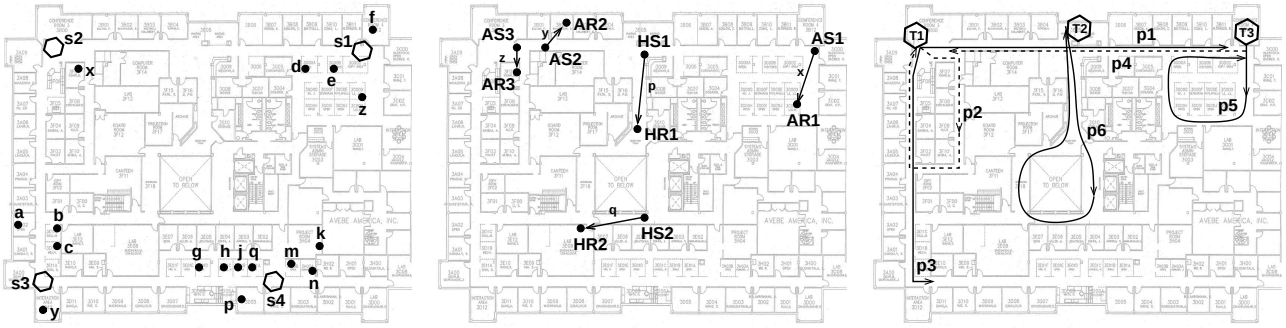

(a) Static experiments (b) Problematic scenarios (c) Mobility experiments

Figure 4.5: Experimental setup.

to or better than that of the REF phase the probability of returning to that power level and to any higher power level will converge to 1 due to the multiplicative increase of the probability. Otherwise, the probability of returning to that power level will converge to a small value due to the multiplicative decrease of the probability. In our implementation, we bound the probability on the lower side to $\frac{1}{64}$ to be responsive to changing channel conditions and mobility.

4.4 Evaluation

To demonstrate the achievement of the design goals outlined earlier, in this section, we carry out a systematic and extensive set of experiments, in both controlled and uncontrolled environments. In what follows, we first describe our experimental setup and then present our evaluation results.

4.4.1 Experiment Design

Figure 4.5 shows our setups for different experiments. We describe the relevant setups when presenting the specific experiments. We perform most experiments on 802.11a channels to avoid disturbing our office's WLAN that uses all three 802.11g channels (1,6,11). We perform a few experiments on 802.11g to emulate interactions with realistic WLAN scenarios. The clients are placed randomly in different office cubicles, and the APs are placed at locations where there are operational WLAN APs (on 802.11g channels) to provide a realistic WLAN environment for our experiments. However, for showing that Symphony addresses asymmetric channel access and receiver-side interference, we carefully choose the AP locations.

The APs and clients are Dell laptops with the Atheros PCMCIA cards, and the APs are

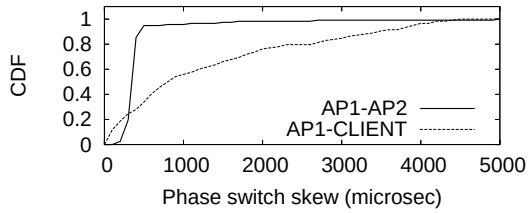

Figure 4.6: Skew between phases on two APs and an AP and a client as seen by an external monitor.

connected by a 100Mbps wired network. The Atheros cards transmit at a default (maximum) power level of 18dBm[2] for 54Mbps data rate, and allow Symphony to change transmit power at the desired granularity of 3 dB. In all experiments, the APs and clients run on the same 802.11 channel. While both clients and APs can execute Symphony, we perform most of our experiments with APs as senders and clients as receivers. This is because of the current limitation of Atheros cards that do not implement per-packet power control for ACK and CTS packets. To overcome this limitation for proper evaluation, unless specified otherwise, we run Symphony on APs, and disable per-packet transmit power control on the clients. Further, APs append transmit power information to each outgoing frame and we make a minor modification in the client driver to extract this information and set the card to this power level (similar to how `iwconfig athN txpower $val` sets power). This ensures that all frames, including ACKs are returned at the configured power.

For AP synchronization and configuration, instead of using a different central controller, we just use one AP as the master AP, and synchronize the others with the master through NTP on the wired network. For all the experiments in this section, we use 200ms and 800ms as the length of REF and OPT phase resp. All APs are configured to start the REF phase at the beginning of one second boundary.

4.4.2 Experiment Results and Analysis

In this section, we describe several experiments that demonstrate the efficacy of Symphony.

[2]Although the technical specs. specify a maximum transmit power of 15$\pm$2dBm, we observed that the card can use up to 18dBm.

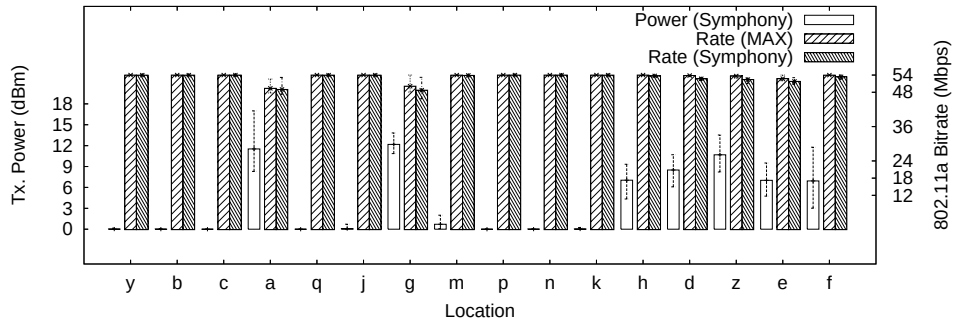

Figure 4.7: Transmit power reduction for clients in several locations.

Two-phase Synchronization

To test that Symphony indeed executes synchronously on APs and clients, we consider a network of two APs with one client each. In this experiment, both clients and APs run Symphony. We place a monitoring laptop close to the APs with its wireless card running in the monitor mode. We start two UDP transfers of 200 pkts/sec between each (AP, client) pair and collect packets on the monitor. The packets are appended with the transmit power used to send them and the executing phase for use at the monitor. Figure 4.6 plots the CDF of skew between two APs and one AP and its client as seen by the monitor. The graph shows that the nodes make corresponding phase transitions within 3 ms of each other more than 80% of the time.

Transmit Power Reduction

In this experiment, we consider the setup in Figure 4.5(a), with clients placed in different cubicles and office rooms (represented by lower case letters), and APs placed close to the operational WLAN APs (at s1, s3 and s4). We setup ten 1 minute VOIP calls for each client. From the associated APs, the calls start at different times (separated by 5 minutes) for each client. The white bars in Figure 4.7 show the average transmit power used by the Symphony APs for each client; in most typical user locations, the required transmit power can be substantially lower than the default 18dBm. Further, the error bars plot the minimum and maximum of the average transmit power per call, showing that the optimal power can vary with time for even static locations. The other two bars show the rate chosen by Symphony and the rate when transmitting at maximum power. Symphony causes minimum effect on the data rate chosen. *In this setup, for the three cells with S1, S3 and S4 as APs, a per-cell solution would operate all links*

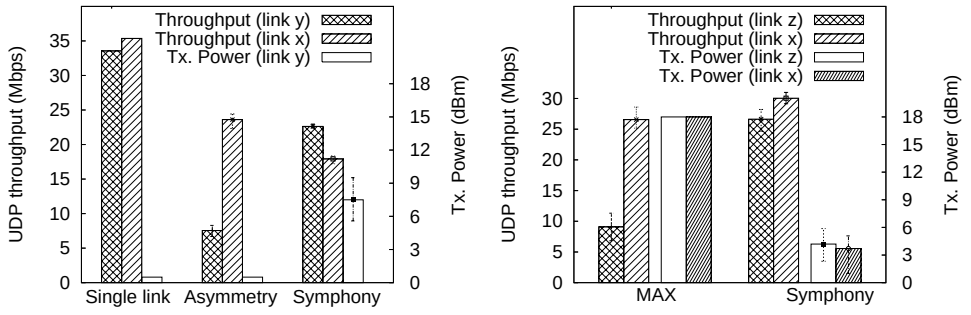

(a) (b)

Figure 4.8: Preventing asymmetric channel access. (a) shows Symphony's ability to detect and avoid channel access asymmetry, and (b) shows power control removing inherent link asymmetry.

at the worst client's transmit power, which is about 12 dBm. In contrast, Symphony enables 75-80% of the clients in the cells to settle at 3 to 12 dB (i.e. 50% to 94%) lower transmit power than 12 dBm.

Avoiding Channel Access Asymmetry

In this set of experiments, we demonstrate that Symphony avoids channel access asymmetry by intelligently adapting the transmit power of a link. In the first experiment, we consider links x and y as in Figure 4.5(b) (where AS1 and AS2 are senders and AR1 and AR2 are receivers), and setup backlogged UDP traffic to measure the link throughput. Link x always operates at the maximum power. We consider several cases—(i) links x and y running one-at-a-time with y running Symphony, (ii) links x and y running simultaneously with y at a fixed 0 dBm, and (iii) links x and y running simultaneously with y running Symphony. Figure 4.8(a) shows the results over ten runs. In case (i) Symphony enables link y to operate at 0 dBm and still achieve full throughput. If link y is operated at 0 dBm together with link x, however, the throughput of link y drops significantly compared to x, as shown in case (ii) due to asymmetric channel access. We validate this by observing the difference in ETTs observed by links x and y. In case (iii), we show that Symphony increases the transmit power of link y to between 6 and 9 dBm to avoid asymmetry. Observe that link y didn't have to operate at the maximum power to let link x perceive its transmission.

In the second experiment, we show that when there is inherent asymmetry in the environment at default power, power control is beneficial in removing it and increasing the throughput

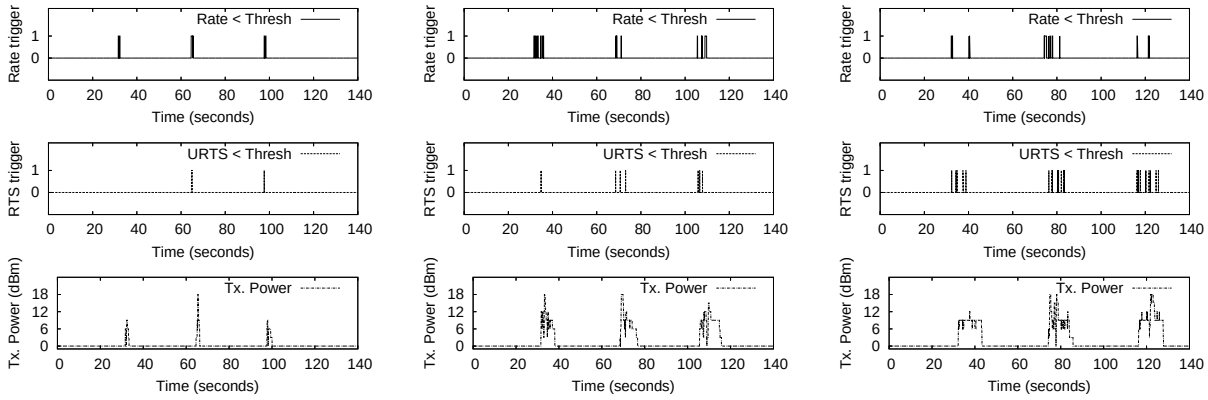

Figure 4.9: Efficacy in detecting and avoiding receiver side interference.

and fairness of links. As in Figure 4.5(b), we found the location AS3 where a sender gets significantly affected when running in conjunction with a sender at AS1. For this experiment, we use links x and z and consider two cases: (i) both x and z running together at maximum power, and (ii) x and z running together with Symphony. Figure 4.8(b) shows results averaged over ten runs. The graph shows that when operating at maximum power, link z gets significantly lower throughput than link x. When running x and z with Symphony, both links achieve greater throughput because they are able to operate independently at the lower transmit powers, thereby also demonstrating increased spatial reuse due to power control. *In both experiments, Symphony increases the throughput of asymmetry-affected links by three times.*

Avoiding Receiver Side Interference

For this experiment, we consider links p and q in Figure 4.5(b). Link q operates at maximum power, whereas link p operates with Symphony. The setup is such that link p operates at 0 dBm when run individually, HS1 and HS2 share the channel at maximum power, and HS2 does not perceive transmission on link p if p operates at 0 dBm and hence destroys packets at HR1 (i.e. causes receiver side interference) . In each run, we start a 20 Mbps UDP transfer on link p for 3 minutes, and start a 5 Mbps transfer on link q for a short periods (1 second, 5 seconds and 10 seconds) of time at different times during the 3 minutes. The bottom graphs in Figure 4.9 show that in response to link q's entry and exit, Symphony on link p increases and decreases transmit power respectively to avoid the adverse affect of receiver side interference. The graph shows that *Symphony is responsive to receiver side interference even at short timescales of 1*

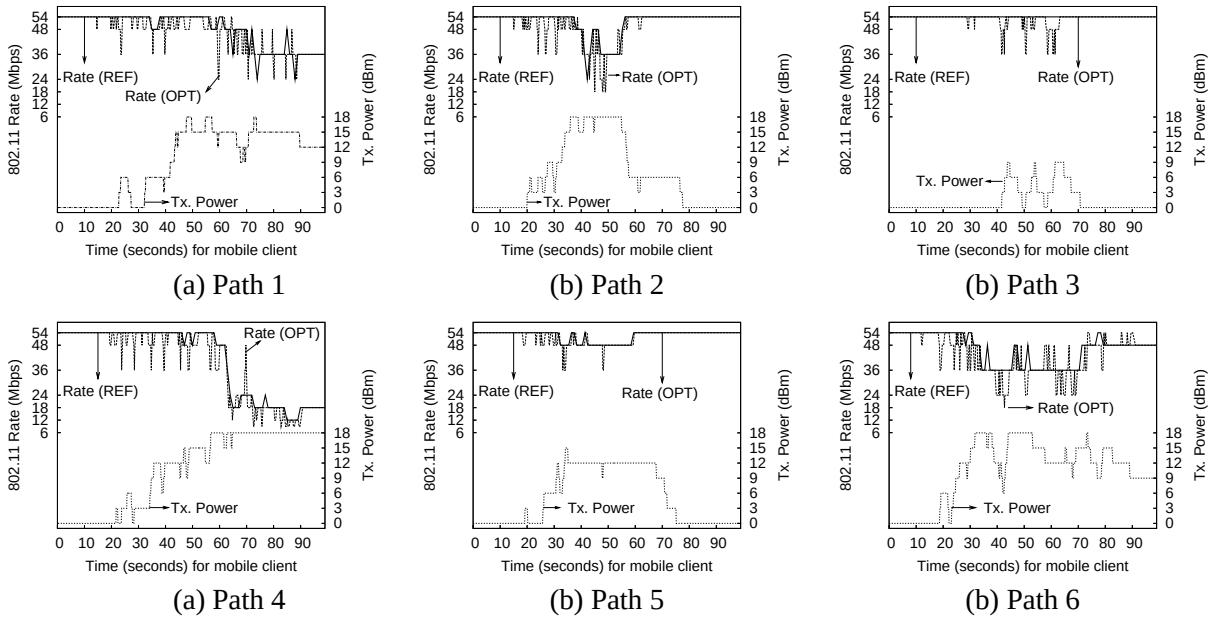

(a) Path 1 (b) Path 2 (b) Path 3

(a) Path 4 (b) Path 5 (b) Path 6

Figure 4.10: Adaptation to mobility: Symphony's behavior of rate and power in different paths.

second. The top two graphs in each of the cases (a), (b) and (c) show the rate and RTS triggers that identify the condition that triggered increased power as in line 1 of Algorithm 3. We note that the utility of RTS is not always sufficient to detect receiver side interference, primarily because even RTS is sent at the chosen (lower) transmit power with Symphony, which may not be perceived by HS2. In such a case, rate drops and leads to increased transmit power, thereby letting RTS reach HS2. Recall that if RTS is useful, it reduces unnecessary reduction in data rate.

Agility to Client Mobility

Figure 4.10 shows Symphony's behavior with client mobility. For this experiment, we consider three AP locations and six client paths as shown in Figure 4.5(c). The AP is at location T1 for paths p1,p2 and p3, at T2 for path p6, and at T3 for paths p4 and p5. On each of the paths, the client is mobile at a speed of 0.75m/s. We again setup VOIP calls from the AP to the client.

Figure 4.10(a) shows for path p1 that moving the client away starts affecting the bitrate in the OPT phase, and hence Symphony increases transmit power to maintain the bitrate to the same level as in the REF phase. As the client moves farther, even the rate in REF phase falls. Figure 4.10(b) shows that moving the client away on path p2 makes Symphony increase transmit power, but it also reduces the power when the client returns to the AP location. *Overall,*

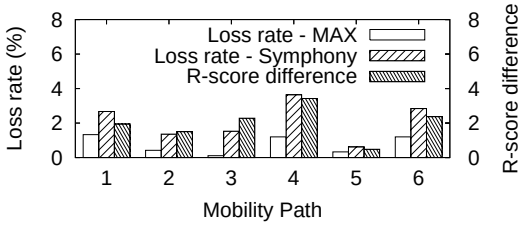

Figure 4.11: Application layer loss rate with mobility.

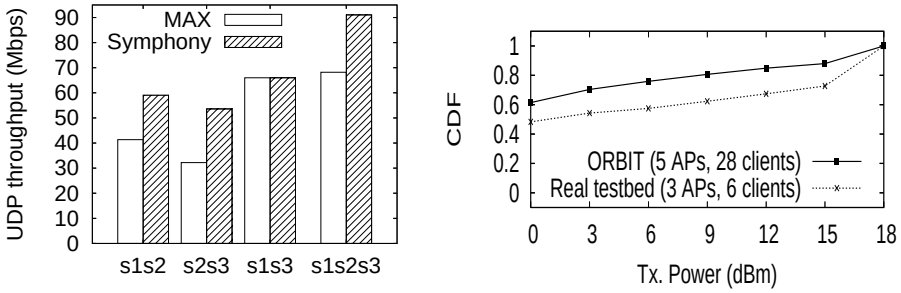

(a) (b)

Figure 4.12: (a) Spatial reuse experiment. (b) Large scale experiments.

Symphony opportunistically enables a link to operate up to 18dB lower than the default. Similar observations can be made on other paths.

Figure 4.11 shows the application level loss rate for Symphony in comparison to using default maximum transmit power, and the difference in R-score. R-score is a popular performance metric for the quality of voice calls [149]. An R-score of 70 or more is considered good voice quality. While R-score depends on several factors [149], in brief, the difference in R-score can be simplified to $40 \times (log(1 + 10e_m) - log(1 + 10e_s))$, where e_m and e_s represent the loss rate with maximum power and Symphony respectively. The graph shows that Symphony incurs little extra impact on application level loss on all the paths. Further, *in the worse case (path 4), the R-score using Symphony deteriorates only by 3.4 and the average R-score deterioration using Symphony is 2.* While the actual R-score also depends on the end-to-end delay, we note that with even the maximum loss rate (as in path 4) using Symphony, in order to achieve R-score as low as 70, the acceptable end-to-end delay is over 300ms. Further, in-depth analysis of the losses shows that significant part of the loss occurs when the client is far away from the AP. In a real mobility enabled WLAN, mobile clients handoff to closer APs for better quality, and hence we believe that even this application-level loss will not occur in practice.

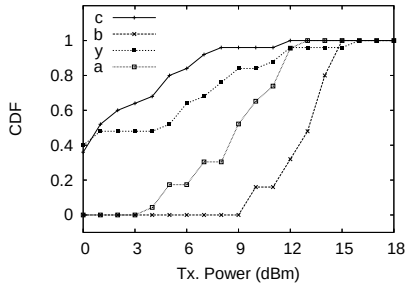

Figure 4.13: Interaction with operational network that does not run Symphony.

Spatial Reuse

Several studies have demonstrated that transmit power control leads to increased spatial reuse [1, 99]. Many of these works have shown simulation results over large topologies that can not be easily realized in a prototype testbed. However, for the sake of completeness, we perform a small-scale spatial reuse experiment with Symphony. In this experiment, we consider links (S1, z), (S2, x), and (S3, y) as in Figure 4.5(a) and make a subset of them operate simultaneously. We consider two cases: when the links operate at maximum power and when they operate with Symphony. *Figure 4.12(a) shows the aggregate throughput of the links, and clearly demonstrates 30-50% increased throughput for different combinations.*

Large-scale experiments

To assess the effect of high node density on Symphony, we emulate dense deployments on the indoor ORBIT testbed [131] with 5 APs and 28 clients, and on our office testbed with 3 APs and 6 clients. The inter-AP distance is 5 meters on ORBIT and 15m in our office testbed. Clients are within 15 meters of each AP in both cases. We setup bi-directional traffic between each client and its associated AP, and enable Symphony on all nodes. Figure 4.12(b) shows the CDF of the average transmit powers used by clients in each second over a period of 60 seconds. We observe that *Symphony enables clients to settle at much lower power levels in both sets of experiments. For instance, in the ORBIT experiment, clients settle at transmit power of 0dBm over 60% of the time and within 9dBm over 80% of the time.*

Experiments with an Operational Network

Finally, we perform experiments with Symphony operating in conjunction with an operational network. We choose channel 6 in 802.11g for our experiments, on which there are other APs and clients transmitting. We consider the setup (Figure 4.5(a)) with s3 as the AP location, and a,b,c,y as four client locations, with s3 running Symphony. By observing beacons at the location s3, we determine that there are at least 13 APs on channel 6. We perform this experiment over a period of 12 hours mostly during regular office hours when the network is active. For each client, the AP makes a 2 Mbps transfer to every client every half an hour, to estimate the power level appropriate for the clients. Figure 4.13 shows the CDF of transmit powers (average) chosen during each run for each client. The graph shows that *Symphony opportunistically reduces the transmit power on all links even when operating in conjunction with a non-Symphony-compliant network.* For instance, links (s3, c) and (s3, y) operate 6dB lower than the default transmit power 85% of the time. This result demonstrates Symphony's incremental deployability.

Chapter 5

The Sonata Framework: Design and Implementation

This chapter describes the design and implementation of Sonata, a framework that intelligently combines directionality, diversity and bit rate adaptation to maximize the uplink throughput of a mobile client. Sonata aims to reduce data transfer times, and increase connectivity durations.

5.1 Overview

In this section, we address each of the issues raised earlier in Chapter 3 (Section 3.7) through three key features:

1. We leverage the observation that several resource parameters at a location remain unchanged for long enough timescales (such as minutes to hours) [122], and can be learnt and reused across several mobile clients passing the same location. We split the process of discovering the right resource parameter settings across several clients, and store the information in a centralized easily accessible location (BeamManager). The BeamManager can be used by subsequent clients for rapidly converging to the appropriate settings. The discovery process is continued as a short phase at a subset of clients passing a location, to ensure that the parameter settings are up-to-date.

2. We leave to the client the choice of the exact beam for reaching a base station to make the learning process independent of client capabilities. The learning part only identifies the *direction* in which a beam should be formed.

3. We apply short-term adaptation to locally tune the parameters for better throughput to ensure that the algorithm makes appropriate choices to adjust to fine-timescale link fluctuations and diverse client capabilities. We name the location-based adaptation algorithm on the client R2D2, to stand for Robust Rate with Directionality and Diversity. R2D2 is the core and novel component of Sonata that we explore in detail in this dissertation. In

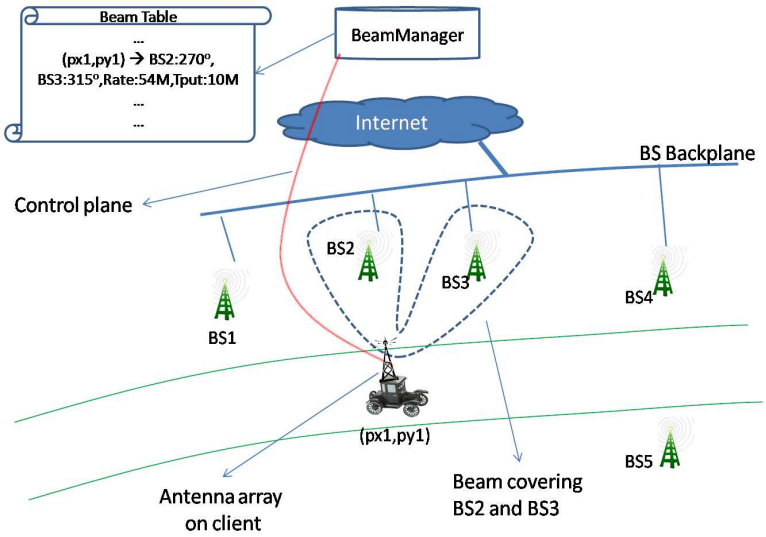

Figure 5.1: Sonata Framework Overview.

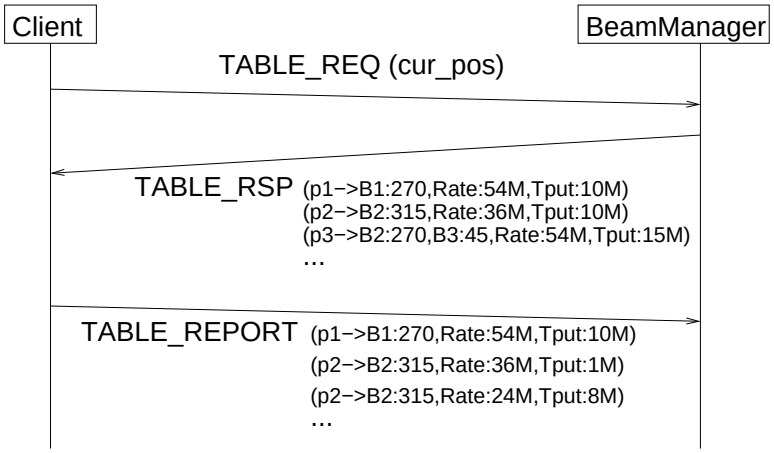

Figure 5.2: Sonata Protocol Overview.

what follows, we provide an overview of the complete solution, and in the next section, we expand on R2D2.

Figure 5.1 shows a block diagram of the overall solution. For a client at a given location that intends to do uplink transmission, our goal is to choose a set of receivers, form a transmit beam that covers all the receivers, and choose a bit rate that is appropriate to make packets decodable at all the receivers. The BeamManager is accessible via a control plane that can be assumed to be reliable and always available (such as a low rate but long-distance cellular network).

Algorithm 5 BeamManager pseudo-code

```
1:  FUNCTION on_TABLE_REQ (pos)
2:  for each posi close to pos do
3:     if exists unexplored combo then
4:        add posi, combo to ret_table;
5:     else if random() <= 0.9 then
6:        add combo with max. throughput;
7:     else
8:        randomly pick non-max. combo;
9:     end if
10: end for
11: send_TABLE_RSP (ret_table);
12:
13: FUNCTION on_TABLE_REPORT (table)
14: for each posi in table do
15:    if observed throughput > throughput_in_table then
16:       update combo and throughput in table;
17:    end if
18: end for
19:
```

Figure 5.2 shows the protocol between a client and the BeamManager to exchange and update beam configuration table, and Algorithm 5 shows the pseudo code of the BeamManager. For every, TABLE_REQ from a client, the BeamManager returns a TABLE_RSP containing a set of close-by locations that the client may traverse during its mobility, and the corresponding parameters. Lines 2 and 3 in the on_TABLE_REQ pseudo code represent the initial phase when the BeamManager makes clients try out all possible parameter combinations. Lines 5-7 ensure that periodically the BeamManager table is updated with recent best configurations.

The base stations in Figure 5.1 are connected through the backplane to the Internet. For a set of consecutive locations, we designate an *anchor* base station to collect packets forwarded by all the receiving base stations and to determine packet loss rate summary. The summary is sent back from the anchor base station to the client through one of the currently receiving base stations (which is easily determined from the receiving packets). Using one of the receiving base stations ensures that the client is not *deaf* to transmissions and the client indeed has a beam focused towards that base station. The anchor base station is less frequently changed than the beams and receivers themselves to reduce the effect of handoffs; this is primarily made possible because of the assumption that the base stations are connected on a backplane.

Algorithm 6 R2D2 pseudo-code.

```
1:  FUNCTION on_position_change ()
2:  t = get_mapping_table(pos);
3:  rcvr_angles = lkup_table(t, pos);
4:  beam = form_beam (rcvr_angles);
5:  configure_beam(beam);
6:  tx_rx_packets();
7:
8:
9:  FUNCTION on_pkt_loss_summary ()
10: if throughput < expected throughput then
11:     (PER,SNR) = get_per_snr();
12:     action = lkup_run_time_tbl(PER,SNR);
13:     perform(action);
14: end if
15: tx_rx_packets();
16:
```

The client reports back periodically to the BeamManager the performance of suggested parameter settings and any new resource parameter settings that the client tried out for inclusion into the database using the TABLE_REPORT message (Figure 5.2). The BeamManager performs a weighted update to resource parameters to ensure that a single observation due to momentary fluctuations will not change the settings significantly.

5.2 R2D2 Design

Our adaptation algorithm, R2D2, performs two functions (pseudo code shown in Algorithm 6).

1. At each new location, it uses the settings suggested by the BeamManager for transmission; the server chooses the settings that are the best in terms of throughput among the settings that are tried across clients in different sensing phases. The map of expected throughput at each location is also given to the client by the server. The client looks up the angles at which a beam should be formed at the current location, uses its antenna elements to form the appropriate transmit beam, and transmits and receives packets till either the location changes or a packet loss summary information is obtained from one of the base stations. The exact method of forming beams is client implementation dependent, and we discuss one such implementation in the prototype implementation section (Section 5.4).

Table 5.1: Run-time adaptation.

SNR ↓ / PER→	High	Low
High	↑ Diversity or ↓ Rate	↑ Rate
Low	↑ Directionality or ↓ Rate	Do nothing

2. At the existing location, if the obtained throughput is lower than the expected throughput indicated by the server, and if the packet loss summary depicts that the resource parameters are under-performing, the algorithm performs further run-time adaptation to better tune the parameters.

Run-time adaptation performed by the client can be easily understood by looking at Table 5.1. The goal of this adaptation is to maintain the link at as high a rate as possible by first adapting directionality and diversity; rate adaptation is done only when neither directionality nor diversity can help improve the PER. We derive this table using the observations made in our tradeoff study. The key idea of using this table is that the reduction in uplink throughput for a client manifests as variations in SNR and PER. If we consider the current resource parameter settings as a state, the next state we transition to depends on whether SNR and PER are high or low compared to a threshold. If SNR is already high and crosses the required threshold for packet decodability, and PER is still high, increasing diversity by adding additional receivers is a better option than increasing directionality. However, if the SNR is low and PER is high, there is increased chance of making some of the existing links to the receivers better by increasing directionality. If the PER is low, however, the client tries to increase rate to obtain better throughput.

5.3 R2D2 Evaluation

In this section, we evaluate the performance of R2D2 using trace-driven analysis. We collect several packet delivery traces through an extensive set of measurements in two different realistic settings, as shown in Figure 5.3. The first set of four runs are taken on a circular path in a parking lot around a building, and the second set of four runs are taken on a road with a speed limit of 40 MPH. The four receivers for each setting are placed as shown by the stars close to

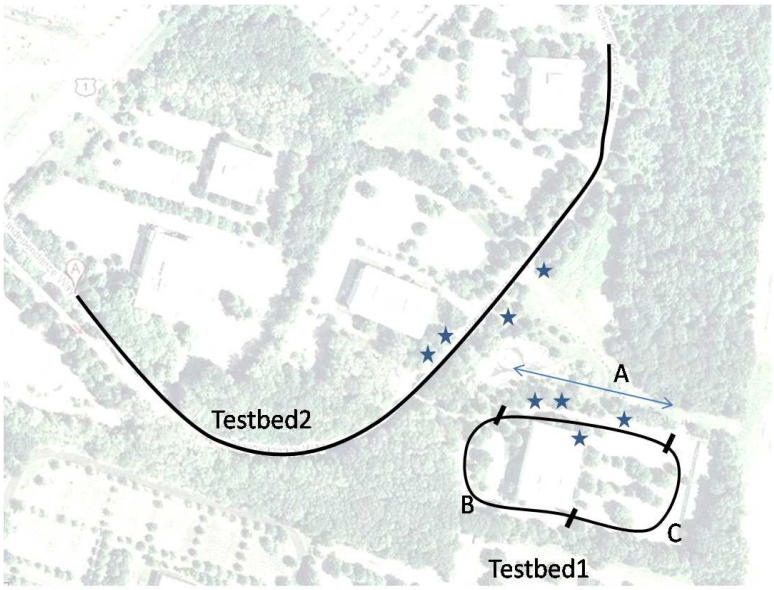

Figure 5.3: Vehicular testbeds.

the paths. We consider two distances of the receivers in each setting from the path to emulate both high SNR conditions when the receivers are close to the path, and low SNR conditions when the receivers are a little away from the path.

5.3.1 Methodology

To make data collection tractable, we use the following evaluation methodology. We use an omni-directional transmitter, and use packets with different transmission power for emulating the effect of creating beams of different beamwidths. We observe (a) with a phase array directional antenna [150] available with us, and (b) theoretically with MATLAB simulations, that a beam covering two well separated receivers (i.e. having two main lobes) will have approximately 3dB lower gain than a beam with single main lobe pointing to a receiver, and a beam covering three receivers will have 3dB lower gain than the two receiver case. Base on this observation, we pick four transmit power levels: 17dBm, 14dBm, 11dBm and 8 dBm. When transmitting at 17dBm, we assume that the beam is pointed to any one of the receivers, at 14 dBm to any combination of two receivers, at 11dBm to any combination of three receivers, and at 8 dBm to cover all four receivers.

The transmitter broadcasts 200 ICMP packets (1350 bytes payload) per second using different 802.11g PHY rates (18, 36, 54Mbps). The use of broadcast mode suppresses MAC-level features such as retransmissions, acknowledgments and RTS/CTS and enables us to measure the packet error encountered due to impairments suffered at the physical (PHY) layer. The receivers operate in monitor mode, in which the node can passively listen to all data on a particular channel without being associated with any AP. Receivers utilize the tcpdump [151] utility, which gives it relevant information on a per-packet basis from both the PHY and MAC layers. In addition, all the nodes continuously log their location and speed information using a GPS device. The system time on each node is set to the GPS time so that the system clocks of all nodes are synchronized. The transmitter includes its timestamp in the ICMP packet's payload so that the receiver can correlate the location from which each packet was transmitted.

5.3.2 Algorithms

We compare the performance of R2D2 with five alternate techniques outlined below. Some of these algorithms represent existing solutions, while some of them are our enhancements to the existing algorithms to demonstrate that even after fixing the existing algorithms there is still scope for improvement, and R2D2 is effective in improving the throughput close to what an ideal adaptive algorithm would achieve.

- *R2D2-LOC*: A version of R2D2 that just uses parameter combination as suggested by the location-based database, and does not perform any further run-time adaptation in response to changing network conditions.
- *ViFi*: A system using omni directional transmissions from the transmitter, as in ViFi [93], and that transmits at the highest possible rate and does not perform any rate adaptation.
- *AR-ViFi*: An enhanced version of ViFi that adapts rate based on observed packet loss to choose the best rate that is appropriate for maximizing uplink throughput.
- *Mobisteer*: A system using location-based beam adaptation, with the beam pointing to a single receiver at any point of time, and rate adaptation being done independently, as in Mobisteer [109].
- *MAX*: An oracle system that chooses the best combination on a per packet basis to emulate the maximum achievable throughput in the given conditions.

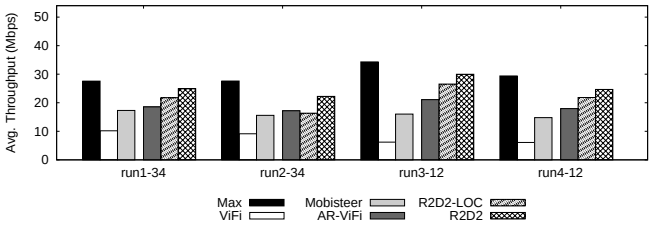

(a) Parking Lot

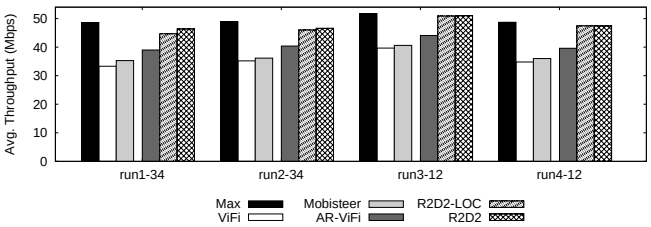

(b) Parking Lot, High SNR

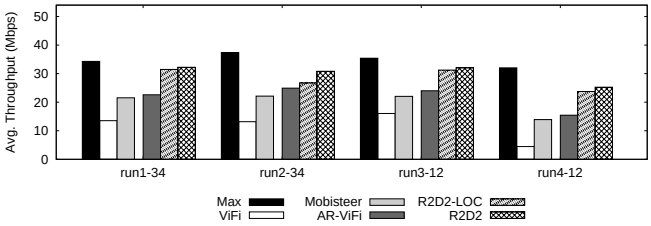

(c) Independence Way

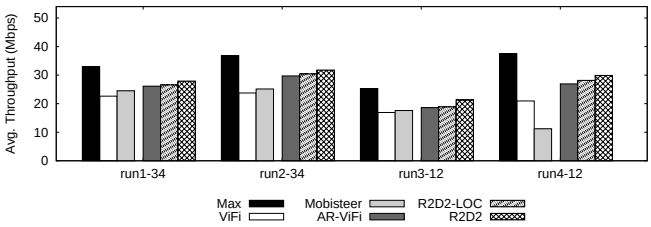

(d) Independence Way: High SNR

Figure 5.4: Average throughput obtained by R2D2 compared to several algorithms.

5.3.3 Results

Performance Improvement

Figure 5.4 shows the throughput obtained by the six algorithms in several runs in each of the different settings. For the algorithms that use the location-based database for parameter selection, we use two runs other than the run under consideration for training the database. For instance, run1-34 indicates the performance of the algorithms in run1 when using 3rd and 4th runs for training the database. Figure 5.4(a) shows the throughput for the parking lot case when the

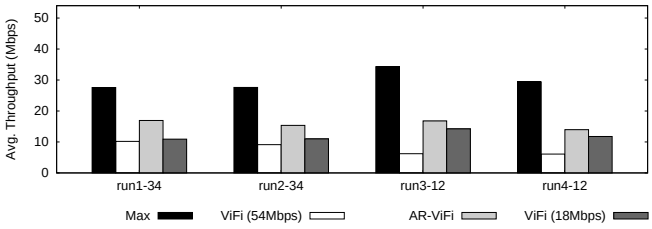

(a) Parking lot

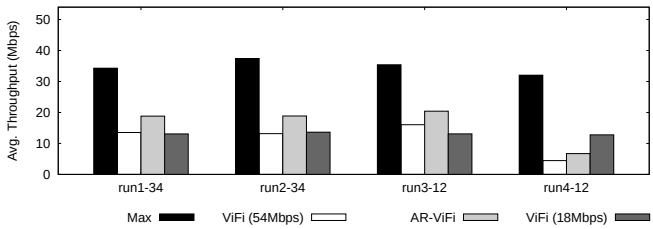

(b) Independence Way

Figure 5.5: Variants of Vifi.

average SNR is low, i.e. the receivers are a little farther away (about 25 meters) from the path. In this case, R2D2 performs significantly better than existing algorithms and even their extensions. The difference between ViFi and AR-ViFi (also see Figure 5.5) shows the additional benefit of rate adaptation in diversity based solutions, an extension often acknowledged but not implemented in previous works [93]. The difference between AR-ViFi and MAX shows that even after extension to ViFi, there is still a significant scope for improvement. The difference between R2D2 and R2D2-LOC (also see Figure 5.6) shows the requirement for additional runtime adaptation over location-based beam and rate selection. Finally, the difference between AR-ViFi, Mobisteer and R2D2 demonstrates that carefully trading off directionality and diversity and jointly performing rate adaptation can take the uplink throughput significantly closer to MAX. Similar observations can be made on the Independence Way path in Figure 5.4(c).

Figures 5.4(b) and 5.4(d) plot the throughput under high SNR conditions, i.e. when the receivers are very close to the paths, emulating road-side AP deployments. Even in this case, R2D2 performs better than other algorithms and is close to what MAX can achieve. Observe, however, that diversity-based algorithms (ViFi and AR-ViFi) perform much better in this scenario relative to the low-SNR case compared to directionality-based algorithm (Mobisteer). This is in line with the theoretical predictions as in Figure 3.17. In setups with high average SNRs, higher directionality would not benefit as much as higher diversity that helps mask off

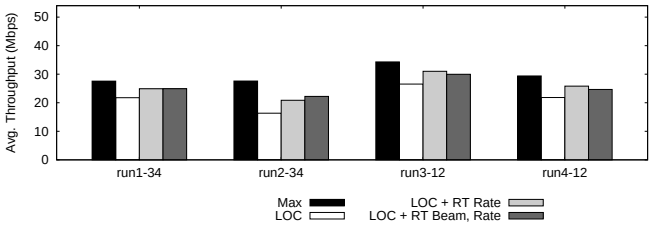

(a) Parking lot

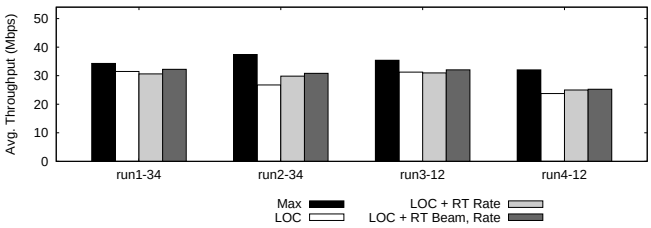

(b) Independence Way

Figure 5.6: Variants of R2D2.

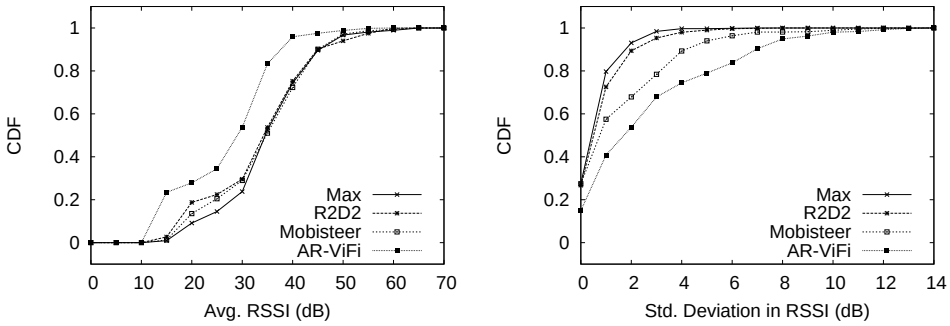

(a) Mean SNR distribution

(b) Std. deviation in SNR

Figure 5.7: The distribution of mean, and std. deviation of, SNR for each scheme relative to the maximum. R2D2 improves link robustness by increasing the mean SNR and reducing the variance.

variance in SNR. [93] observe that in mobile settings deep fades or variance in SNR can be very high. On the other hand, in low SNR regimes, increased diversity further reduces gain to an individual receiver thereby hurting packet reception probability; increased directionality increases the average SNR (at least to a subset of the receivers) thereby increasing packet reception probability.

To highlight the effect on SNR better, Figure 5.7 shows the distribution of the mean and standard deviation in the received signal strength indicator (RSSI) achieved by each of the

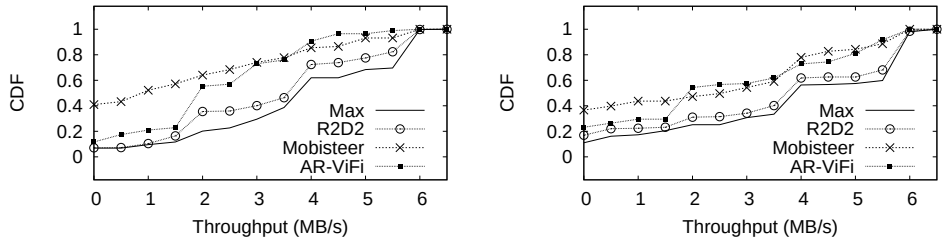

(a) CDF of throughput for parking lot (b) CDF of throughput for Independence way

Figure 5.8: CDF of throughput obtained by several algorithms at several client locations.

schemes relative to MAX. RSSI is an estimate of the signal energy at the receiver during packet reception, measured during the PLCP headers of arriving packets and reported on proprietary (and different) scales. The Atheros cards we use, for example, report RSSI in dB relative to the noise floor [152]. For a single link, since SNR = RSSI / noise floor, and since the noise floor is relatively constant, here, we use RSSI to represent SNR. The CDFs shown in Figure 5.7 use the average and standard deviation in SNR calculated over successive 100-millisecond intervals. As RSSI is available only if the corresponding packet was successfully decoded, and since we need the SNR of lost packets to calculate the variance, we make the assumption that lost packets have an SNR just below the receive threshold for the bit-rate used to send that packet. Using this assumption, we report the "best-case" mean and standard deviation of SNR. Figure 5.7(a) shows that both R2D2 and Mobisteer come close to approximating the maximum achievable mean SNR. The median for both schemes is 5 dB higher than the pure diversity based scheme. However, Figure 5.7(b) shows that both pure directionality and diversity based schemes are unable to effectively deal with SNR variance (due to deep fades). The distribution of standard deviation for both schemes has heavy tails with a maximum deviation of 14 dB. The graph also shows that even MAX observes standard deviations of up to 4dB. R2D2 is able to closely approximate the behavior of MAX (and eliminate a large part of the heavy tails of the distribution).

Figure 5.8 shows the CDF of throughput for R2D2, Mobisteer and AR-ViFi for run 3, using runs 1 and 2 for training the database. The CDF clearly demonstrates that R2D2 has high throughput at a higher percentage of locations. For instance, in the parking lot case, the median throughput with R2D2 is 30 Mbps, whereas that with AR-ViFi and Mobisteer is 14 and 16 Mbps respectively.

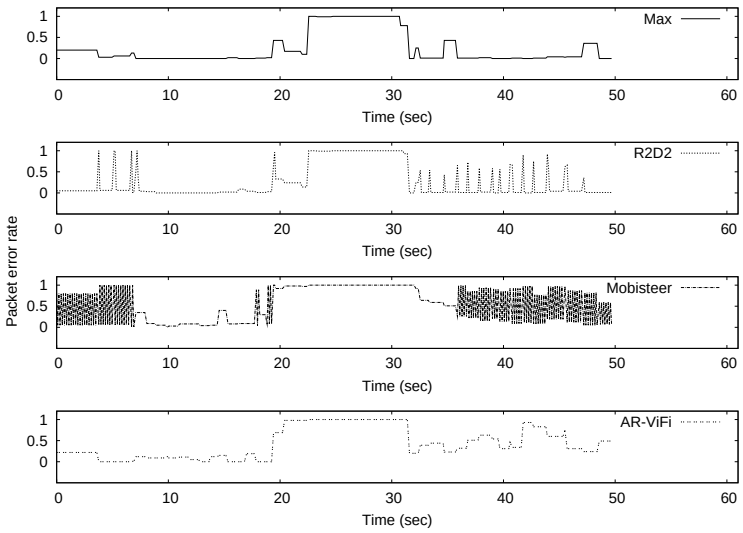

Figure 5.9: Link packet error rate over time (in seconds) for the three schemes relative to the achievable maximum. Each point is an average over 100 milliseconds. R2D2 comes closest to Max in both LOS and NLOS regions.

Figure 5.9 shows that the throughput improvements of R2D2 are obtained by reducing the packet loss significantly. This is crucial to ensure that the channel is used effectively by the transmitter for uplink transmission. The graph shows packet error rate over time for the three schemes Mobisteer, AR-WiFi and R2D2 relative to MAX. The lines indicate averages over successive 100-millisecond intervals.

Algorithm Behavior

Figure 5.10 expands on the behavior of the algorithms to demonstrate the efficacy of R2D2 over others in adapting to fluctuating channel conditions and mobility. Figure 5.10(a) shows the throughput with time (averaged every 200 milliseconds) for the parking lot case with the low SNR setup for run3 using runs 1 and 2 for training the database. The graph clearly shows that in many regions R2D2 makes a better choice of parameters than AR-ViFi or Mobisteer. Figure 5.10(b) shows similar result for the Independence Way setting with the low SNR case.

In Figure 5.11, we zoom into the behavior of R2D2 and compare its performance with R2D2-LOC and MAX, to demonstrate the efficacy of additional run-time adaptation over using location-based database for parameter selection. The graph clearly shows that location mis-predicts the best choice in several instances making R2D2-LOC perform worse than MAX;

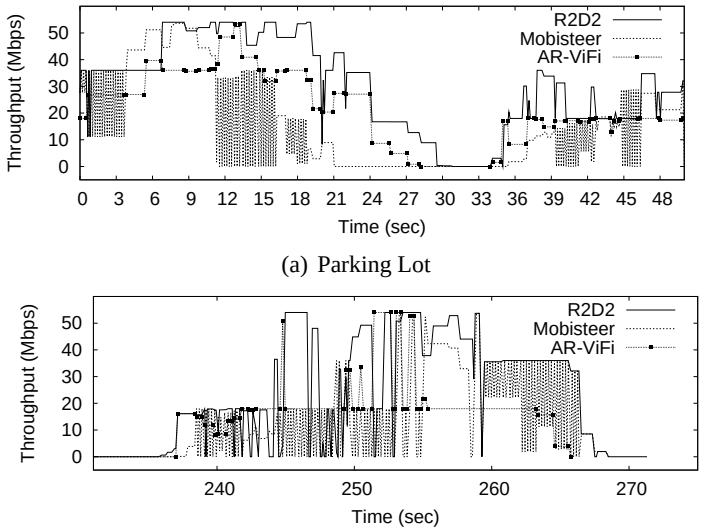

(a) Parking Lot

(b) Independence Way

Figure 5.10: Instantaneous throughput obtained by R2D2 compared to a variety of algorithms.

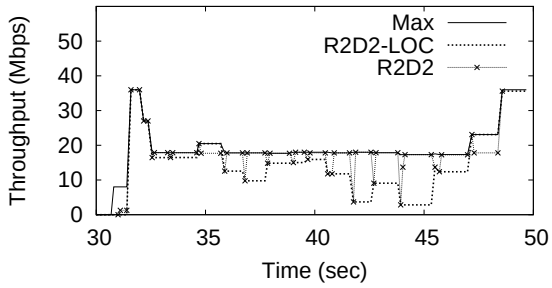

Figure 5.11: Parking Lot: zoomed-in.

whereas R2D2 detects these conditions and adapts well to match the performance of MAX.

Figure 5.12 shows the combination of receivers chosen with time by the transmitter when using R2D2, in comparison to the extremes of choosing all receivers (as in AR-ViFi) and only one receiver (as in Mobisteer). We use information from run3-12 for the parking lot case and run4-12 for the Independence Way case. The graph demonstrates that in several locations a middle-ground between the extremes, i.e. choosing a subset of visible receivers is the best choice. Note that while Mobisteer and R2D2 overlap in Figure 5.12(a) after 20 seconds, the choice of rate is different due to run-time adaptation, and hence R2D2 performs better than Mobisteer, as observed in Figure 5.10(a).

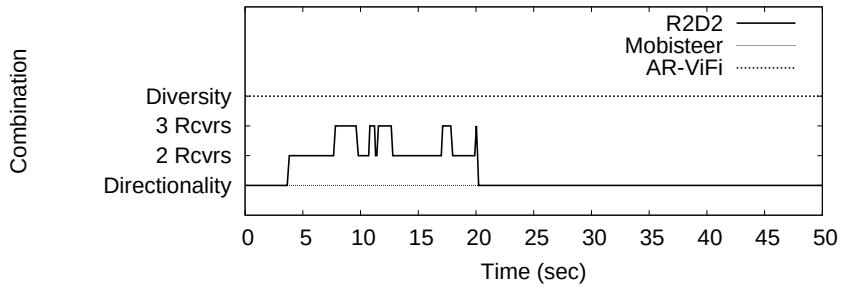

(a) Parking lot

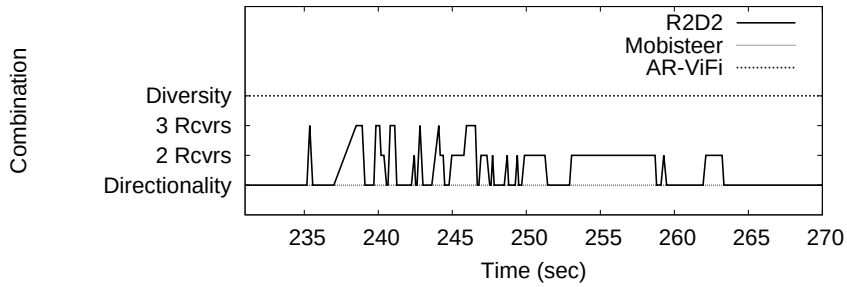

(b) Independence way

Figure 5.12: Receiver combinations chosen in different settings.

5.4 System Implementation

5.4.1 Overview

Our prototype system (Figure 5.13) consists of a car carrying a Phocus Array beamforming node [150] as the client that does uplink transmissions (on 802.11g, channel 6), a set of receivers that are small form-factor PCs with 6dBi gain external antennas [131]; one of the receivers acts as an anchor base station to which the other receivers forward packets through additional wireless interfaces on channels that the transmitter does not use (802.11g, 1 and 11). Using wireless backhaul avoids setting up wired connectivity outdoors. The Phocus Array is a full-fledged 802.11 node with a beamforming antenna. This antenna consists of an array of eight elements arranged in a regular octagon. The antenna is electronically steerable, i.e., a specific beam pattern out of the several precomputed beams can be chosen from software on the fly.

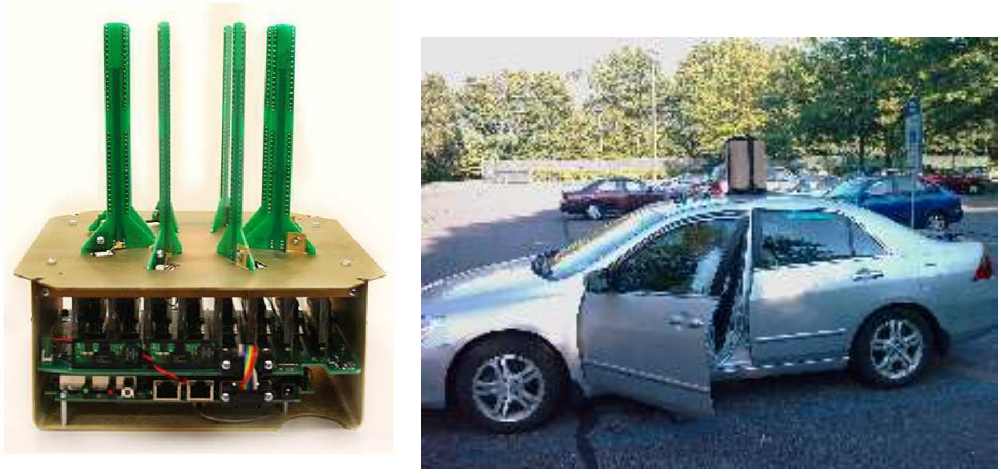

(a) (b)

Figure 5.13: (a)Transmitter with beamforming antenna, (b) Transmitter enclosed in a box and mounted on a car.

5.4.2 Beamforming

We generate several beams offline by setting antenna element weights (phase and magnitude) appropriately. The exact element weights to be set for each beam pattern can be found in our technical report [153]. In all, we have (a) one omni-directional beam, (b) 8 beams with single main lobe, with each beam shifted clockwise by 45^o from the previous beam such that they together cover the entire 360^o around the transmitter, (c) 28 beams with two main lobes at different angles relative to the direction of travel of the car, and (d) 49 beams with three main lobes at different angles relative to the direction of travel of the car. The single lobe beams are adopted from our previous work [40], and possess the characteristic of having very low side lobes (a front-to-side lobe ratio of 18dB), and about 8 dB extra gain over the omnidirectional pattern. The multi-lobe beams are derived in MATLAB by super-imposing the single lobe beams and deriving the combined weights using conventional antenna theory [154].

Figure 5.14 shows an example set of beams with different number of main lobes at different angles. These beams are configured on the Phocus Array, and a single command can change the pattern that the node uses to transmit at run-time. Changing the pattern of transmission to one of the configured patterns on the Phocus Array node takes about 150 μs. Observe that we make an assumption here that the gain is equal for both lobes in the two-lobe beams and for all three lobes in the three-lobe beams. We make this simplifying assumption primarily due to lack of

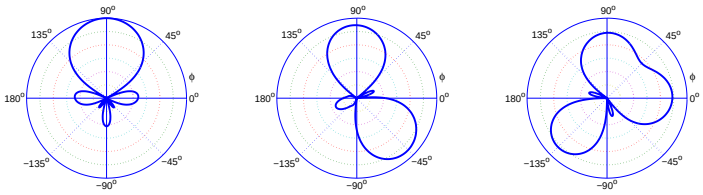

(a) 1 lobe at 90^o (b) 2 lobes at 90^o,-45^o (c) 3 lobes at 0^o, 90^o, -135^o

Figure 5.14: Example beams with different number of main lobes at different angles. The motion of the car is along zero.

rapid adaptive beamforming support on the existing Phocus Array hardware. For rapid adaptive beamforming, we envision that arbitrary beams with different gains towards each receiver can be formed, which we leave as a part of our future work.

5.4.3 Mapping beams to receivers

Mapping a beam to a set of receivers indicated by the BeamManager at each location involves two steps.

1. Angular Localization

First, the client needs to know the beam angle required for each receiver for maximum throughput. Note that this angle need not be the physical angle between the transmitter and the receiver relative to the direction of the client's motion; a beam in a completely different direction might be the best beam due to multipath reflections. The information provided by the BeamManager itself for each receiver includes $<$*Location, RecevierID1, BeamAngle1, ReceiverID2, BeamAngle2,...*$>$. In the training phase, the client receives beacons (or pilots) from the base stations when using the eight single lobe beams in turn. The beam that gets the beacons at the highest RSSI determines the angle of the receiver.

2. Calibration

While angular localization enables the client to know the direction in which it has to create a beam to reach a receiver relative to the client's direction of motion, it does not know the exact orientation of the client antenna. This is essential to ensure that the beam is actually

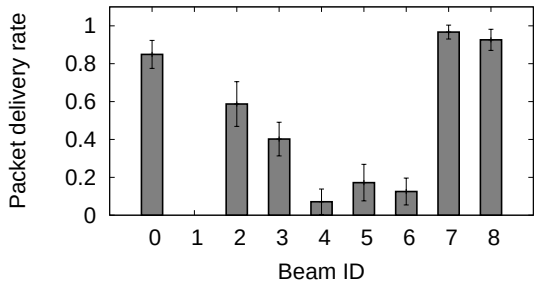

Figure 5.15: Angular Localization.

formed in the desired direction. For addressing this problem, since we do not have control of the antenna placement on the clients, we design a simple calibration procedure to calculate the offset to be applied while mapping beams to receiver positions. The offset ensures that the beam is selected not just relative to the direction of motion of the car, but also relative to the placement of antenna array on the client. The procedure is as follows. For each of a small set of receivers that the BeamManager indicates the positions and the beam directions to be used, the client uses the eight single lobe beams in turn and finds out if a beam selected in the direction indicated by the BeamManager for a given receiver indeed gives the maximum throughput, and if not which beam gives the maximum; the difference in angle between the beam that gives the maximum throughput and the beam indicated by the BeamManager is the offset. We use a subset of receivers instead of one for greater confidence in the offset estimated.

5.4.4 Data Transfer Protocol

We use the idea similar to MRD [112] and ViFi [93] to provide packet loss summary reports from the anchor base station to the client. The client uses the summary reports for run-time adaptation. We do not currently perform the optimization that ViFi [93] proposed—of using probabilistic relaying of packets to the anchor base station—to minimize traffic on the backhaul. Instead, we over-provision the backhaul by using two 802.11g channels, to concentrate the work on the client to base station link and the tradeoff involved between directionality and diversity.

5.4.5 Evaluation

We perform our prototype evaluation in the parking lot setup, as shown in Figure 5.3. The BeamManager is trained with two client runs on the same path and the best set of receivers at each location are determined. A location is uniquely determined across all clients and the BeamManager by considering only the first four digits after the decimal point of the GPS coordinates (latitude and longitude). This method defines unique rectangles that enables sharing beam configuration and rate selection information across clients. We perform angular localization for the same set of locations to determine the beam angle that is best at each location for each receiver, very similar in methodology to Mobisteer [109].

5.5 Limitations and Discussion

Compared to an omni-directional beam, the cost of choosing a "wrong beam" can be very high, for example due to lack of enough information at a location or due to inaccurate position estimation due to GPS errors [122]. This can be easily appreciated by looking at Figure 5.15, where six beams out of the eight provide significantly lower packet delivery rate than the omni-directional beam. This observation indicates that the training phase is a very important component of the system, requiring a significant number of training samples to gain enough confidence in the parameter setting at each location, contrary to our initial assumption. This problem is especially challenging with high mobility of clients; clients at a speed of 45MPH take about 500ms to cross a 10 x 10 meters GPS location. However, we envision that resource parameter settings at a location can be used across different clients, and hence the training phase can be distributed across the large client-base.

In Section 3.7, we discussed the effect of number of antenna elements of the transmitter on the theoretical tradeoff between directionality and diversity. An additional systems issue with number of elements that we do not address in this paper is that while greater number of elements can form thinner beams, the duration of validity of a beam reduces because of client mobility. To keep the overhead low, the duration of validity of the beam should be taken into consideration for choosing the number of elements to use in forming beams. This observation applies to WiMAX or LTE as well, where to minimize the feedback from the client and to track the fast-moving client, beamforming can be performed based on average instead of instantaneous statistics as is the case in statistical eigen beamforming [155].

Our evaluation uses WiFi devices primarily due to their ready availability. However, we believe that the solution itself is equally applicable to other wireless technologies such as WIMAX [156] and LTE [107], as long as other extensions such as the ability of connecting to multiple base stations are provided. For instance, the current CDMA standards already allow data combining from multiple base stations on the uplink [157].

While we focus on uplink connectivity, diversity and directionality can also be leveraged in the downlink direction by using multiple base stations to transmit to a single client; however, this requires additional functionalities on the base stations such as synchronization. While we do not expand on this direction in this paper, this topic is already of significant interest in the LTE and WIMAX standards, and forms an interesting problem for future work.

Finally, while it is unlikely that mobile devices such as cell phones and PDAs will incorporate multi-antenna systems with more than three elements due to their small form-factor, we envision the evolution of relay devices on moving vehicles that can afford to carry such large form-factor antenna systems; user devices can connect to the relay devices that in turn connect to the base stations. Similar proposal for a different purpose was made by Rodriguez et al. [121]. As opposed to the existing approach of overloading client devices with multiple radios (e.g. GPS, WiFi, 3G, etc.), this architecture aims to shift significant complexity from the client devices to relays. By radically reducing client device complexity, this approach has the potential to help reduce energy consumption for hand-held devices, in addition to increasing network capacity.

Chapter 6

Conclusions

In this dissertation, we identify the emerging requirements of high throughput, increased battery life, and support for mobility for next generation CSMA wireless networks. We demonstrate how parameter adaptation has to possess two fundamental, but conflicting properties, *agility*, and *robustness*, to satisfy these requirements. Adaptation techniques need to be responsive to the vagaries of the wireless medium, while not compromising on link reliability and network stability. We show that striking this *agility-robustness tradeoff* effectively is the key challenge for resource parameter adaptation in these networks.

We then show that several unique practical issues in existing incarnations of CSMA wireless networks make it challenging to incorporate these properties in the design of parameter adaptation techniques. These include: (a) the use of unlicensed spectrum, (b) CSMA/CA scheduling's inability to prevent interference, (b) the lack of mechanisms for accurate interference measurement and feedback, and (c) the requirement for distributed parameter adaptation. To address these challenges, we focus on joint resource parameter adaptation since no single parameter is sufficient to satisfy all the requirements. We also realize the need to select different parameters for indoor and outdoor environments because of their disparate characteristics.

To show the applicability of this work to real network deployments, in addition to an extensive set of experiments, we design and implement two resource management frameworks, *Symphony* and *Sonata*, for indoor and outdoor environments respectively. Indoors, *Symphony* increases network capacity and battery life for mobile clients by addressing the classical problem of joint, per-link, transmit power control and rate adaptation. For improved robustness, *Symphony* uses novel mechanisms based on measuring the expected transmission time (ETT), and the *utility* of RTS/CTS frames, while relying on a *learning* approach to converge quickly to the right resource parameter choice. These innovations allow *Symphony* to increase network throughput by up to 50% across four realistic deployment scenarios, and battery lifetime by up to 46%. Further, *Symphony*'s original, two-phase design makes it readily deployable even

in the presence of non-compliant nodes. Outdoors, the Sonata framework introduces a novel and fundamental tradeoff between directionality and base station diversity for uplink transmissions. Using a new location-based approach for improved parameter convergence, *Sonata* is able strike the agility-robustness tradeoff effectively. Extensive experiments demonstrate that *Sonata* improves link throughput by over 2x relative to state-of-the-art, and is incrementally deployable.

Together, these frameworks prove that, achieving the right balance between agility and robustness can enable CSMA wireless networks to transition to the wireless broadband era.

6.1 Future Directions

We propose addressing the following open issues as part of future work.

- In this work, the *Symphony* framework has focused mainly on supporting applications that do not distinguish between individual packets (e.g. VoIP and TCP). While these applications cover a broad portion of today's traffic types, it does not cover applications that use packets with different priorities (e.g. video). An interesting future direction is to understand the effect of power and rate adaptation on such applications.

- Another open challenge is that of the integration of parameter adaptation with QoS provisioning. How can parameter adaptation be designed with QoS guarantees in mind?

- At present, the *Sonata* framework mainly addresses the connectivity concerns of a single link. Its main focus is on ensuring fast-convergence for parameter adaptation. From a network-wide perspective, although its use of directionality as a key building block promises the potential for greater spatial reuse, future work will have to look at integrating the property of interference-awareness into this framework.

References

[1] T.-S. Kim, H. Lim, and J. C. Hou, "Improving spatial reuse through tuning transmit power, carrier sense threshold, and data rate in multihop wireless networks," in *MobiCom '06: Proceedings of the 12th annual international conference on Mobile computing and networking*. New York, NY, USA: ACM, 2006, pp. 366–377.

[2] "Smartphone market forecast (japanese shipments). source: Research and markets," http://www.linuxdevices.com/news/NS4728745041.html.

[3] [Online]. Available: http://en.wikipedia.org/wiki/Guglielmo_Marconi

[4] [Online]. Available: http://www.news.com/2100-1039_3-6159491.html

[5] [Online]. Available: http://www.usatoday.com/tech/wireless/phones/2006-08-20-cellphone-recycling_x.htm

[6] [Online]. Available: http://tinyurl.com/2ffm3c

[7] "Report of NSF workshop on Overcoming Barriers to Disruptive Innovation in Networking," January 2005. [Online]. Available: http://www.arl.wustl.edu/netv/noBarriers_final_report.pdf

[8] [Online]. Available: http://tinyurl.com/texdw

[9] [Online]. Available: http://tinyurl.com/gx5qw

[10] [Online]. Available: http://tinyurl.com/29d2ay

[11] [Online]. Available: http://en.wikipedia.org/wiki/Windows_Mobile

[12] [Online]. Available: http://www.google.com/intl/en_us/mobile/index.html

[13] [Online]. Available: http://www.apple.com/iphone/

[14] "Google phone." [Online]. Available: http://www.google-phone.com/

[15] IEEE, *Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE-SA Standards Board, 1999.

[16] IEEE WiMax Standard, *Standard for Local and metropolitan area networks Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems Amendment for Physical and Medium Access Control Layers for Combined Fixed and Mobile Operation in Licensed Bands.*, IEEE-SA Standards Board.

[17] [Online]. Available: http://hspa.gsmworld.com/

[18] CDMA2000 1xEV-DO, "Qualcomm Inc." [Online]. Available: http://web.archive.org/web/20061104082805/http://www.qualcomm.com/technology/1xev-do.html

[19] E. Dahlman, H. Ekstrm, A. Furuskr, Y. Jading, J. Karlsson, M. Lundevall, and S. Parkvall, "The 3G Long-Term Evolution - Radio Interface Concepts and Performance Evaluation," in *IEEE VTC*, 2005.

[20] "WiMAX and WiFi Together: Deployment Models and User Scenarios," 2007. [Online]. Available: http://tinyurl.com/2tk2t7

[21] [Online]. Available: http://tinyurl.com/yq4ucs

[22] [Online]. Available: http://tinyurl.com/2cpa9b

[23] [Online]. Available: http://www.apple.com/itunes/

[24] [Online]. Available: http://www.zune.net/en-US/

[25] "UMA," 2007. [Online]. Available: http://www.umatoday.com/umaOverview.php

[26] IEEE 802.11p, *Wireless Access in Vehicular Environment (WAVE) in Standard 802.11 Information Technology Telecommunications and Information Exchange Between Systems, Local and Metropolitan Area Networks, Specific Requirements, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications (802.11p)*, February 2006.

[27] P1609.3/D17, *IEEE Wireless Access in Vehicular Environments (WAVE) Networking Services*, November 2005.

[28] P1609.4/D05, *IEEE Wireless Access in Vehicular Environments (WAVE) Channel Coordination*, November 2005.

[29] October 2007. [Online]. Available: http://softwarecommunity.intel.com/articles/eng/1556.htm

[30] A. Akella, G. Judd, S. Seshan, and P. Steenkiste, "Self-management in chaotic wireless deployments," in *MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking*. New York, NY, USA: ACM, 2005, pp. 185–199.

[31] P. Thornycroft, Aruba Networks, "WLAN RF Architecture Primer," http://www.arubanetworks.com/pdf/technology/whitepapers/wp_RFARCH.pdf.

[32] "Cisco Aironet 350 Series Access Points." http://www.cisco.com/univercd/cc/td/doc/product /wireless/airo_350/accsspts/ap350hig/ap350ch2.htm.

[33] V. Mhatre, K. Papagiannaki, and F. Baccelli, "Interference mitigation through power control in high density 802.11 wlans," in *Proc. of IEEE INFOCOM*, 2007.

[34] 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 '06: Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications*, Pisa, Italy, 2006, pp. 51–62.

[35] A. Kashyap, S. Ganguly, and S. R. Das, "A measurement-based approach to modeling link capacity in 802.11-based wireless networks," in *ACM MobiCom '07: Proceedings of the 13th annual ACM international conference on Mobile computing and networking*, Montreal, Quebec, Canada, 2007, pp. 242–253.

[36] D. Niculescu, "Interference map for 802.11 networks," in *IMC '07: Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*. New York, NY, USA: ACM, 2007, pp. 339–350.

[37] X. Yang, "Efficient packet scheduling in wireless multihop networks," Ph.D. dissertation, University of Illinois at Urbana-Champaign, Champaign, IL, USA, 2005, adviser-N. Vaidya.

[38] J. Zhu, B. Metzler, X. Guo, and Y. Liu, "Adaptive csma for scalable network capacity in high-density wlan: A hardware prototyping approach," in *INFOCOM*, 2006.

[39] A. Mishra, V. Shrivastava, D. Agrawal, S. Banerjee, and S. Ganguly, "Distributed channel management in uncoordinated wireless environments," in *ACM MobiCom*, 2006, pp. 170–181.

[40] M. Blanco, R. Kokku, K. Ramachandran, S. Rangarajan, and K. Sundaresan, "On the Effectiveness of Switched-beam Antennas in Indoor Environments," http://www.nec-labs.com/∼ravik/RESEARCH/directional.pdf.

[41] "Notebooks surpass pass desktops in U.S. retail sales." http://news.cnet.com/2100-1044_3-6033967.html.

[42] "Infonetics research," http://www.infonetics.com.

[43] 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.

[44] M. Lacage, M. H. Manshaei, and T. Turletti, "IEEE 802.11 rate adaptation: a practical approach," in *Proc. of ACM MSWiM*, 2004, pp. 126–134.

[45] MadWifi, "Multiband Atheros Driver for WiFi," http://www.madwifi.org/.

[46] J. C. Bicket, "Bit-rate selection in wireless networks," Master's thesis, M.I.T., Cambridge, MA, February 2005.

[47] S. H. Y. Wong, S. Lu, H. Yang, and V. Bharghavan, "Robust rate adaptation for 802.11 wireless networks," in *Proc. of ACM MOBICOM*. New York, NY, USA: ACM Press, 2006, pp. 146–157.

[48] K. Ramachandran, H. Kremo, M. Gruteser, P. Spasojevic, and I. Seskar, "Scalability analysis of rate adaptation techniques in congested ieee 802.11 networks: An orbit testbed comparative study," in *Proceedings of 8th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, Helsinki, Finland, 2007.

[49] G. Holland, N. Vaidya, and P. Bahl, "A rate-adaptive MAC protocol for multi-hop wireless networks," in *Proc. of ACM MOBICOM*, Rome, Italy, 2001.

[50] J. Pavon and S. Choi, “Link Adaptation Strategy for IEEE 802.11 WLAN via Received Signal Strength Measurement,” in *Proc. of IEEE ICC*, May 2003.

[51] I. Haratcherev, K. Langendoen, R. Lagendijk, and H. Sips, “Hybrid rate control for IEEE 802.11,” ACM International Workshop on Mobility Management and Wireless Access, October 2004, pp. 10–18.

[52] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, “Opportunistic media sccess for multirate ad hoc networks,” Proc. of ACM MOBICOM, 2002, pp. 24–35.

[53] J. Zander, “Performance of optimum transmitter power control in cellular radio systems,” *IEEE Transactions on Vehicular Technology*, vol. VT-41, no. 1, pp. 57–62, February 1992.

[54] G. J. Foschini and Z. Miljanic, “A simple distributed autonomous power control algorithm and its convergence,” *IEEE Transactions on Vehicular Technology*, vol. VT-42, no. 4, pp. 641–646, April 1993.

[55] R. D. Yates, “A framework for uplink power control in cellular radio systems,” *IEEE JSAC*, vol. 13, no. 7, pp. 1341–1347, September 1995.

[56] N. Bambos, S. C. Chen, and G. J. Pottie, “Radio link admission algorithms for wireless networks with power control and active link quality protection,” in *IEEE INFOCOM*, 1995, pp. 97–104.

[57] D. Mitra and A. Morrison, “A distributed power control algorithm for bursty transmissions in cellular, spread spectrum wireless networks,” in *Proc. of 5th WINLAB Workshop, Rutgers University*, New Brunswick, NJ, 1995.

[58] C. Saraydar, N. Mandayam, and D. Goodman, “Pricing and power control in a multicell wireless data network,” *IEEE JSAC*, vol. 19, no. 10, pp. 1883–1892, October 2001.

[59] T. Javidi, “Decentralized rate assignments in multi-sector cdma networks,” in *IEEE GLOBECOM*, December 2003.

[60] M. Chiang and J. Bell, “Balancing supply and demand of wireless bandwidth: joint rate allocation and power control,” in *IEEE INFOCOM*, March 2004.

[61] D. O’Neill, D. Julian, and S. Boyd, “Adaptive management of network resources,” in *IEEE VTC*, 2003.

[62] M. Xiao, N. B. Shroff, and E. Chong, “A utility-based power control scheme in wireless cellular systems,” *IEEE/ACM Transactions on Networking*, vol. 11, no. 2, pp. 210–221, April 2003.

[63] P. Hande, S. Rangan, and M. Chiang, “Distributed uplink power control for optimal sir assignment in cellular data networks,” in *INFOCOM*, 2006.

[64] [Online]. Available: http://www.3gpp.org/ftp/tsg_ran/WG4_Radio/TSGR4_06/Docs/Pdfs/r4-99362.pdf

[65] R. Ramanathan and R. Rosales-Hain, “Topology control of multihop wireless networks using transmit power adjustment,” in *IEEE INFOCOM*, 2000, pp. 404–413.

[66] V. Rodoplu and T. Meng, "Minimum energy mobile wireless networks," *IEEE JSAC*, vol. 17, no. 8, pp. 133–1344, August 1999.

[67] R. Wattenhofer, L. Li, P. Bahl, and Y.-M. Wang, "Distributed topology control for power efficient operation in multihop wireless ad hoc networks," in *IEEE INFOCOM*, 2001, pp. 1388–1397.

[68] T. Elbatt, S. V. Krishnamurthy, D. Connors, and S. Dao, "Power management for throughput enhancement in wireless ad-hoc networks," in *Proc. of IEEE ICC*, 2000, pp. 1506–1513.

[69] T. Moscibroda and R. Wattenhofer, "Minimizing interference in ad hoc and sensor networks," in *ACM DIALM-POMC '05: Proceedings of the 2005 joint workshop on Foundations of mobile computing*, Cologne, Germany, 2005, pp. 24–33.

[70] M. Burkhart, P. V. Rickenbach, R. Wattenhofer, and A. Zollinger, "Does topology control reduce interference?" in *MobiHoc '04: Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing*. New York, NY, USA: ACM, 2004, pp. 9–19.

[71] S. Narayanaswamy, V. Kawadia, R. S. Sreenivas, and P. R. Kumar, "Power control in ad hoc networks: Theory, architecture, algorithm and implementation of the compow protocol," in *Proc. of European Wireless Conference*, 2002.

[72] V. Kawadia and P. R. Kumar, "Power control and clustering in ad hoc networks," in *IEEE INFOCOM*, 2003.

[73] V. Kawadia and P. R. Kumar, "Principles and protocols for power control in ad hoc networks," *IEEE JSAC*, 2005.

[74] P. Karn, "MACA - a new channel access method for packet radio," in *Proc. of 9th ARRL Computer Networking Conference*, 1990, pp. 134–140.

[75] M. B. Pursley, H. B. Russell, and J. S. Wysocarski, "Energy-efficient transmission and routing protocols for wireless multi-hop networks and spread spectrum radios," in *Proc. of EUROCOMM Conference*, 2000, pp. 1–5.

[76] E.-S. Jung and N. H. Vaidya, "A power control mac protocol for ad hoc networks," in *MobiCom '02: Proceedings of the 8th annual international conference on Mobile computing and networking*, Atlanta, Georgia, USA, 2002, pp. 36–47.

[77] J. Gomez, A. T. Campbell, M. Naghshineh, and C. Biskidian, "PARO: supporting dynamic power controlled routing in wireless ad hoc networks," *ACM/Kluwer Journal on Wireless Networks*, vol. 9, no. 5, pp. 443–460, 2003.

[78] S.-L. Wu, Y.-C. Tseng, and J.-P. Sheu, "Intelligent medium access for mobile ad hoc networks with busy tones and power control," *Proc. of IEEE JSAC*, vol. 18, no. 9, pp. 1647–1657, 2000.

[79] J. Monks, V. Bharghavan, and W.-M. Hwu, "A power controlled multiple access protocol for wireless packet networks," in *Proc. of IEEE INFOCOM*, 2001, pp. 219–228.

[80] A. Muqattash and M. Krunz, "Power controlled dual channel (pcdc) medium access protocol for wireless ad hoc networks," in *Proc. of IEEE INFOCOM*, 2003, pp. 470–480.

[81] Alaa Muqattash and Marwan Krunz, "A single-channel solution for transmission power control in wireless ad hoc networks," *IEEE JSAC*, 2005.

[82] X. Yang and N. Vaidya, "On the physical carrier sense in wireless ad hoc networks," in *Proc. of IEEE INFOCOM*, 2005.

[83] D. Qiao, S. C. A. Jain, and K. G. Shin, "Adaptive transmit power control in IEEE 802.11a wireless LANs," in *IEEE VTC (Spring)*, vol. 1, April 2003, pp. 433–437.

[84] I. Broustis, J. Eriksson, S. V. Krishnamurthy, and M. Faloutsos, "Implications of power control in wireless networks: A quantitative study," in *PAM*, 2007, pp. 83–93.

[85] M. Kubisch, H. Karl, A. Wolisz, L. C. Zhong, and J. M. Rabaey, "Distributed algorithms for transmission power control in wireless sensor networks," in *IEEE WCNC*, March 2003.

[86] D. Son, B. Krishnamachari, and J. Heidemann, "Experimental study of the effects of transmission power control and blacklisting in wireless sensor networks," in *Proc. of the First IEEE Conference on Sensor and Adhoc Communication and Networks*. Santa Clara, California, USA: IEEE, October 2004, pp. 289–298.

[87] T. He, S. Krishnamurthy, J. A. Stankovic, T. Abdelzaher, L. Luo, R. Stoleru, T. Yan, L. Gu, J. Hui, and B. Krogh, "Energy-efficient surveillance system using wireless sensor networks," in *MobiSys '04: Proceedings of the 2nd international conference on Mobile systems, applications, and services*. New York, NY, USA: ACM, 2004, pp. 270–283.

[88] J. Jeong, D. E. Culler, and J.-H. Oh, "Empirical analysis of transmission power control algorithms for wireless sensor networks," EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2005-16, Nov 2005. [Online]. Available: http://www.eecs.berkeley.edu/Pubs/TechRpts/2005/EECS-2005-16.html

[89] G. W.-Allen, K. Lorincz, M. Ruiz, O. Marcillo, J. Johnson, J. Lees, and M. Welsh, "Deploying a wireless sensor network on an active volcano," *IEEE Internet Computing, Special Issue on Data-driver Applications in Sensor Networks*, vol. 10, pp. 18–25, March 2006.

[90] J. Jeong, D. E. Culler, and J.-H. Oh, "Empirical analysis of transmission power control algorithms for wireless sensor networks," in *IEEE INSS*, 2007.

[91] S. Lin, J. Zhang, G. Zhou, L. Gu, J. A. Stankovic, and T. He, "Atpc: adaptive transmission power control for wireless sensor networks," in *SenSys '06: Proceedings of the 4th international conference on Embedded networked sensor systems*. New York, NY, USA: ACM, 2006, pp. 223–236.

[92] R. Draves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks," in *MobiCom '04: Proceedings of the 10th annual international conference on Mobile computing and networking*. New York, NY, USA: ACM, 2004, pp. 114–128.

[93] A. Balasubramanian, R. Mahajan, A. Venkataramani, B. Levine, and J. Zahorjan, "Interactive wifi connectivity for moving vehicles," in *SIGCOMM*, 2008.

[94] P. Hande, S. Rangan, M. Chiang, and X. Wu, "Distributed uplink power control for optimal sir assignment in cellular data networks," *IEEE/ACM Transactions on Networking*, 2008.

[95] A. Sheth and R. Han, "A Mobility-Aware Adaptive Power Control Algorithm For Wireless LANs," in *IEEE CAS Low Power Workshop*, 2002.

[96] A. Sheth and R. Han, "An Implementation of Transmit Power Control in 802.11b Wireless Networks," Dept. of Comp. Science, Univ. of Colorado, Tech. Rep., 2002.

[97] D. Qiao, S. Choi, A. Jain, and K. G. Shin, "Miser: an optimal low-energy transmission strategy for ieee 802.11a/h," in *ACM MobiCom*, 2003, pp. 161–175.

[98] P. Chevillat, J. Jelitto, and H. L. Truong, "Dynamic Data Rate and Transmit Power Adjustment in IEEE 802.11 Wireless LANs," *Intl. Journal of Wireless Information Networks*, 2005.

[99] V. Navda, R. Kokku, S. Ganguly, and S. Das, "Slotted Symmetric Power Control in managed WLANs, http://www.nec-labs.com/∼ravik/RESEARCH/contour.pdf," NEC Laboratories America, Tech. Rep., 2007.

[100] A. Sheth and R. Han, "Shush: Reactive transmit power control for wireless mac protocols," in *WICON '05: Proceedings of the First International Conference on Wireless Internet*. Washington, DC, USA: IEEE Computer Society, 2005, pp. 18–25.

[101] V. Shah, E. Gelal, and S. V. Krishnamurthy, "Handling asymmetry in power heterogeneous ad hoc networks," *Comput. Networks*, vol. 51, no. 10, pp. 2594–2615, 2007.

[102] J. Zander, "Distributed co-channel interference control in cellular radio systems," *IEEE Transactions on Vehicular Technology*, vol. VT-41, pp. 305–311, August 1992.

[103] J. Heidemann and W. Ye, "Energy conservation in sensor networks at the link and network layers," USC/Information Sciences Institute, Tech. Rep. ISI-TR-2004-599, 2004. [Online]. Available: http://www.isi.edu/ johnh/PAPERS/Heidemann04c.html

[104] J. Heidemann and W. Ye, *Wireless Sensor Networks: A Systems Perspective*. Artech House, Inc., 2005, ch. Energy Conservation in Sensor Networks at the Link and Network Layers, pp. 75–86, Chapter 6. [Online]. Available: http://www.isi.edu/ johnh/PAPERS/Heidemann05a.html

[105] W. Ye, J. Heidemann, and D. Estrin, "An energy-efficient mac protocol for wireless sensor networks," in *Proceedings of the IEEE Infocom*, USC/Information Sciences Institute. New York, NY, USA: IEEE, June 2002, pp. 1567–1576. [Online]. Available: http://www.isi.edu/ johnh/PAPERS/Ye02a.html

[106] "IEEE-SA Get IEEE 802.16 LAN/MAN Broadband Wireless LANs," http://standards.ieee.org/getieee802/802.16.html.

[107] "Technical Specificati on Group Radio Access Network: Physical Layer Aspects for Evolved Universal Terrestrial Radio Access (UTRA)," 3GPP TR 25.814 v.7.0.0 3GPP.

[108] "IEEE 802.15 WPAN Millimeter Wave Alternative PHY Task Group 3c (TG3c)," http://www.ieee802.org/15/pub/TG3c.html.

[109] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das, "Mobisteer: using steerable beam directional antenna for vehicular network access," in *MobiSys '07: Proceedings of the 5th international conference on Mobile systems, applications and services*. New York, NY, USA: ACM, 2007, pp. 192–205.

[110] R. Ramanathan, "On the performance of ad hoc networks with beamforming antennas," in *MobiHoc '01: Proceedings of the 2nd ACM international symposium on Mobile ad hoc networking & computing*. New York, NY, USA: ACM, 2001, pp. 95–105.

[111] R. R. Choudhury, X. Yang, R. Ramanathan, and N. H. Vaidya, "Using directional antennas for medium access control in ad hoc networks," in *MobiCom*. New York, NY, USA: ACM, 2002, pp. 59–70.

[112] A. Miu, H. Balakrishnan, and C. E. Koksal, "Improving loss resilience with multi-radio diversity in wireless networks," in *MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking*. New York, NY, USA: ACM, 2005, pp. 16–30.

[113] A. Viterbi, A. Viterbi, K. Gilhousen, and E. Zehavi, "Soft handoff extends cdma cell coverage and increases reverse link capacity," *IEEE JSAC*, vol. 12, no. 8, pp. 1281–1288, Oct 1994.

[114] S. Biswas and R. Morris, "Exor: opportunistic multi-hop routing for wireless networks," *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 4, pp. 133–144, 2005.

[115] S. Chachulski, M. Jennings, S. Katti, and D. Katabi, "Trading structure for randomness in wireless opportunistic routing," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 169–180, 2007.

[116] A. Miu, H. Balakrishnan, and C. E. Koksal, "Improving loss resilience with multi-radio diversity in wireless networks," in *MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking*. New York, NY, USA: ACM, 2005, pp. 16–30.

[117] Z. Tao, T. Korakis, F. Liu, S. Panwar, J. Zhang, and L. Tassiulas, "Cooperation and directionality: Friends or foes?" in *Communications, 2008. ICC '08. IEEE International Conference on*, 2008.

[118] V. Bychkovsky, B. Hull, A. K. Miu, H. Balakrishnan, and S. Madden, "A Measurement Study of Vehicular Internet Access Using In Situ Wi-Fi Networks," in *In MOBICOM*, Los Angeles, CA, September 2006.

[119] J. Ott and D. Kutscher, "A disconnection-tolerant transport for drive-thru internet environments," in *In IEEE Infocom*, March 2005.

[120] J. Eriksson, H. Balakrishnan, and S. Madden, "Cabernet: vehicular content delivery using wifi," in *MobiCom '08: Proceedings of the 14th ACM international conference on Mobile computing and networking*. New York, NY, USA: ACM, 2008, pp. 199–210.

[121] P. Rodriguez, R. Chakravorty, J. Chesterfield, I. Pratt, and S. Banerjee, "Mar: a commuter router infrastructure for the mobile internet," in *MobiSys '04: Proceedings of the 2nd international conference on Mobile systems, applications, and services*. New York, NY, USA: ACM, 2004, pp. 217–230.

[122] D. Hadaller, S. Keshav, T. Brecht, and S. Agarwal, "Vehicular opportunistic communication under the microscope," in *MobiSys '07: Proceedings of the 5th international conference on Mobile systems, applications and services*. New York, NY, USA: ACM, 2007, pp. 206–219.

[123] A. J. Nicholson and B. D. Noble, "Breadcrumbs: forecasting mobile connectivity," in *MobiCom '08: Proceedings of the 14th ACM international conference on Mobile computing and networking*. New York, NY, USA: ACM, 2008, pp. 46–57.

[124] P. Shankar, T. Nadeem, J. Rosca, and L. Iftode, "CARS: Context Aware Rate Selection for Vehicular Networks," in *IEEE ICNP*, 2008.

[125] R. Mahajan, J. Zahorjan, and B. Zill, "Understanding wifi-based connectivity from moving vehicles," in *IMC '07: Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*. New York, NY, USA: ACM, 2007, pp. 321–326.

[126] C. W. S. Gaudin, "Intel unveils new low-power chip/building block," http://tinyurl.com/4wkxdr.

[127] G. Bianchi, L. Fratta, and M. Olivieri, "Performance evaluation and enhancement of the CSMA/CA protocol for 802.11 wireless lans," IEEE PIMRC, vol. 2, Taipei, Taiwan, October 1996, pp. 392–396.

[128] S. Choi, K. Park, and C. Kim, "On the performance characteristics of WLANs: Revisited," in *Proc. of SIGMETRICS*, 2005.

[129] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proc. of IEEE INFOCOM*, vol. 2, 2003, pp. 836–843.

[130] J. Kim, S. Kim, S. Choi, and D. Qiao, "CARA: Collision-aware rate adaptation for IEEE 802.11 WLANs," in *Proc. of IEEE INFOCOM*, Barcelona, Spain, March 2006.

[131] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh, "Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols," Proc. of IEEE WCNC, March 2005.

[132] MADWiFI, "RSSI in MadWifi," http://madwifi.org/wiki/UserDocs/RSSI.

[133] Z. Wu, "Orbit traffic generator," http://www.orbit-lab.org/wiki/Documentation/OTG, WINLAB, Rutgers University.

[134] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer, "Understanding congestion in IEEE 802.11b wireless networks," in *Proc. of USENIX Internet Measurement Conference (IMC)*, 2005, pp. 279–292.

[135] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 3, pp. 535–547, May 2000.

[136] R. Jain, D. Chiu, and W. Hawe, "A quantitative measure of fairness and discrimination for resource allocation in shared computer systems," DEC Tech Report TR-301, Tech. Rep., September 1984.

[137] J. Bellardo and S. Savage, "802.11 denial-of-service attacks: Real vulnerabilities and practical solutions," in *Proc. of USENIX Security Symposium*, Washington D.C., 2003.

[138] A. Kochut, A. Vasan, A. U. Shankar, and A. Agrawala, "Sniffing out the correct physical layer capture model in 802.11b," in *Proc. of IEEE ICNP*. Washington, DC, USA: IEEE Computer Society, 2004, pp. 252–261.

[139] G. Tan and J. Guttag, "The 802.11 MAC Protocol Leads to Inefficient Equilibria," in *Proc. of IEEE INFOCOM*, Miami, FL, March 2005.

[140] P. Gopalakrishnan, P. Spasojevic, L. Greenstein, and I. Seskar, "A method for predicitng the performance of rate-adaptice WLANs," WINLAB, Rutgers University, Piscataway, NJ, Tech. Rep. TR-259, September 2004.

[141] D. Malone, P. Clifford, and D. J. Leith, "MAC Layer Channel Quality Measurement in 802.11," in *IEEE Communication Letters*, February 2007, vol. 11.

[142] M. Heusse, F. Rousseau, R. Guillier, and A. Duda, "Idle sense: an optimal access method for high throughput and fairness in rate diverse wireless lans," *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 4, pp. 121–132, 2005.

[143] S. Avallone, S. Guadagno, D. Emma, A. Pescape, and G. Ventre, "D-ITG distributed internet traffic generator," in *Proc. QEST 2004*, September 2004, pp. 316–317.

[144] K. Ramachandran, R. Kokku, H. Zhang, and M. Gruteser, "Synchronous Two-phase Rate and Power Control in 802.11 WLANs," http://www.nec-labs.com/∼ravik/RESEARCH/symphony-tech.pdf.

[145] [Online]. Available: http://bcm-v4.sipsolutions.net/802.11

[146] M. Buettner, E. Anderson, G. Yee, D. Saha, A. Sheth, D. Sicker, and D. Grunwald, "A phased array antenna testbed for evaluating directionality in wireless networks," in *ACM MobiEval*, 2007, pp. 7–12.

[147] Z. Yi and P. P. Waskiewicz, "Enabling Linux Network Support of Hardware Multiqueue Devices," in *Linux Symposium*. Red Hat Inc., June 2007.

[148] V. Shrivastava, D. Agarwal, A. Mishra, S. Banerjee, and T. Nadeem, "Understanding the limitations of transmit power control for indoor wlans," in *IMC*, 2007.

[149] R. Cle and J. Rosenbluth, "Voice over IP performance monitoring," *ACM Computer Communication Review*, vol. 31, no. 2, April 2001.

[150] "Phocus Array Antenna by Fidelity Comtech," http://www.fidelity-comtech.com/.

[151] "tcpdump." [Online]. Available: http://www.tcpdump.org/

[152] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, “Measurement-based models of delivery and interference in static wireless networks,” *ACM SIGCOMM CCR*, vol. 36, no. 4, pp. 51–62, 2006.

[153] K. Ramachandran, R. Kokku, K. Sundaresan, M. Gruteser, and S. Rangarajan, “R2D2: Regulating Beam Shape and Rate as Directionality meets Diversity, http://www.nec-labs.com/∼karthiks/papers/r2d2-tech.pdf,” NEC Laboratories America, Tech. Rep., 2008.

[154] S. J. Orfanidis, *Electromagnetic Waves and Antennas*. ONLINE BOOK: http://www.ece.rutgers.edu/ orfanidi/ewa/, 2004.

[155] “WiMAX Antennas Primer - A guide to MIMO and beamforming,” http://www.wimax.com/commentary/spotlight/wimax-antennas-primer-a-guide-to-mimo-and-beamforming/.

[156] “WiMAX Forum.” [Online]. Available: http://www.wimaxforum.org/home/

[157] “UE Radio Transmission and Reception (FDD),” 3GTS 25.1010 3GPP.

Curriculum Vita

KISHORE RAMACHANDRAN

EDUCATION

- Ph.D., Computer Engineering, Jan. 2009
  Rutgers University, New Brunswick, NJ., USA
- M.S., Computer and Systems Engineering, May 2002
  Rensselaer Polytechnic Institute, Troy, NY., USA
- B.E., Electronics and Telecommunications Engineering, July 2000
  V.E.S. Inst. of Technology, University of Mumbai, Mumbai, India

WORK EXPERIENCE

- WINLAB, Rutgers University, Piscataway, NJ., 06/2003 - 12/2008
  Graduate Research Assistant
- NEC Laboratories America, Inc., Princeton, NJ., 05/2008 - 08/2008
  Summer Intern
- NEC Laboratories America, Inc., Princeton, NJ., 05/2007 - 09/2007
  Summer Intern
- Bell Laboratories, Lucent Technologies, Holmdel, NJ., 06/2005 - 08/2005
  Summer Intern
- Formation Inc., Moorestown, NJ., 07/2002 - 02/2003
  Member of Technical Staff
- Electrical and Systems Engineering, RPI, Troy, NY., 07/2001 - 06/2002
  Graduate Research Assistant

PUBLICATIONS

- K. Ramachandran, R. Kokku, H. Zhang, and M. Gruteser. Symphony: Synchronous Two-phase Rate and Power Control in 802.11 WLANs. In the 6th International Conference on Mobile Systems, Applications and Services (MobiSys), Colorado, USA, June 2008. [Acceptance Rate: 18%]
- M. Blanco, R. Kokku, K. Ramachandran, S. Rangarajan, and K. Sundaresan. On the Effectiveness of Switched Beam Antennas in Indoor Environments. In the 9th Passive and Active Measurement (PAM) Conference, Ohio, USA, April 2008.

- K. Ramachandran, M. Gruteser, R. Onichi and T. Hikita. Experimental Analysis of Broadcast Reliability in Dense Vehicular Networks. In the 1st IEEE Symposium on Wireless Vehicular Communications (WiVec), Baltimore, MD., September 2007. Among the top 5 papers forwarded to IEEE Vehicular Tech. Magazine.

- K. Ramachandran, H. Kremo, M. Gruteser, P. Spasojevic and I. Seskar. Scalability Analysis of Rate Adaptation Techniques in Congested IEEE 802.11 Networks: An ORBIT Testbed Comparative Study. In the 8th IEEE Symposium on a World of Wireless Mobile and Multimedia Networks (WoWMoM), Helsinki, Finland, June 2007. [Top 10% of all submitted papers].

- S. Kaul, K. Ramachandran, P. Shankar, S. Oh, M. Gruteser, I. Seskar and T. Nadeem. Effect of Antenna Placement and Diversity on Vehicular Network Communications, In the 4th IEEE Communications Society Conference on Sensor, Mesh, and Ad Hoc Communications and Networks (SECON), San Diego, CA., June 2007. [Acceptance Rate: 20%].

- K. Ramachandran, S. Rangarajan and J. C. Lin. Make-Before-Break MAC Layer Handoff in 802.11 Wireless Networks. In the 41st IEEE Intl. Conference on Communications (ICC), Istanbul, Turkey, June 2006.

- D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu and M. Singh. Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols. In the 7th IEEE Wireless Comm. and Networking Conference (WCNC), LA., USA, March 2005.

- M. A. Ergin, K. Ramachandran, and M. Gruteser. An Experimental Study of Inter-cell Interference Effects on System Performance in Unplanned WLAN Deployments. In Computer Networks (Elsevier), 2008.

- K. Ramachandran, M. Gruteser, R. Onichi and T. Hikita. Experimental Analysis of Broadcast Reliability in Dense Vehicular Networks. In IEEE Vehicular Technology Magazine, December 2007.

- Y. Xia, D. Harrison, S. Kalyanaraman, K. Ramachandran and A. Venkatesan. Accumulation-based Congestion Control. In IEEE/ACM Transactions on Networking, February 2005.

- S. Ganu, K. Ramachandran, M. Gruteser, I. Seskar and J. Ding. Methods for Restoring MAC Layer Fairness in IEEE 802.11 Networks with Physical Layer Capture. In the 2nd International Workshop on Multi-hop Ad Hoc Networks: from Theory to Reality (REALMAN, held with MobiHoc), Florence, Italy, May 2006. [Acceptance Rate: 18%].

- K. Ramachandran, S. Kaul, S. Mathur, and M. Gruteser. Towards Mobility Emulation Through Spatial Switching on a Wireless Grid. In E-WIND: Workshop on Experimental Approaches to Wireless Network Design and Analysis (held with SIGCOMM), Philadelphia, PA, August 2005.