A Multi-Tier Network Architecture for Long Distance Rural Wireless Networks in Developing Regions

Rabin Krushnachandra Patra

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2009-92
http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-92.html

June 10, 2009

A Multi-Tier Network Architecture for Long Distance Rural Wireless Networks in Developing Regions

by

Rabin Krushanchandra Patra

B.Tech. (Indian Institute of Technology, Kharagpur) 2002

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Computer Science

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:
Professor Eric Brewer, Chair
Professor Jean Walrand
Professor AnnaLee Saxenian

Spring 2009

The dissertation of Rabin Krushanchandra Patra is approved:

Chair Date

Date

Date

University of California, Berkeley

Spring 2009

A Multi-Tier Network Architecture for Long Distance Rural Wireless Networks in Developing Regions

Abstract

A Multi-Tier Network Architecture for Long Distance Rural Wireless Networks in Developing Regions

by

Rabin Krushanchandra Patra

Doctor of Philosophy in Computer Science

University of California, Berkeley

Professor Eric Brewer, Chair

Today, in spite of technological advances, large disparities still exist between the industrialized world and the developing world. Information and communication technologies, in particular, have immense potential to provide better health, education and economic opportunities to people in the developing world, but most currently available communication solutions are often not financially viable in rural areas.

We believe that wireless infrastructure that starts off by offering connectivity to targeted locations is the most practical way to extend communication coverage to rural areas. To that end, we propose WiLDNet, a new multi-tier network architecture composed of high-bandwidth long-distance point-to-point backbone wireless links and medium-distance point-to-multipoint access links. We also propose to use cheap and widely available off-the-shelf WiFi-based radios as the base technology to build this architecture. Unfortunately,

standard WiFi radios suffer from low throughput, high packet loss rates and poor spectrum efficiency in the real world long-distance environments.

In this dissertation, we demonstrate the feasibility of the WiLDNet architecture. We first characterize the underlying problems behind the poor performance of WiFi in long-distance scenarios. To overcome these problems, we build WiLDMAC, a novel time-division based MAC-layer that increases channel utilization and eliminates packet collisions at long distances. To achieve high end-to-end multi-hop throughput in the point-to-point backbone of our architecture, we use a combination of ARQ and FEC-based loss recovery mechanisms. Our measurements show 2–5 fold improvement in throughput on single-hop links as long as 382km. To achieve high capacity scaling and to support dynamic traffic demands in the point-to-multipoint part of our architecture, we implement and evaluate three techniques – a) dynamic power adaptation to minimize interference and maximize spectrum usage b) dynamic channel width adaptation to increase the number of simultaneous clients and c) physical antenna combination to decrease the cost of installation of base-stations.

Finally, we deploy our long-distance wireless links in several real world networks. Our rural telemedicine network in Aravind Eye Hospitals, India that connects ten village clinics, has already enabled 90,000 remote video consultations, showed operational sustainability and is in the process of expanding coverage to treat 500,000 people/year.

Professor Eric Brewer
Dissertation Committee Chair

To my parents.

Contents

List of Figures

List of Tables

Acknowledgments

I will first start off by saying that this dissertation would never have happened without the contributions of numerous people – some, who I know personally, but there are many others, who I will never know. I have realized how incredibly fortunate I am to have been able to come to Berkeley, to get associated with the TIER research group, and to know such outstanding and brilliant people over the course of my Ph.D. Indeed, I could not have dreamt of a better graduate student experience than working in TIER, to research some really challenging problems, while also being aware of the real people and their real problems that we were ultimately trying to solve.

For making all this possible, I first of all thank Eric Brewer, my adviser and the visionary who kick-started a new collaboration between computer scientists and social scientists to work on technology and development in Berkeley. He has been a great adviser who supported me both financially and intellectually over the last six years. Eric has an uncanny ability to master various fields, put diverse concepts together into a big picture, and come up with ideas that makes everybody wonder – why did anybody not think of that before?

My work and personal life in Berkeley would not have nearly been as enriching or fun without our own band of the three musketeers (Sergiu, Joyojeet and I), and I would not have nearly become the same person as I am today. We have had some great adventures together, and I would like to thank both of them for always being there for me as friends, colleagues and co-conspirators.

I met Sergiu on my first day in my office in Soda Hall, and was I immediately

struck by this dapper Romanian, who is still probably the most well dressed man in the whole Computer Science department. One of smartest and talented computer scientists I know, his spirit is as competitive as his heart is kind. All throughout the last six years, we have worked together on numerous class projects, prelims, research proposals, dissertations and now, on realizing our vision to start our own company. Outside work, he has been my permanent guide who taught me varied skills ranging from skiing, skating, tennis, ice-climbing, windsurfing and car repairing (with less enthusiasm from me) among others. We have traveled together to the plains of India, to the tops of some of the most scenic mountains and also underwater to see some of the most stunning coral reefs. Joyojeet, who is one of the most perceptive and insightful social scientists I have met, has been more than a friend to me – he has often been like the bullying big brother, or like the doting father who never hesitated to put me to task whenever I went my wayward ways (which was often). His life experiences have never ceased to amaze me, are definitely going to include writing a few books, making a few movies or even acting in them!

The greatest thing about all the TIER people is that not only are they outstanding collaborators, we were also best friends together. The high-point of the TIER days were our travel trips – starting from the unforgettable hot summer trip to India with Renee, Jen, Arianna, Melissa, Sergiu, Joyojeet, Kevin, Bowei and Eric, to subsequent trips to India for more wireless work, to my trip to Venezuela with Sonesh.

A constant companion through my years at Berkeley has been Sonesh, one of the most articulate and clear thinking persons I know, and I thank him for playing a crucial role to realize the potential of our work and for all the years together, since the time we shared neighboring offices in Soda Hall, to the times when we clumsily climbed towers in India, to when we both despaired over India loosing in cricket to Australia in Mohali, to when we celebrated our world record wireless link in Venezuela. I have had many memorable times with Melissa, the intrepid “ICTD chick”, whether we were traveling on scary Indian roads or writing papers late through the nights at Intel. Someday, we might conclude our yet unfinished philosophical discussions about our place in the universe. I would probably never have graduated from Berkeley, if I had not come to know Lakshmi. A multi-faceted personality, Lakshmi joined our group at the right time, gave momentum and direction to our work, and taught us how to plan our research, to write about it in a compelling manner and to present it with confidence. There are few people I know who are more capable or colorful than Yahel, indeed a “fountain of knowledge”, and I hope I will be climbing more towers (or rock-faces) and braving more high waves to more islands (after Farallones) to erect more “pucca” wireless installations (or diving under water) in the times to come. There are also few people who are as incisive or pragmatic as Sushant, who taught me not only how to write cast-iron papers but also how one can be reckless by trying to climb Mt. Shasta in one day or to jump from a plane at 15,000 ft, but nevertheless, emerge without a scratch. I greatly appreciate my friendship with Divya, who has the uncanny ability to come up with ever more creative ways of mangling my name. I tried to learn Tamil from her, but she is the one who mastered Hindi in the end. I thank Janaki also, who opened my eyes to a whole new world of long-distance train travel (not wireless) in the US. Of course, my TIER experience cannot be accounted fully without mentioning the other jokers – Bowei, who always managed to dig up the most obscure stories, Mike, who fed me the wonders

of Kashi, and RJ, who I could never beat in any kind of board game. I am happy that I also had the chance to work with Anmol Sheth, Omar Bakr, Michael Demmer, Rowena Luk, Renee Kuriyan, Jen Bussell, Madelaine Plauche, Matt Podolsky, Manuel Ramos and Sandra Salmeron among others.

I spent most of my PhD years at the Intel research labs, and I thank Intel and in particular my good friends there – Kevin, Sylvia, Gianlucca and Paul among others for providing us support and help with our research. I spent my first few years as a PhD student in the unmistakable Soda hall, and I thank all my architecture buddies - Nemanja, Mark, Sukun, Yatish, Steve and Yury for providing all the moral support a new graduate student needs. I especially thank the people who were always there behind the scenes – Amrita (at Intel) for handling our equipment and refreshment requests, La Shana (at Soda) for guiding me through all my requirements and Damon (at Soda) for bearing with my perennially late reimbursement requests.

Life in Berkeley would not have been the same without my other friends – Rodrigo and Paula who went on some great boat and train trips with me in India, Shaji and Indu who climbed the peaks of New Mexico with me, Roycie and Rosemary who taught me the most effective dance moves. I was also privileged to know JK, Matulya, Biswo, Akhila, Dave, Rajesh, all my ex-housemates and my friends from IIT Kharagpur.

Finally, I would forever be thankful to my parents for making me what I am. It is some times easy to forget that but for their vision and enterprise, I would still be living in a sleepy town in rural Orissa. They were always supportive and patient with me while I went on interminably (as it looked to them) with my PhD. My brother, Dhiren, I thank you for always keeping tabs on my research and even helping out with some of our field work in India!

In the end, I am grateful to live in the wonderful city of Berkeley, with its myriad idiosyncrasies and beautiful people, which never made me feel like an outsider, and which I have made my home.

Bibliographic Notes

This dissertation is based on research that I have performed between 2002 and 2008 with a number of outstanding colleagues. Part of the material presented in Chapters 3 and 4 appear in three papers [104, 129, 135] co-authored with Sergiu Nedevschi, Sonesh Surana, Anmol Sheth, Lakshminarayanan Subramanian and Eric Brewer, and also appears in the dissertations of Sergiu Nedevschi [92] and Sonesh Surana. Some of the work on dynamic power adaptation, antenna combination and variable width channels presented in Chapter 5 appears in three papers [47, 48, 105] co-authored with Ramakrishna Gummadi, Sergiu Nedevschi, Sonesh Surana, Eric Brewer and Hari Balakrishnan. The description of the deployment of our rural telemedicine network in Chapter 6 appears in a paper [137] co-authored with Sonesh Surana, Sergiu Nedevschi and Eric Brewer.

Chapter 1

Introduction

Today, we are in the 21st century, but in spite of tremendous advances, the world still faces many challenges. The biggest challenge is the persistent existence of widespread poverty around the world. This is further exemplified by the stunning disparity between the industrialized countries and the developing countries in all the key indicators of development - average income, mortality rates, education levels, life expectancy and quality of life. Table 1.1 clearly shows that the 27 richest countries in the world (OECD countries) outperform the 50 poorest countries by overwhelming margins in many measures of development.

Information and communication technologies have been key drivers for rapid economic growth in the industrialized nations over the latter part of the last century, leading to vast improvements in the quality of life for their citizens. The digital revolution, including the Internet, has completely transformed our way of life, by increasing access to information, simplifying communication and collaboration across countries and continents, and improving productivity and business efficiency.

It is in this context that the role of information and communication technologies (or ICTs) in development has been widely recognized. The *Millennium Development Goals* defined by United Nations to spur development in the world's poorest countries by improving social and economic conditions explicitly include spread of ICTs as an important aim: "make available the benefits of new technologies — especially information and communications technologies" (ICT) [143].

The key question is, how can we use ICTs to create better quality of life and economic growth in developing countries? Unfortunately, the mechanics of economic development are such that they often reinforce existing disparities. Only nations that already have advanced technological societies are able to best use and absorb further innovations to increase productivity and efficiency, thus maintaining and even widening the gap between the "haves" and the "have-nots".

Attempts that try to transplant technologies that were designed for and sold in the industrialized world directly in the markets of the developing countries often fail because they ignore the stark differences in the ground realities with respect to purchasing power, usage models and supporting infrastructure (power, roads) that exist in these environments.

The aim of this dissertation is to address one aspect of this challenge – *how do we build affordable communication networks in developing regions by adapting technologies originally designed for the industrialized world to be appropriate and sustainable in the developing world environments?*

We propose an architecture consisting a high-bandwidth backhaul network of long-distance wireless links combined with medium-distance wireless links for local distribution

Development indicator	High income OECD countries	Least developed countries
Population (million)	965	799
GDP per capita in US$ (PPP)	36,657	1173
Life expectancy (years)	79	55
Child mortality rate, under 5 years (per 1000)	6	141
Internet users (per 100 people	65.5	1.8
Phone penetration (per 100 people)	145	10

Table 1.1: Disparity in key developmental indicators between the 27 high income OECD countries and the 50 least developed countries in the world. Source: World Development Indicators 2008, World Bank [152]

to end users. By using off-the-shelf, mass produced WiFi technology, we have brought down the cost of each link to be less than US $800. The complete solution includes innovative network algorithms to optimize the performance and robustness of the network as well as the ease of installation and remote management.

This dissertation's work was done jointly with other members of the Technology and Infrastructure for Emerging Regions (TIER) research group at UC Berkeley [1] who worked on various aspects of this network architecture. The specific focus of this dissertation is to create novel MAC-layer mechanisms implemented on off-the-shelf WiFi hardware to achieve high performance and reliability in long-distance environments.

In addition to developing the technologies, we also spent considerable efforts on deploying and testing them in real world environments. Right now, the principal deployments of our long-distance wireless links include the Aravind telemedicine network for doctor-patient video-conferencing [10] and other links in Uganda, Ghana and the Philippines.

In the rest of this chapter, we start off with the motivation behind our work on ICTs for developing countries (Section 1.1). Since the focus of this dissertation is on communication technologies, we provide an overview of their potential in socio-economic development (Section 1.2), the current gaps in availability of communication (Section 1.3) and the recent trends that have given a fresh impetus to innovation (Section 1.4). We then articulate the goal and contributions (Section 1.6) of this dissertation as part of a new WiFi based connectivity solution for rural areas of developing regions (Section 1.5).

1.1 Information and Communication Technologies for Development

The relationship between technology and development has long been a principal subject of academic research. With the advent of the digital revolution and the Internet, there has been a renewed interest in the potential of ICTs in development. This included work on the new ideas of a networked society from the academic community [25, 94], and also by international development agencies [42, 51, 97, 148].

This interest has been partly driven by the rapid growth of the high-tech software services industry in some developing countries (like India) where it has been looked at as important engine for export-led economic growth. For instance, in India, the highly public image of the high-tech industry has also lead to the emergence of aspirational value associated with computer literacy among people at all levels of society. The ability to use computers has acquired immense symbolic value that is more tied to social and economic ascendancy than actual functional use [101].

In addition, management gurus like C.K. Prahalad have theorized on the "bottom of the pyramid" markets [108] suggesting that companies choosing to design products and value chains for currently untapped markets can also develop viable profit making business models. Thus the promise of development through ICTs need not be just through philanthropic ventures.

Overall, it was a burst of interest among engineers and scientists in issues of technology in the developing world that created a significant push and a subsequent slew of projects in this arena. As a consequence, by the 1990s, there were already a multitude of ICT initiatives ranging from pilot projects to large-scale deployments from international aid agencies, local governments and non-governmental agencies. Indeed, anecdotal evidence suggests that there is positive economic benefit from access to technology and communication in particular. For instance, the World Bank's *infoDev* site catalogs hundreds of ICT projects [151].

These ICT initiatives target various areas of development ranging from healthcare, agriculture, education, governance and business development. In healthcare, the use of ICTs have been explored to bring telemedicine to remote areas [21], to develop low-cost diagnosis equipment and for improved information gathering and patient monitoring [58, 81]. In agriculture, ICTs have been proposed for information sharing among farmers [43], to provide price information [54, 75] and for efficient supply chain management [68]. In the field of education, ICTs have been used to improve learning among children through educational games on low-cost computers [107, 109], cellphones [73, 79] and shared computing in schools [88, 147]. ICTs have also been used to improve transparency in governance through e-payments and e-voting [52, 76, 117] and to enable small entrepreneurship through micro-credit [102].

However, the common attribute of some these early initiatives was the over reliance on existing off-the-shelf technology first developed for the users and markets of the industrialized world without proper assessment of what solutions were appropriate and needed. Although well intentioned, a significant number of these initiatives remained in the pilot stage implementation stage and failed to evolve into sustainable models that could be replicated or scaled up to have widespread impact.

These early experiences from the 1990s led to subsequent academic work looking at the impact of ICTD projects in India and elsewhere, to identify and learn best practices and apply them in future efforts [53]. There was an increase in interest starting in the early 2000s with several leading academics from key universities proposing ICT for Development (ICTD) as a new area of multi-disciplinary research that combined social scientists, educationists, economists, public health professionals and technology researchers from various fields such as Computer Science, Electrical, Mechanical and Civil Engineering. Several research groups were formed at UC Berkeley (TIER research group) [22], University of Washington [102], MIT, Carnegie Mellon University, IIT Kanpur [35] and IIT Chennai [64] among others to work in the area of technological innovation for the needs of the developing world. In addition, concrete international research partnerships in technology and development emerged, a notable example being MIT's Media Lab Asia, which set up shop in India. Also around this time, research wings in major technology groups, including Hewlett Packard Labs and Microsoft Research, set up establishments in the developing world, which naturally prompted thinking about application ideas relevant to the needs of these markets, adding even further momentum.

1.2 Communication Technologies in Developing Regions: The Potential

Although reliable communication infrastructure is taken for granted in the industrialized world, good quality communication facilities are widely unavailable in most parts of the developing world. This is akin to the unavailability of other basic infrastructure components such as roads, power and water in developing countries. Needless to say, affordable communication infrastructure has been recognized as a critical driver that could enable a large number of other applications.

The most visible of communication technologies is cellular telephony which seen tremendous growth in the past few years. With 1.15 billion phones sold worldwide in 2008, cellphone penetration has grown fastest in the poorest countries of the world. The flagship example of the spread of cellphones is Grameen Telecom (GT) [46], where women from villages could use microcredit to buy a cellphone and then operate it as a pay phone franchisee for their neighbors. At its peak, Grameen Telecom had 95,000 franchisees, who covered over 50,000 of the 68,000 villages in Bangladesh, and some 60 million people.

Apart from satisfying basic communication needs (phone calls and Internet access), communication can be effectively used for more targeted development oriented outcomes also. In healthcare, a number of research efforts have tried to build networks to enable telemedicine [21], providing video conferencing to remote rural areas with no doctors. Because of a lack of existing connectivity options, people have used either satellite (VSAT) links (India's space research organization, ISRO has been in this area for over 20 years [122]) or high-bandwidth long-distance wireless links. Other efforts that involve information gathering for epidemiological research or healthcare aid impact assessment [27, 58] would also benefit from connectivity options, even if they are intermittent. High-bandwidth connectivity would also enable remote learning and vocational training using interactive virtual classrooms in rural areas [67]. ITC's e-choupal project that introduced direct marketing to farmers to sell their produce used satellite terminals (VSAT) to connect its kiosk in villages [144]. A number of state governments in India have started projects to build village centers that offer services such as payment of bills, access to land records and health information to citizens. Various technologies such as VSAT and wireless have been used to connect the village centers together [8, 99]. Inveneo uses long-distance wireless links to connect refugee camps together in Uganda [66].

1.3 Communication Technologies in Developing Regions: The Gap

Although the potential of communication technologies is well acknowledged, there still are vast areas and populations, mostly in the developing world, that are still underserved.

Cellphones have seen fantastic growth rates, but the penetration of cellphones in the least developing countries is still at a low 10 out of 100 [152]. The picture is even more striking if we look at the difference between rural and urban areas. For example, in India while there are about 45 cellphones for every 100 people in urban areas, only 4.5 people (out of every 100 people) have cellphones in rural areas (in 2007). In sub-Saharan Africa, the penetration of cell phones into rural areas is even lower.

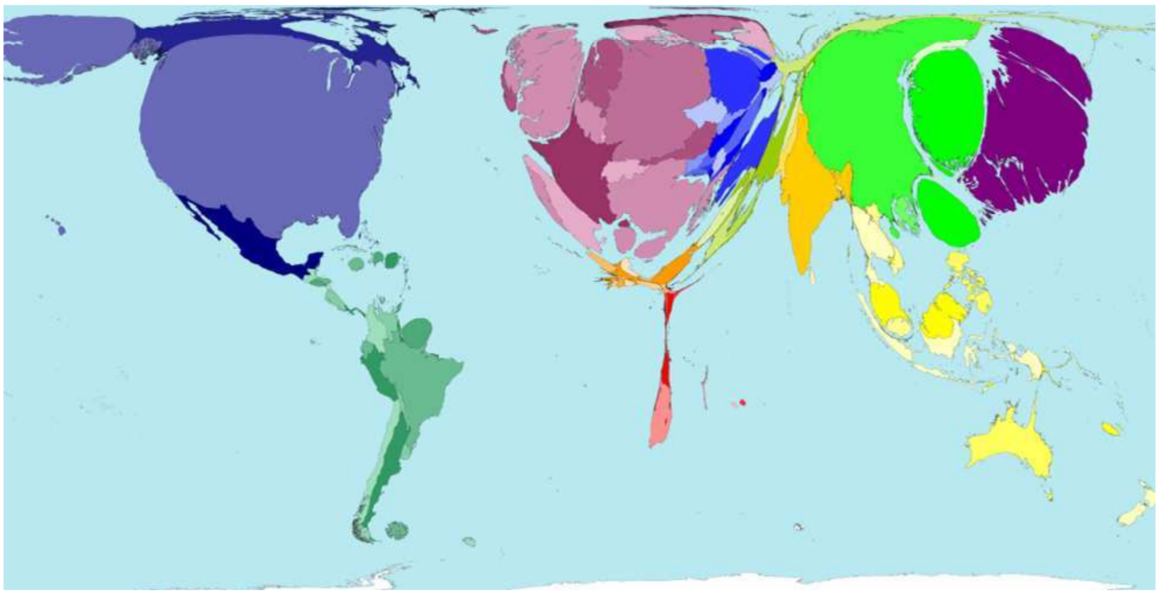

Figure 1.1: World map showing the number of Internet users in various countries. The area of a country is proportional to the number of Internet users in the country. Source: http://www.worldmapper.org

A similar situation exists with Internet usage. Figure 1.3 shows a modified map of the world where the area of each country is drawn to be proportional to the number of Internet users in that country. Not surprisingly, this map is remarkably similar to a map that uses average GDP per capita instead. In fact, the total number of Internet users in the whole of sub-Saharan Africa is just 35 million, which is about the same as the number of Internet users in South Korea, a country with just one-sixteenth the population.

The reasons for these kinds of disparities are many. More often than not, most rural areas in developing countries simply do not have required density of users or purchasing power capacity that can support cellular or other traditional Internet infrastructure. In addition, in many developing countries, because of archaic regulations, existing monopolies can often restrict access to network infrastructure (e.g. fiber) or charge prohibitive fees.

The default choice for connectivity in industrialized countries are wired solutions such fiber or copper wire. However, the high costs of laying the wires (at least US $1000/km in India for laying fiber) are not economically justified in the absence of high initial demand from users. The total cost of installing a phone line in the US is about US $500 which is viable given that more than 90% of households can afford US $30 a month on telephone service. However, in India more than 60% of households can afford to spend at most US $5 on communications every month [71].

Satellite networks provide fantastic coverage, but are very expensive. VSAT equipment installation costs over US $10,000 with a recurring monthly cost of over US $2,000 for a 1 Mbps link. In low user-density regions, VSAT is affordable only for businesses or wealthy users [103].

Networks with a base-station model such as WiMAX, and cellular networks like GPRS and CDMA, have an asymmetric design philosophy where expensive base-stations are amortized by large number of cheap clients over many users. In low-density regions, such base-stations simply do not cover enough users to be economically viable. Also, in a typical network, more than 70% of the cost is in the access network, not in the backbone.

The expectation that cellular solves the connectivity problem for developing regions is thus somewhat of a myth: cellular success in developing countries is an urban phenomenon, with a few exceptions. Bangladesh has good rural coverage because it is actually a very high density country, and base-stations that cover roads and rail lines also cover many villages. China has dictated good coverage as policy, despite the economic issues.

Other countries either subsidize rural users through taxation, much like the US universal access tax, or require some rural coverage as part of spectrum allocation. Hence urban areas tend be covered by multiple carriers, while rural areas typically have one or none. In its intended deployment model, with expensive base-stations covering many users, WiMax also shares the shortcomings of other cellular technologies. We expect that cellular and WiMAX networks will expand over time, but there is an immediate need for new rural connectivity options.

Finally, the existing business and deployment models for connectivity often ignore actual usage patterns and requirements. Providing network access does not necessarily have to be associated with Internet access. In many developing regions, basic local communications infrastructure is absent. A wireless network within a city or a district can enable a wide range of applications including telephony, essential services and health care. Solutions that focus on licensed spectrum and carrier-based deployment limits their usefulness to the kind of *grass roots* projects typical for developing regions.

1.4 Communication Technologies in Developing Regions: The Opportunity

In spite of the vastly inadequate network coverage in developing regions there are several indicators that point towards new opportunities today – the *impact of Moore's law*, the *growth of wireless communication* and a more *supportive business environment.*

- *Moore's Law:* The impact of Moore's law has been felt across all aspects of the technology industry. The exponential growth in the density of integrated circuits

(ICs) have brought the cost of computing down substantially. This cost reduction applies mostly to infrastructure that be shared such as servers and data centers, but it also partly applies to personal devices such as personal computers (PCs), laptop computers and smartphones. A PC can be bought in India for less than US $200 today. Full featured but small form factor laptops (called netbooks) such as the Asus Eee PC are available at less than $300. The same scaling advantages have also led to projects like the Intel Classmate PC [65] and the OLPC laptop [109] targeted towards children's education. With a focus on shared server infrastructure, the cost of computing and storage becomes realistic even for the poorest users.

- *Spread of wireless communication:* The second trend is the high-volume production of wireless communication equipment such as cellular and WiFi, which has brought the cost of devices down. This opens up the possibility of adapting them to be used in developing regions as well.

  WiFi technologies in particular have been used both for long-distance links using directional antennas and for local access with devices such as WiFi phones as well. There has been a slew of research on optimizing the performance of WiFi links [18, 19, 31, 56, 77, 86, 119, 153], and providing quality of service guarantees [9, 83, 119, 141]. We will cover WiFi technologies and related innovations in more detail in Chapter 2.

  In the non-WiFi market, a number of manufacturers like Airspan [6] and WiLAN [149] produce outdoor wireless equipment that work in the same unlicensed spectrum as WiFi. The TeNet group at IIT Madras developed corDECT building upon low-cost cordless phone technology to operate in wireless local loop providing data and voice services [72].

There are also other efforts to manufacture lower cost and lower power cellular base-stations that can be used to coverage targeted areas such as office buildings or villages. Such base-stations called micro cells [145], pico cells [90] or femto cells [7] are a few orders of magnitude cheaper than full fledged cellular base-station solutions. They also have much lower operating cost by obviating the need for expensive power backup solutions and by replacing expensive microwave backhauls with lower cost IP based wireless links. For example, GSM pico cells built by nanoGSM [90] transmit 23 dBm of power (200 mW), and can provide coverage in multi-story office environments, hospitals or small villages [90]. Even smaller femto cells such as ones built by 3Way [7] support 3G and can even use a standard DSL connection as backhaul.

- *Supportive business environment:* The third important trend is that the worldwide diffusion of technology and the growing access to capital have created a favorable environment for entrepreneurship and experimentation. This environment, combined with the success of franchising as a way to deploy large-scale ICT projects, means that there is a viable path from research to large-scale impact.

All these trends have not gone unnoticed by the research and business community. Early projects that took advantage of these opportunities include the Akshaya network [8] deployed by the Indian government in 2004 to connect more than 600 kiosks in a district of 3.5 million people to provide computer literacy, e-governance and other services. The network used a two-tier architecture consisting backhaul links (upto 30 km long) and local access links (between 1-10 km) and used proprietary outdoor wireless technologies [6, 149].

Other projects such as AirJaldi [5], CRCNet [34] and EHAS [38] used standard WiFi equipment to build community wireless networks.

1.5 WiLDNet Network Architecture

We believe that for rural areas in developing countries, wireless infrastructure appears to be the first kind of infrastructure that is affordable. We hypothesize that successful wireless infrastructure may lead to sufficient increases in rural incomes to make other infrastructure investments viable, such as water and power distribution.

Thus, we argue that for the low density of users typically seen in rural areas, approaches that provide full coverage are not feasible. The right strategy is to cover only those few places where connectivity is required, by employing long-distance wireless links. Such links can rely on WiFi, WiMax or other technologies that can offer reasonable throughput in both point-to-point and point-to-multipoint configurations.

Until now, for practical and cost-related reasons, we have chosen to use WiFi-based Long Distance (WiLD) links. WiFi radio cards are cheap and highly available, enjoying economies of scale. In our existing WiLD deployments, the cost of a WiLD link is approximately US$800 (excludes the cost of tower) with no recurring cost. Because they operate in unlicensed spectrum, WiLD links are easy to deploy and experiment with, and spectrum license costs are eliminated. With the latest WiFi standards, 802.11a and 802.11g that use OFDM modulation, we can achieve theoretical bandwidth of up to 54 Mbps. By using directional antennas and high-power radio cards, we have demonstrated that direct line of sight wireless links as long as 382 km can be achieved [39]. Further, manufacturers of

WiFi chipsets (e.g. Atheros) often support open-source drivers, allowing us to completely subvert the stock 802.11 MAC protocol and tailor the protocol to meet our needs.

In this section, we first provide a basic overview of WiLDNet, our proposed multi-level network architecture for providing connectivity to rural regions using WiFi as the main technology platform. We propose this architecture as an alternative to current solutions such as cellular and wireline networks. The main challenge in building this architecture comes from the fact that despite high theoretical bandwidth, currently available off-the-shelf WiFi devices fail to provide good performance and reliability in rural settings. This is primarily because the standard WiFi protocols were designed for short-distance broadcast environments inside offices and homes. Subsequently, we identify the research challenges at different layers of the network stack to increase the performance and reliability of WiFi based long-distance networks, and describe our efforts in addressing these challenges.

1.5.1 Design Principles

The requirements and operating constraints for networking systems in rural areas of the developing world are sometimes very different than the conditions in the developed world. Some of the key principles we consider in our WiLDNet architecture are as follows.

Low cost: Our solution should have both low capital and low operational costs. Capital cost includes both equipment cost and installation cost. Operational cost is often neglected when people fail to consider the overhead of providing stable power, and replacement of equipment and software maintenance. Finally, given the low purchasing power and sparsity of users, it is essential that the network can operate with low usage charges.

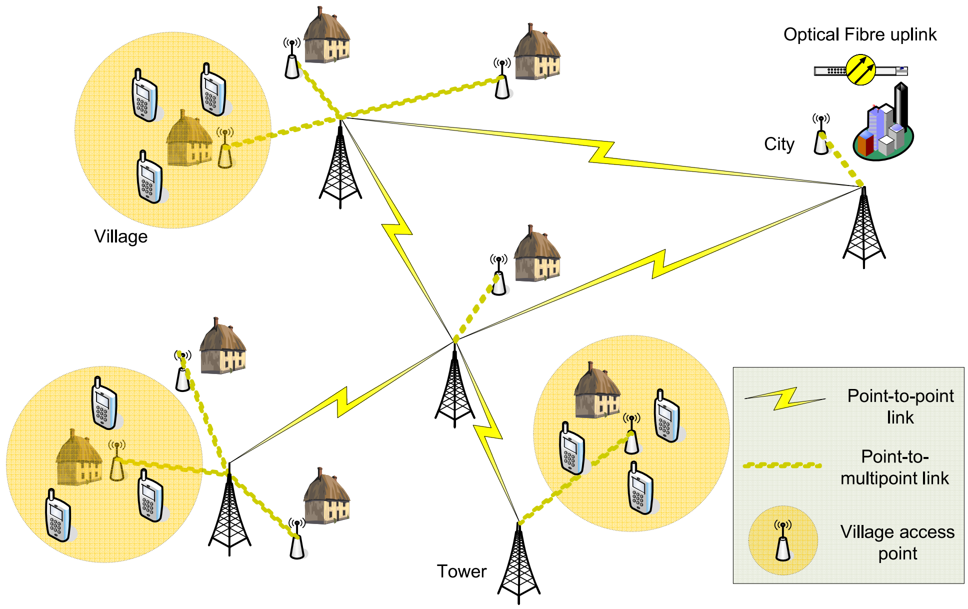

Figure 1.2: WiFi based rural network: Showing a combination of long-distance point-to-point backhaul links, medium-range point-to-multipoint access links, local access wireless networks and WiFi-enabled devices.

Grassroots deployment: Our solution should allow even small organizations such as non-profits or hospitals to start a network deployment of a small size, that can be later scaled up as demand from users and revenues increase. Technologies that require expensive base-stations need a large user base right from the beginning to amortize costs. Again we prefer technologies that work in the unlicensed spectrum because they can be set up by grassroots organizations as needed, avoiding dependence on a telecom carrier.

Usage patterns: Our solution should support various usage patterns and not be restricted to provide just Internet connectivity. By providing strong connectivity between villages and cities within an area, we would enable many promising applications and also encourage sharing and exchange of information within local communities.

Local maintenance: A key challenge for long-term sustainability of rural networks is lack of trained local manpower. Reliance on remote experts to fix problems whenever the network faces downtime is expensive and could take a lot downtime. Even when local expertise is available, staff turnover is a problem as trained manpower often leave to find better avenues for their skills. Thus we need to design systems that allow easy monitoring. We also need to run training programs to build capacity among local administrators so that they can diagnose problems, replace equipment and perform as much as possible of the maintenance locally without intervention from remote experts.

1.5.2 Architecture Components

With the above design principles in mind, we propose a network architecture (shown in Figure 1.2) consisting of long-distance WiFi based point-to-point backbone links and point-to-multipoint access links terminated by local access points and user WiFi devices.

Point-to-Point backhaul network: Long-distance point-to-point links (ranging from 15–150 km) provide high-bandwidth backhaul capabilities. Each network node is equipped with multiple high-power radios with directional antennas. The aim is to provide a reliable and stable backbone network that interconnect bigger population centers to optical fiber uplinks.

Because of radio propagation requirements, all the links need to have direct line of sight, and the nodes usually have be mounted on towers or masts on top of existing structures. In fact, the tower cost could be a significant component of the total cost of the installation. All these factors necessitate careful planning of the network topology. The choice of locations and selection what links to connect can be framed as a network wide optimization problem to minimize the overall cost of all the towers and equipment in the network [123].

The key challenges that we tackle here include the design of novel MAC protocols and loss recovery mechanisms for WiFi that can deliver high end-to-end throughput and low delay over long distances on both single hop and multihop paths.

Point-to-Multipoint access network: Short to medium-distance point-to-multipoint links provide connectivity using multiple sectorized antennas. The aim is to distribute bandwidth from the backbone nodes to stations in villages and other targeted locations (schools, hospitals) which might be at distances between 1–15 km away. The base-station (co-located at a backbone node) is equipped with multiple wireless radios and multiple sector antennas, while the stations are equipped with directional antennas pointing towards the base-station. The total bandwidth at the base-station is shared among all the stations. The nature of the connectivity implies that the set of client stations and the traffic demand might be dynamic.

The medium distances of these links still implies that directional antennas would be required at each of the stations, leading to hidden node effects among different stations. The key challenges that we tackle here include the design of time-division based MAC protocols to allocate bandwidth to stations, and to use various other techniques such as transmit power control and variable channel width adjustment to increase capacity and spectrum efficiency of the base-station.

Local access mesh networks: Each of the client stations at the end of the point-to-multipoint network might further provide local omni-directional wireless coverage in a limited area such as a village or inside a hospital or school. This part of the network could use off-the-shelf access points that run standard WiFi protocol without modifications.

End-user WiFi devices: This part of network consists client devices such as PCs or laptops with WiFi cards using the network for data access, or WiFi phones running voice over IP and other applications.

1.6 Contributions

The WiLDNet architecture has the potential to the most cost effective networking solution for rural areas, but the practical implementation in real world environments is not straightforward.

Our own early experience with long links and the experience of other groups such as Digital Gangetic Plains [35] and the Akshaya project [8] showed that WiFi links performed very poorly even at medium distances. For example, a 60 km 802.11b link we installed in rural Ghana in 2005 showed TCP bandwidth of only 600 Kbps although the raw capacity of the link was closer to 6 Mbps. This was particularly surprising as there were no other interfering WiFi sources in the vicinity of this link. In urban environments, we found WiFi links to be especially susceptible to interference from external sources resulting in losses as much as 60%. Finally, we saw that, whenever a node is configured to operate multiple directional links on the same channel because of spectrum considerations, the multihop performance dropped drastically.

The primary contribution of this dissertation is to answer the question, *what are right MAC-layer protocols to achieve high throughput and low delay in rural point-to-point* *and point-to-multipoint long-distance wireless networks?*

The approach of the dissertation research follows a logical progression where we first try to understand the root causes behind the sub-optimal performance of the standard WiFi protocol at long distances. We then build appropriate mechanisms to overcome the challenges. In particular, we implement a new time-division based MAC protocol (WiLD-MAC) to replace the existing CSMA MAC protocol that improves both throughput and spectrum efficiency over long-distance links by using loose time synchronization. Finally, we use this new MAC-layer to optimize the performance of both point-to-point backbone links and point-to-multipoint access links.

The set of techniques developed by us are not limited to just WiFi but could be complimented with many other kinds of physical layer wireless technologies.

We now summarize the main contributions of this dissertation.

1.6.1 Characterize the performance of WiFi in long-distance settings

The first contribution is to perform a systematic study to investigate the commonly cited sources of packet loss induced by the wireless channel and by the 802.11 MAC protocol [129].

To understand channel induced losses, we study different sources such as external WiFi interference, non-WiFi interference and multipath interference. We find out that outdoor wireless links are susceptible to external interference and can show highly variable packet loss.

We also study losses induced from the deficiencies of the 802.11 MAC protocol and identify three specific shortcomings: (a) the default 802.11 link-level recovery mechanism results in low utilization; (b) at long distances frequent collisions occur because of the failure of CSMA/CA; and (c) nodes with multiple links on the same channel experience inter-link interference.

1.6.2 Implement novel time-division based MAC-layer for long-distance links

The second contribution is to develop a novel MAC-layer (WiLDMAC) that replaces WiFi's standard CSMA with a time-division based access layer. To ensure per hop loss recovery, the stop-and-wait acknowledgment protocol of WiFi is replaced by a combination of sliding-window based flow-control with *bulk acknowledgments* (ARQ) and forward error correction (FEC).

For nodes with multiple radios (each with its own directional antenna), WiLDMAC imposes synchronization among them. The nodes use a largely *interference-free* mode of operation termed as Simultaneous Synchronized Operation (*SynOp*) where they either transmit simultaneously (*SynTx*), or receive simultaneously (*SynRx*) [112]. A simple loose time synchronization mechanism is used where during each time slot along each link, the sender acts as the master and the receiver as the slave.

We implement WiLDMAC in Linux on top of the madwifi driver for Atheros radio cards using the Click modular router framework.

1.6.3 High end-to-end performance in point-to-point backbone networks

For the long-distance point-to-point backbone part of our architecture, we use the time-division based MAC protocol (WiLDMAC) to achieve high end-to-end throughput and low delay over multihop wireless paths [104].

- We show that the TDMA-based WiLDMAC protocol breaks the dependence of throughput with distance and achieves close to optimal throughput at any distance.

- We use two different loss recovery mechanisms to deal with channel losses. A first approach uses retransmissions with minimal throughput overhead but at the expense of increased delay. A second forward error correction (FEC) based approach incurs additional throughput overhead but does not incur lower delay penalties. We show that we can combine these two mechanisms to trade off between throughput, loss and latency depending upon application traffic demands.

We deploy and validate the performance of our modified MAC-layer on various real life testbeds including the in Bay Area, India and Venezuela.

1.6.4 Capacity scaling in point-to-multipoint access networks

In the point-to-multipoint part of our architecture, the base-station has multiple radios with sectorized antennas, each service multiple client stations at medium to long distances.

We use the same time-division based MAC protocol (WiLDMAC) to allocate time-slots to clients from the base-station based on fairness and bandwidth demand. We then develop the following techniques to improve spectrum efficiency, allow capacity scaling at lower costs, and to support high number of clients in the point-to-multipoint part of our network architecture.

- *Dynamic transmit power control:* To achieve maximum capacity with limited availability of non-overlapping wireless channels, we need to operate as many radios as possible from different sectors on the same channel. However, operating co-located radios on the same channel can result in substantial interference especially with the current practice of operating all radios at maximum power.

  We formulate an LP optimization problem that maximizes throughput by computing optimal transmit schedules, optimal allocation of clients to base-station radios, and optimal radio power levels [105].

- *Antenna combination:* We show that by combining multiple radios in to the same physical antenna, we can scale up capacity at the base-station without increasing the number of antennas simultaneously.

  We demonstrate a simple yet practical multiplexing design that uses cheap RF combiner/splitter devices for multiplexing several radios onto a single antenna while achieving high throughput. We also examine a more general design that uses RF switches, and which allows us to allocate radios to antennas dynamically based on client traffic demands [48].

- *Channel width adaptation:* We show that by using variable width channels, instead of just the small number of fixed wireless channels, we can increase the number of simultaneous transmissions from the base-station and significantly improve cumulative capacity [47].

1.7 Roadmap and Organization

We first start by providing a comprehensive overview of the WiFi technologies in Chapter 2. In Chapter 3, we present a detailed measurement study to understand the root causes behind the poor performance of outdoor WiFi links. Subsequently, we present the design and implementation of WiLDMAC, the novel time-division based MAC-layer and evaluate it on various point-to-point testbed networks in Chapter 4. For the point-to-multipoint part of our architecture, we propose three techniques to increase capacity and spectrum efficiency of base-stations (Chapter 5).

Finally, in Chapter 6, we present results from testing our MAC-layer modifications on real world long-distance links. We also describe our experience from deploying the Aravind telemedicine network to show the real world impact of our work. We conclude by summarizing the contributions from this dissertation, the limitations of our approaches and the most promising future directions of our work.

Chapter 2

Background

The 802.11 group of standards (or WiFi) are the most ubiquitous wireless technology in use today. Not only is WiFi omnipresent in offices and homes, but it is also the technology of choice for outdoor networks such as wireless hotspots, university campuses, city-wide wireless mesh networks and long-distance networks to connect remote areas and provide emergency response.

The widespread popularity of WiFi and its rapid adoption all around the world in the last decade has been driven not only by its excellent performance, but also from commodification and standardization. Because of commodification, now there are numerous firms around the world that manufacture WiFi equipment in very large volumes (387 million WiFi chipsets were sold in 2008), bringing prices of WiFi devices down. The standardization of the protocols on the other hand has also ensured that WiFi equipment from different manufacturers interoperate with each other.

It is therefore not surprising that the extent of WiFi usage has far exceeded the goals and intentions of the original designers. WiFi has become the first choice for building low-cost community networks, rural wireless networks where distances between are orders of magnitudes longer than the distances inside offices, and other innovative applications in developing regions.

These efforts were lead by both non-profit community groups as well as academic research groups all over the world. In the United States, these include the Bay Area Research Wireless Network (BARWN) that runs a mesh of wireless nodes using directional antennas all over the Bay Area providing free Internet access [16], the Champaign-Urbana Community Wireless Network (CuWiN) that provides support to build decentralized, community-owned networks using open source technology [26], and lot of other similar groups in New York [98], Seattle [124] and numerous other cities. Outside of the United States, groups like AirJaldi [5] that runs a wireless mesh network in Dharamsala (India), Enlace HispanoAmericano de Salud (EHAS) Foundation [38] that uses WiFi to build wireless networks for health applications in Colombia [116] and Peru [131], CRCNet that helps connect rural communities in New Zealand [34] and Nepal Wireless [95] that connects a number of villages in the foothills of the Himalayas have pioneered the spread of WiFi.

A number of other groups also engaged in training of local communities around the world so that wireless networks can be installed and managed by local talents who are best positioned to respond to the needs of their communities while maintaining a viable network. This is very important to break the dependence on outside expert help, which has been the bugbear of many ICTD projects in the past. For example, the Abdus Salam International Center for Theoretical Physics (ICTP) in Trieste (Italy) organizes workshops to train administrators from various developing countries in the basics of radio link planning, WiFi link installation and IP networking [3]. AirJaldi runs a network academy where they impart knowledge that is relevant to professionals and operators in Dharamsala, where their network is located. All these efforts have also resulted in a lot of books and resources in the open domain including the book, *Wireless Networking in the Developing World* that is free for download in six languages [41].

While WiFi can provide high performance, it has also been used to create intermittent networks - projects such as the Wizzy Digital Courier [150] and KioskNet [125] have proposed the use of WiFi-enabled mobile vans that wirelessly synchronize data with village kiosks whenever they happen to be in communication range.

The common challenge that all these efforts that use standard off-the-shelf WiFi equipment for outdoor long-distance links face is that the WiFi protocol was not designed to work in these environments, and as a result, often performs poorly.

Our objective in this chapter is to provide a broad overview of different techniques that have been proposed and are being currently being used at different layers of the network and radio stack of WiFi (or 802.11) within the context of outdoor long-distance WiFi links.

We first present the basic set of innovations are independent of the 802.11 protocol or equipment, but are necessary to increase the range of WiFi from a few hundred feet to hundreds to kilometers using higher power radios and directional antennas (Section 2.2). We then discuss the various physical layer optimizations that have been proposed to improve throughput and loss on long-distance links in Section 2.3. Next, we introduce the concept of adaptive smart antennas that can significantly increase flexibility in installing and running wireless networks (Section 2.4). Finally, how we can present how we can overcome some of the more fundamental problems of performance, spectrum allocation spatial reuse in the existing 802.11 MAC-layer by either tweaking parameters within the existing standard or by devising new MAC protocols (Section 2.5).

2.1 Overview of WiFi 802.11

The 802.11 family of standards (also known in many circles as Wi-Fi) include various over-the-air modulation techniques that use the same basic protocol. The standards describe the supported physical layers and the common media access control (MAC).

The original version of the standard IEEE 802.11, that was released in 1997 specified two bitrates of 1 and 2 Mbps, but is obsolete today. The first widely accepted variant of 802.11 was 802.11b [61], ratified by IEEE in 1999 followed by 802.11a (also ratified in 1999) [60], 802.11g (ratified in 2003) and recently 802.11n (still in draft) [62]. In this section, we provide a brief overview of the basic physical layer and different MAC-layer mechanisms defined by the 802.11 family of standards.

2.1.1 Physical Layer

At the physical layer, IEEE 802.11 uses various modulation techniques including Direct Sequence Spread Spectrum (DSSS), Frequency Hopping Spread Spectrum (FHSS) and Orthogonal Frequency Division Multiplexing (OFDM).

The 802.11b variant uses DSSS, in a portion of the ISM band from 2.400 to 2.495 GHz. It has a maximum physical data rate of 11 Mbps, with actual usable data speeds up to about 7 Mbps. It is probably the most widely wireless technology in the world today and millions of devices supporting it have shipped since 1999.

The 802.11a variant of 802.11 uses OFDM that splits the signal across 52 separate sub-carriers to provide transmission at physical data rates as high as 54 Mbps. 802.11a operates in the ISM band between 5.745 and 5.805 GHz, and in a portion of the UNII band between 5.150 and 5.320 GHz. 802.11g uses the same OFDM modulation as 802.11a but works in the ISM range of 802.11b. It also has a maximum data rate of 54 Mbps (with usable throughput of about 22 Mbps). Because of its higher frequency range, 802.11a is incompatible with 802.11b or 802.11g and also has shorter range compared to 802.11b/g at the same transmit power.

Although the 802.11b standard uses a channel width of 22 MHz, the standard defines 11 channels spaced at only 5 MHz apart within the ISM band. However the spectral mask used by the modulation requires that the signal be attenuated by at least 30 dB from its peak energy at 11 MHz from the center frequency. As a consequence, stations can only use every fourth or fifth channel without overlap, typically channels 1, 6 and 11 for 802.11b.

2.1.2 Medium Access Control

The 802.11 standard for wireless LANs defines two different ways to coordinate transmissions among stations - DCF (distributed coordinations function) and PCF (Point Coordination Function). DCF which is predominantly used in wireless networks, uses a Carrier Sense Multiple Access protocol with Collision Avoidance (CSMA/CA) that considers all stations as equals (client stations and access points) and is the only one has been widely implemented in real hardware. Here we describe the basics of both DCF and PCF modes of operation but more details can be found in the full standards [59–61].

DCF mode of operation:

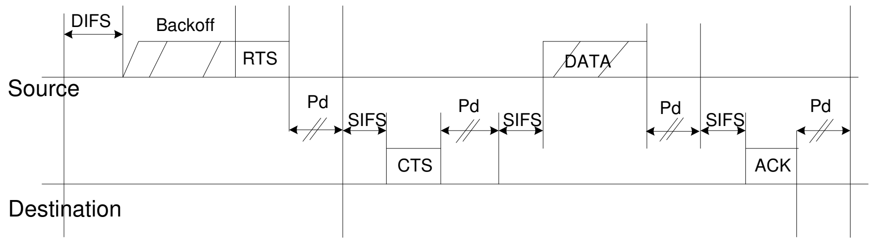

Figure 2.1: Timeline for DCF operation.

DCF employs the CSMA/CA mechanism and works as follows (timeline of DCF operation is shown in Figure 2.1). A station (including the AP) with a packet ready for transmission senses whether or not the channel is busy. If the channel is idle for a time period called DCF Inter Frame Space ($DIFS$) interval, the station starts packet transmission. Otherwise, the station continues to monitor the channel busy or idle status. After finding the channel idle for a $DIFS$ interval, the station starts to treat channel time in units of slot time ($SlotTime$). It generates a random backoff interval in units of slot time, but continues to monitor whether the channel is busy or idle. For each slot time where the channel remains idle, the backoff interval is decremented by one. When the interval value reaches zero, the station starts packet transmission. During this backoff period, if the channel is sensed busy in a slot time, the decrement of the backoff interval stops (i.e., is frozen) and backoff is resumed only after the channel has been idle for a full $DIFS$.

DCF uses an exponential backoff mechanism. The number of slots in the contention window is calculated as a uniform random variable in $[0, CW_{i+1}]$ where CW_{i+1} grows exponentially each time the transmission is unsuccessful, starting at $CW_{min} + 1$ and ending at $CW_{max} + 1$. After a successful transmission CW, is reset to the minimum value.

The receiver is required to send an ACK packet for each successfully received packet. A simple stop-and-wait protocol is used. The sending station is expected to receive the ACK within a Short Inter Frame Space ($SIFS$) interval, after the packet transmission is completed. If the ACK does not arrive at the sending station within a specified ($ACKTimeout$) period, or it detects transmission of a different packet on the channel, the original transmission is considered failed and retransmission is performed.

The protocol also implements an RTS/CTS mechanism that solves the hidden terminal problem. It uses a network allocation vector (NAV) to specify the expected duration of the current transfer (including the expected response to this packet).The value of the NAV indicates the amount of time is expected to pass until the channel becomes idle. All packets contain a duration field, specifying the expected duration of the current transfer (including the expected response to this packet). Thus, the NAV acts as a virtual carrier sense mechanism. The MAC uses the combined physical and virtual sensing to avoid collisions.

PCF mode of operation: PCF is a contention-free mode of operation, that assumes an access point node (AP). The AP polls its associated mobile stations one after another, by sending polling messages. If the AP has data to send to the station being polled, this data can be included in the poll message. If the polled station has data or acknowledgments for the AP, it is sent in the response message. This mode of operation is intended to support QoS by enabling the AP to regulate the bandwidth allotted to each of the mobile stations.

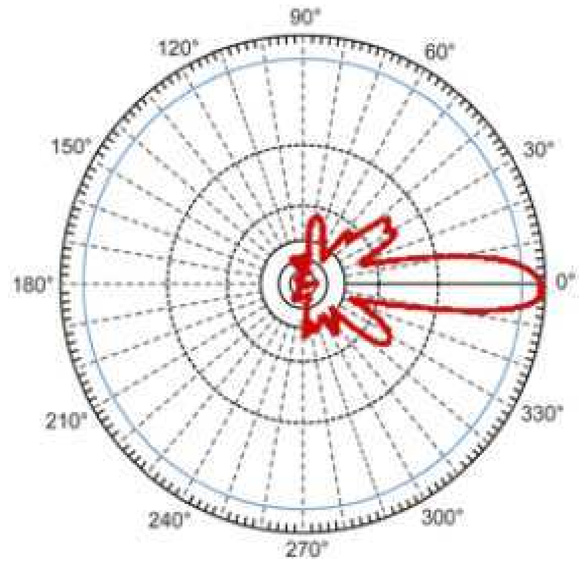

Figure 2.2: Antenna pattern for 27 dBi 2.4 GHz parabolic antenna [57]

2.2 Link Budget for Long-Distance Wireless

The range of most WiFi equipment is limited only a few hundred feet, and mostly in indoor environments. Our objective, however, is to connect rural communities that could be as far as 100 km apart.

The power output of most wireless access points is less than hundred milliwatts. The simplest technique to increase range is by using higher power output wireless cards or external power amplifiers. Although earlier most client wireless cards manufactured for the laptop market had low-power output (< 50 mW), the recently many manufacturers have release high-power wireless cards of upto 1000 mW [37, 142].

The second technique is to increase the EIRP (Effective isotropically radiated power) in a given direction by using highly directional antennas or sector antennas. Directional antennas are characterized by the peak gain and the beam width i.e the angular range

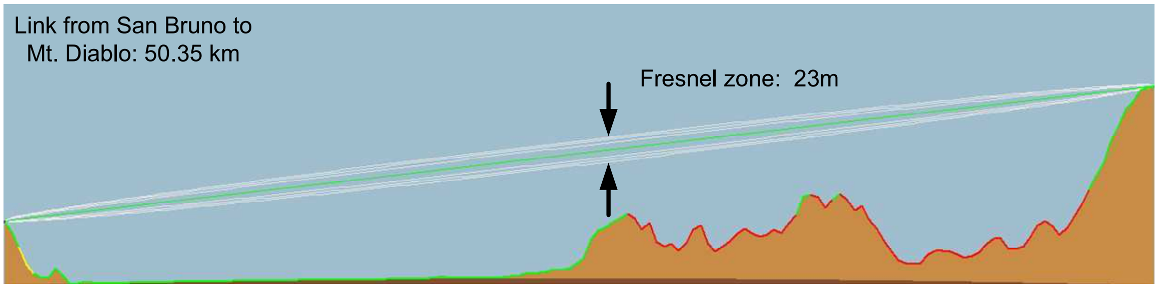

Figure 2.3: Path profile for a link from Mt Diablo peak to the San Bruno mountain in the Bay Area. The link is 50.7 km and the total path loss is 134 dB. The Fresnel zone is 23m at the center point.

where the signal power is at most 3 dB less than the peak gain. The gain of a parabolic shaped antennas depends on the size of the dish and the frequency of operation. Commercially available parabolic dish antennas can have gain of upto 30 dBi in the 2.4 GHz band for a 4 feet dish and 32 dBi in the 5 GHz band for a 3 feet dish. Because of their narrow beam pattern (Figure 2.2 shows the beam pattern for a 27 dBi 2.4 GHz antenna [57]), directional antennas can increase spatial use by allowing simultaneous operation on the same channel of multiple directed antenna links in the same physical vicinity.

The key challenge in operating a long-distance link is then to have sufficient receive signal to noise ratio (SNR) that exceeds the sensitivity of the receiver radio by a reasonable margin. An important caveat is that we have to assume a direct line-of-sight for most long-distance links exceeding a few kilometers. Since we know that radio signals fade with distance at the third power of distance, the path loss (in dB) for a transmitter is given as:

$$Path\_Loss = 92.4 + 20log10(F) + 20log10(D)$$

where F is the frequency in GHz and D is the path length in kilometers.

In addition it has been shown that because of radio wave diffraction, the signal spreads over a band (known as the Fresnel zone), which also has to be obstruction free for

Parameter	Value
Distance between Mt. San Bruno and Mt Diablo	50.7 km (31.5 miles)
True North Azimuth	64.2
Elevation angle	0.6290
Terrain elevation variation	1122.7m
Minimum clearance	0.8F1 at 50.3 km
Free Space Loss	134.3 dB
Total propagation loss	138.1 dB
System gain (symmetrical)	162.0 dB
Receive sensitivity	-90 dB
Margin of SNR	23.9 dB

Table 2.1: Key parameters for the link from Mt. San Bruno and Mt. Diablo shown in Figure 2.3

a good quality signal propagation. The first Fresnel zone can be calculated as follows:

$$H = 5.1\sqrt{D/F}$$

where, H is the width of the First Fresnel Zone (in meters), D is the distance between the antennas (in kilometers) and F is the frequency in GHz. However, because of the shape of the first Fresnel zone, even a line-of-sight path may not be a clean radio path. Infact, the link behaves essentially the same as a clear free-space path only if at least 60 percent of the First Fresnel Zone is clear of obstructions. For a more detailed discussion on characteristics of parabolic antennas and propagation refer to [41].

For example, consider a wireless link between San Bruno mountain and Mt. Diablo peak in the Bay Area. As shown in the path profile of an example link is shown in Figure 2.3, this link of 50 km would have a free space path loss of 134 dB. If we use parabolic antennas of 24 dB gain at either end and a transmitter of 300 mW (26 dB), and other cable losses at 4 dB, the received power is -66 dBm, which is about 24 dB more than the receive sensitivity of a typical wireless card (-90 dB). Also, the width of the first Fresnel zone is 23 m at the center point, which means that we need that much clearance from the direct line of sight path at the middle of the link. The key parameters of this path profile are summarized in Table 2.1.

2.3 PHY-Layer Optimizations

After the first set of 802.11 standards were released, many extensions have been proposed over the years. While some of these extensions such as 802.11n, 802.16 (WiMAX). try to increase performance, other extensions such as 802.11e are aimed to have better quality of service and fairness in wireless networks.

2.3.1 Using variable width channels

By default, 802.11b defines 11 channels in a frequency range of 2.412-2.462GHz in the US, each separated by 5 MHz. In 802.11a, FCC allows 18 default channels are separated by 10 or 20 MHz between 5.18GHz and 5.825GHz. However as the protocol uses 22 MHz of the band, adjacent channels overlap and interfere with each other. That means that there are effectively only three orthogonal channels can be used at any one point of time in 802.11b.

Information theory tells us that the maximum bitrate on a channel is roughly proportional to the channel width. Thus, changing the channel width offers the ability to tradeoff between having more simultaneous transmissions on orthogonal channels and having higher throughput on a single wider channel. A detailed analysis of the tradeoffs for single links is presented in [28].

In essence, using narrower channel widths has two main advantages for wireless networks with long-distance links - a) higher range from better SNR and resilience to multipath and b) higher cumulative throughput by using more orthogonality. We will also use narrower channels to increase capacity in point-to-multipoint access networks in Chapter 5.

1. *Higher range from better SNR:* If we assume that the transmit power of a wireless radio is constant regardless of the channel width, then at narrower widths, the radio can transmit with higher energy per unit Hz. If the noise per unit Hz is assumed to be constant at all channel widths, then the SNR is higher for narrower widths, giving us longer range. Theoretically, with perfect radios, if we use halve the width of a channel, we should get a 3 dB boost in SNR and using a quarter width channel should give a 6 dB boost (that corresponds to a 2x increase in range).

2. *More orthogonal transmissions:* We can also achieve more simultaneous transmissions if we use narrower channels, thereby increasing the total cumulative throughput of a network. It is known that the most optimal throughput on a channel can be achieved by using successive interference cancellation. Unfortunately, this requires complex signal processing on specialized hardware, Instead, it can be proven that with stations that are continuously backlogged, we can always achieve optimal throughput by finding the right allocation of spectrum width to each channel. Although the maximum bitrate of the channel is also reduced, the cumulative throughput of all the narrower channels can equal the optimal throughput, if the channel widths are allocated proportionally to the received SNR on the channel [47, 140].

For practical implementation, the channel width can be configured many of today's off the shelf wireless radios. Atheros cards already support turbo mode where two continuous 20 MHz channels can be bonded together to form a 40 MHz channel that is supposed

Channel width	5 MHz	10 MHz	20 MHz	40 MHz
Symbol Duration	16 μs	8 μs	4 μs	2 μs
SIFS	40 μs	20 μs	10 μs	5 μs
Slot Duration	20 μs	20 μs	20 μs	20 μs
Guard Interval	3.2 μs	1.6 μs	0.8 μs	0.4 μs

Table 2.2: 802.11 timing parameters for different channel widths [28].

to support upto 108 Mbps (802.11n also supports both 20-MHz and 40-MHz channels). Narrower channel widths can be achieved by changing a hardware register that sets the frequency of the reference clock that drives the PLL inside the radio. Changing the clock rate also affects 802.11 timing parameters as summarized in Table 2.2. Since the symbol lengths are different across channel widths, the same modulation scheme that gives 24 Mbps with the default 20 MHz gives only 6 Mbps at 5 MHz and 12 Mbps at 10 MHz. Channel width configuration is also implemented in WiMAX which allows clients to use channels that are multiples of 1.25 MHz, 1.5 MHz and 1.75 MHz.

2.3.2 Using multiple paths

A big problem with wireless networks is that often, a signal reaches a receiver along multiple reflected paths in addition to the primary path. This *multipath* effect can cause significant degradation in the received signal. This can be either because of inter-symbol interference (ISI) where successive symbols in the same packet overlap with each other, or because of destructive interference between the signal received on multiple paths (fading). Measurement studies such as done by MIT's RoofNet group have shown that multipath effects could be primary cause of collisions and high packet loss in outdoor short distance mesh networks [20].

A traditional method to deal with multipath was to use antenna diversity at the receivers where receivers have multiple antennas, but choose the one with the best SNR.

However, recent enhancements such as MIMO i.e. Multiple Input and Multiple Output (as part of 802.11n set of standards) can take advantage of multiple paths by using *spatial multiplexing*. A MIMO radio transmits multiple data streams, each using a different antenna and transmitter to receivers that are similarly also equipped with multiple antennas. Because of the physical placement of the transmit and receive antennas, each signal follows a slightly different path to the receiver. By using sophisticated digital signal processing, the receiver can independently decode the signals from all the parallel data streams. Theoretically, MIMO can dramatically improve SNR and throughput by allowing multiple data streams. Commercially available radios with implementations of the draft 802.11n standard support configurations of upto four transmitters and receivers each [62].

In practice, multiple paths do not cause a lot of performance degradation over long-distance links. As explained later in Chapter 3, as the distance of a link increases, the path difference between the primary line-of-sight path and any ground reflected path decreases and becomes inconsequential.

2.4 Smart Steerable Antennas

Standard antennas that are used in off-the-shelf access points have an omni-directional radiation pattern, that is appropriate for in broadcast environments such as indoor offices, but are less than efficient in utilizing spatial usage. As we saw earlier, directional antennas not only increase the range of a link, but are more spectrum efficient by limiting their radiation in a fixed direction. Thus, they allow us to have multiple co-existing transmissions at a centeral node without interference between them even while sharing the same channel.

The key disadvantage with omni- and directional antennas is the lack of flexibility in controlling the radiation pattern i.e. the direction and gain of the beam. This makes conventional networks especially susceptible to interference and increases the cost of maintenance, as realignment of antennas to reconfigure the topology require manual intervention, and makes them less suitable for mobile and dynamic systems.

Antenna designs that allow software control of the radiation pattern using are called *electronically steerable antennas* or *smart antennas*. Figure 2.4 compares the beam patterns of the different systems. Steerable antennas are a good candidate technology for addressing not only range, coverage and interference issues, but to also make installation and operation of long-distance wireless networks much more easier. These advantages are described below.

- *Automatic alignment:* Installation of long-distance links is typically a tricky proposition. Getting a good quality link requires that the directional antennas at the two ends of the link have to be precisely aligned to each other. This often requires at least two people on either end, one person on the tower trying to change the antenna orientation while a second person reads off SNR values from packets or a spectrum analyzer. In addition, alignment is often not a one time operation, because with time, antennas can get out alignment and might require maintenance to make small adjustments once in a while. We have seen from our experience that on some links,

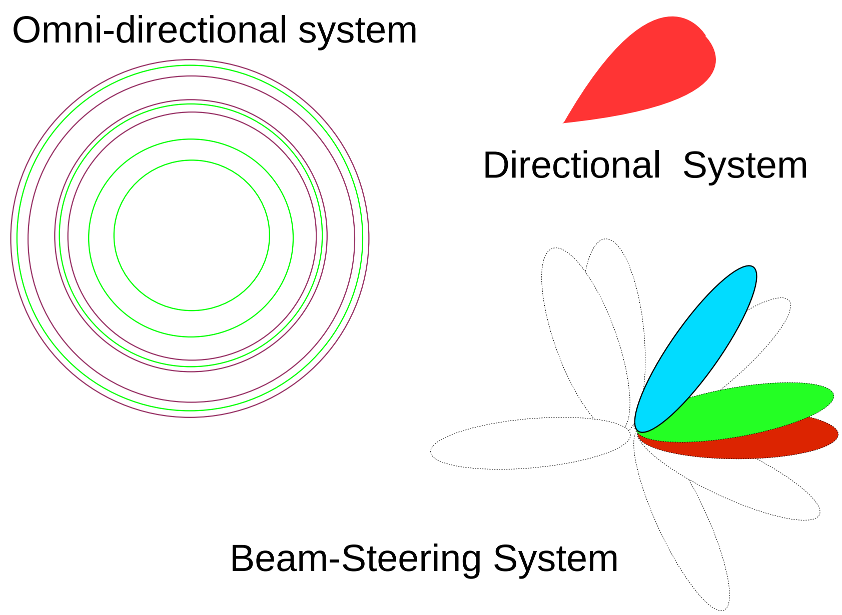

Figure 2.4: Radiation patterns with different antennas

we have seen the SNR drop by 1 to 3 dB in a year [138]. With steerable antennas, the first time installation of link does not require any precise alignment and could be possibly be performed even by untrained personnel. Secondly, since the antennas can dynamically adjust the beam to maintain alignment and link quality, there is no need for frequent manual intervention to maintain beam alignment.

- *Dynamic coverage:* Steerable antenna systems can focus their beams (energy) in many directions and dynamically configure their patterns at very high speeds, and thus potentially combining the best features of both omni-directional and directional systems. With antenna beam steering, radio transmitters can dynamically focus their transmissions on their desired target(s), and thus any radio that is sufficiently far away from the desired target(s) will experience very little interference due to the beamforming transmitters.

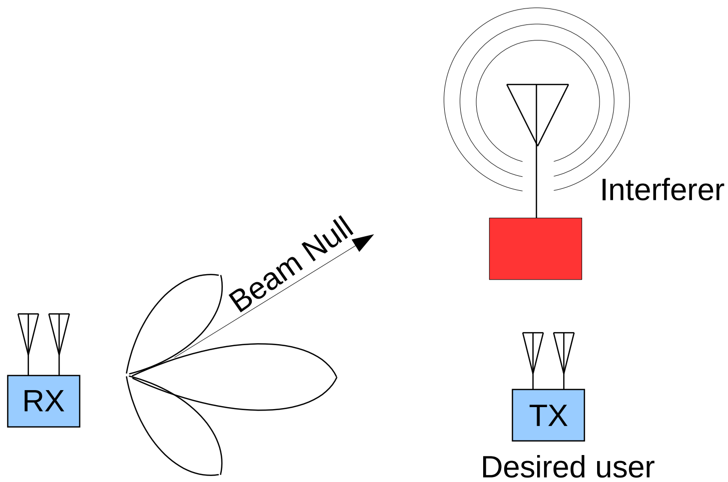

Figure 2.5: Minimizing interference using steerable antennas

The ability to change direction of a link also gives us a powerful tool to reconfigure links in a network to adapt to changing traffic demands or reroute traffic when there are failures in the network. For example, in our point-to-multipoint or access networks with a single base-station, we can steer the beam at the base-station to multiplex among clients while serving them with a better SNR and avoid interference.

- *Minimize interference:* Steerable antennas can modify their radiation pattern to identify and minimize interference from external sources, by creating nulls in the direction of prominent interfering sources (shown in Figure 2.5). For shorter distance links, steerable antennas can also be reduce multipath interference by choosing a pattern that has the least amount of reflected paths.

In the past, steerable antennas have had applications in military radars to track aircraft and moving objects. Recent advances in radio engineering and manufacturing have opened up the possibility of using such technology in lower end commercial wireless equipment as well.

Recently, steerable phased-array beamforming antennas have been used to improve the performance of 802.11 links in the context of communication between a moving vehicle and roadside APs [91]. A lot of research work has also focussed designing MAC protocols for adhoc networks with nodes that are equipped with beamforming antennas [23, 32, 113].

2.5 MAC-Layer Optimizations

The 802.11 MAC protocol was not designed for long-distance operation. Therefore it is not surprising that real life long-distance links show very poor end-to-end performance. In this section, we briefly discuss why the standard DCF and PCF MAC coordination mechanisms are unsuitable at long distances. We then show how the performance of DCF can be improved to some extent by tweaking specific parameters within 802.11. Finally, we present related work on designing completely new MAC protocols and synchronization techniques specifically for optimizing performance and spatial reuse of long-distance multihop 802.11 networks (Digital Gangetic Plains [35] and Roofnet [20]). Some of this work was done by researchers within the TIER group in conjunction with the work presented in this dissertation.

2.5.1 DCF at long distances

A number of analytical models have been proposed for the performance of 802.11 DCF's CSMA in broadcast environments with short distances [19]. Recent work has tried to extend those models to estimate the impact of distance on the throughput of DCF [130].

The main impact of increasing distances between stations in a wireless network are on the retransmission timeout ($ACKTimeout$) and on the contention resolution mechanism. A more through measurement based analysis of these problems is present in Chapter 3.

Since the standard defines the $ACKtimeout$ parameter as the maximum time during which this ACK packet must arrive at the transmitter. Since the default value of $ACKTimeout$ is defined as a constant in the standard, for longer links, the round-trip propagation delay can easily exceed it causing the time to elapse before the reception of the ACK. In that case, the transmitter might discard all ACK frames and retransmits a packet multiple times until the maximum retransmission limit is reached. Thus it can be easily shown that the $ACKTimeout$ needs to be increased to be at least twice the propagation delay.

The contention resolution of DCF is based on the assumption that the maximum distance between any two stations is always less than the slot time so that stations can detect each other's transmissions and backoff. Clearly, this assumption is not true in the case of long-distance links where the propagation delays can be much larger. Analysis shows that as the distance and thereby propagation delay increases and exceeds the default slot time, the stations are less and less likely to detect each others transmissions. If two stations start transmissions with an interval that is less than the propagation delay, their transmissions would definitely result in a collision.

2.5.2 PCF at long distances

Unfortunately, strict timing requirements make the PCF standard also ill-suited for long-distance outdoor links. In particular, the most stringent requirement is that the ACK has to be received from the polled station to the AP within the $SIFS$ time interval, which for 802.11b is only 10 μs. Unfortunately, this corresponds to a round trip of approximately 3 km, limiting links to less than 1.5 km, which is unacceptable in our scenario. Current 802.11 cards do not support setting the $SIFS$ interval to arbitrary values A more practical problem related to PCF is that, because it is optional, the mode is only supported in very few of the current wireless cards in the market.

2.5.3 Optimize DCF parameters

The extended analytical model for long-distance links presented in [130] suggests that losses from collisions can be decreased if we tweak some of the timing parameters of 802.11's DCF.

These parameters are available in EDCA (Enhanced Distributed Channel Access) as a part of 802.11e set of amendments. These extensions try to overcome a key limitation in standard 802.11's channel access mechanism that does not allow any way to differentiate between high priority and low priority traffic. This traffic differentiation is created by configuring a key set of parameters - $AIFSN_i$ (Arbitrary Inter-Frame Space Number), CW_i (Contention Window) and $TXOP_i$ (Transmission Opportunity) differently for the i^{th} class of traffic. In addition, we can also control some non-standard parameters such $SlotTime$ and $ACKTimeout$.

We can use these same parameters to optimize throughput and minimize collisions at longer distances. As shown in [121], the $SlotTime$ must be increased to twice the propagation time for distances longer than 3 km (for a slot time equal to 20 μs) or 1.35 km (for a slot time equal to 9 μs) in order to guarantee that two stations that listen to each other may only collide if they transmit in the same slot. These adjustments are a slight violation of the standard but they are feasible with many available commercial WiFi systems, and compatibility with legacy WiFi stations can still be maintained.

$$\begin{aligned}\sigma &= \begin{cases}\sigma_{std} & 2d_{max} <= c \cdot \sigma_{std} \\ 2\frac{d_{max}}{c} & 2d_{max} > c \cdot \sigma_{std}\end{cases} \\ \alpha &= \alpha_{std} - \sigma_{std} + \sigma\end{aligned} \tag{2.1}$$

where σ is the optimal value for $SlotTime$, σ_{std} is the standard value for $SlotTime$ (20 μs or 9 μs depending on the case), d_{max} is the maximum distance between two stations that can collide in a BSS (Basic Service Set), c is the speed of the light, α is the optimal value for $ACKTimeout$, and α_{std} is the standard value for $ACKTimeout$ [19].

In addition, it can be shown that the optimum congestion window, $CW_{min,i}$ changes with the distance, and there is an optimal value for $CW_{min,i}$ called CW_{opt}, that maximizes the throughput and minimizes the delay.

Although these parameter adjustments improve DCF's performance at long distances to some extent by eliminating almost all collisions, the final throughput is still less than optimal, mainly because of the inefficient stop-and-wait acknowledgment protocol used by 802.11 [121].

2.5.4 Synchronization for spatial reuse

Multihop long-distance wireless networks often have nodes equipped with multiple radios and co-located on the same tower. Although operation on different channels avoids inter-radio interference, there are several practical constraints that may require the operation of co-located radios on the same channel. This can happen due to country-specific

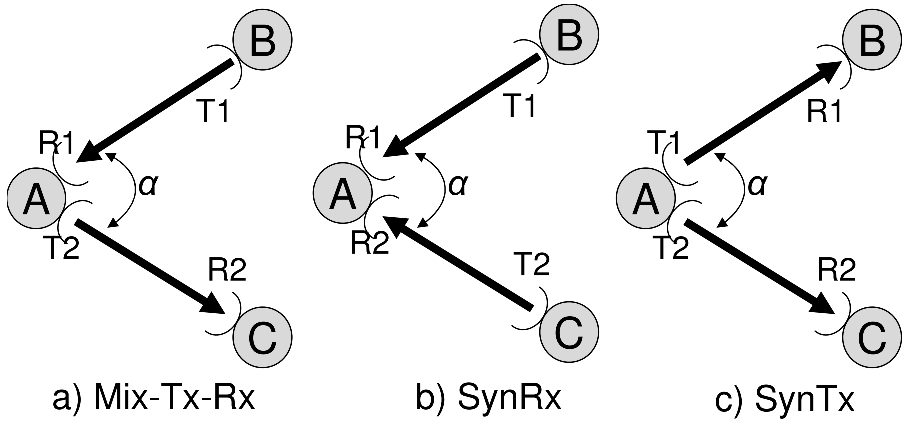

Figure 2.6: SynOp scenarios for interference

regulatory restrictions [135], high spectrum costs, or the limited number of available channels.

However, long-distance links with high-gain directional antennas interfere with each other in a very specific manner. More precisely, as first observed by Raman et al. [111], and reiterated in [104, 112], co-located radios (same physical location) operating on the same wireless channel interfere with each other if one of them transmits while the other receives. However, two adjacent directional links that either transmit simultaneously (*SynTx*), or receive simultaneously (*SynRx*), will be largely *interference-free* – a mode of operation termed as Simultaneous Synchronized Operation (*SynOp*).

We briefly explain the reason for this behavior. Consider the adjacent directional point-to-point links depicted in Figure 2.6, separated by an angle α. Now consider the following three potential interference scenarios:

1. *Mix-Tx-Rx:* In this scenario, depicted in Figure 2.6(a), T_2's transmissions interfere with R_1's reception, due to the physical proximity between the radios and the presence of antenna side-lobes. Therefore, operating the links in this mode is not feasible.

2. *SynRx:* During simultaneous receive, shown in Figure 2.6(b), T_2's transmissions are seen as interference at R_1, and T_1's transmissions are seen as interference at R_2. For the interfering signal to be ignored, the difference between useful signal and interference must be larger than a certain threshold $Th_{isolation}$, which depends on modulation and data-rate; e.g. with 802.11b at 11 Mbps, $Th_{isolation} \approx 12\ dB$ (Chapter 3) [111, 129]. Fortunately, this isolation can usually be ensured through the difference in gain levels provided by the directional antennas, if the links are separated by a sufficiently large angle. If we denote the difference between the antenna gain of the main lobe and the gain at an angle α away from the main lobe by S_{alpha} (also called the rejection level at angle α), then adjacent links are interference free under the following condition [111]:

$$|P_{R1} - P_{R2}| < S_\alpha - Th_{isolation} \tag{2.2}$$

where P_{R1} and P_{R2} are the receive power levels at R_1 and R_2 respectively.

For example, if links use typical 27 dBi grid antennas [57] (also used in our deployments) in horizontal polarization, an angular separation of more than 10° (half the width of the antenna main lobe) translates into an isolation of at least 25 dB (sometimes larger, not monotonically increasing with the separation angle). This means that 802.11b links receiving simultaneously are interference-free if $|P_{R1} - P_{R2}| < 15\ dB$. This can be easily satisfied by a large range of values (e.g. $P_{R1} = P_{R2}$), and even if the path loss of the two links is very different, the condition can be satisfied by adjusting the radio transmit power accordingly (by reducing the TX power on the stronger link).

3. *SynTx:* With simultaneous transmissions, as in Figure 2.6(c), interference may occur at nodes B and C, but not at node A. Once again, R_1 may see interference from T_2, and R_2 from T_1. Given the symmetry of the two links, ensuring non-interference during SynTx can be done by enforcing a similar condition to that in Equation 2.2.

   We note that simultaneous transmission is infeasible using a carrier-sensing MAC, such as 802.11, since radios can hear each others transmission, causing one of the radios to backoff.

In summary, simultaneous synchronized operation (*SynOp*) can allow multiple adjacent WiLD links to simultaneously use the same wireless channel *provided* the links are separated by a sufficiently large angle α and the radio transmit powers are chosen to satisfy the constraint from Equation 2.2. Given the gain pattern of typical grid directional antennas [57], an angular separation α larger than 30° provides generous isolation between adjacent links; this has also been demonstrated experimentally [111, 112] and validated in our deployments [104, 138].

2.5.5 New MAC design for outdoor long-distance links

As CSMA-based MAC protocols have been shown to perform poorly in networks with long-distance links [112, 129], leading to a preference for TDMA-based MAC solutions. 2P [112] was the first to propose a TDMA-based approach for WiLD networks. In our own work, WiLDNet [104] (presented in Chapter 4), we extend the 2P approach with techniques to deal with packet loss and to improve end-to-end performance in multi-hop long-distance networks.

In these MACs, long-distance links alternate between *transmit* and *receive* slots of fixed lengths. Inter-link interference is avoided by eliminating the situation in which a node transmits on one link while receiving on another. Thus a node transmits on all its simultaneously and then receives from all its links simultaneously, thus obeying the *SynOp* constraints as well.

Other TDMA-based approaches such as Overlay MAC Layer (OML) [114] implements a TDMA style MAC on top of the 802.11 MAC using loosely-synchronized clocks to provide applications and competing nodes better control over the allocation of time-slots. SoftMAC [96] is another platform to build experimental MAC protocols.

2.6 Summary

In this chapter, we presented a brief overview of WiFi technologies. We saw how we can increase the range of long-distance outdoor links, and then took a look at various enhancements such as optimizations at the PHY-layer, the use of smart antennas, and novel MAC-layer mechanisms to improve performance, reliability and cost of wireless networks, especially in remote rural areas.

Chapter 3

Characterization of long-distance WiFi (802.11) links

3.1 Background

Despite the promise of low-cost connectivity, the performance of WiFi in Long-Distance (WiLD) networks in the real world has been abysmal. This has been our own experience from links that we deployed in India and Ghana as well as the experience of several other groups such as Digital Gangetic Plains [30], the Akshaya project [8] and Roofnet [4].

For example, a 60 km link we installed in rural Ghana in 2005 showed TCP bandwidth of only 600 Kbps although the raw capacity of the link was closer to 6 Mbps. This was particularly surprising as there were no other WiFi sources in the vicinity of this link. Measurements from other studies have confirmed that even at medium distances, WiFi links often performed poorly.

We also observed high loss variability on our links in urban areas that had high levels interference from other WiFi sources. Figure 3.1 illustrates the loss measurement on two different links in our testbed. The loss rate was averaged over 60-second intervals for a 1 Mbps unidirectional UDP CBR traffic flow with the MAC-layer ACKs turned off and retries set to zero. We find that the loss is highly varying with time and there are bursts of high loss of lengths varying from few milliseconds up to several minutes. We also notice, that there is always a non-zero residual that varies between 1–20%. In contrast, the residual loss rates in our rural links are negligible. In addition, the loss characteristics along a single link are often highly asymmetric. For example, we observe that average loss rate from S to P was lower (10%) than the loss from P to S (20%).

Although Figure 3.1 shows only two links in our testbed, the above behavior is characteristic of all our links, with the urban links showing more loss variability than the rural ones. The key motivation behind the work presented in this chapter is to understand the underlying causes behind these losses.

.

3.2 Contributions

In this chapter, we perform a detailed measurement study to analyze the packet loss characteristics and the sources of packet loss in WiLD network settings. We categorize the sources of packet loss into two broad categories: (a) *channel losses* induced by the long-distance wireless channel; (b) *protocol-induced losses* due to shortcomings in the 802.11 MAC protocol. For each of these, we show that just manipulating driver level parameters is insufficient to achieve good performance over long-distance links.

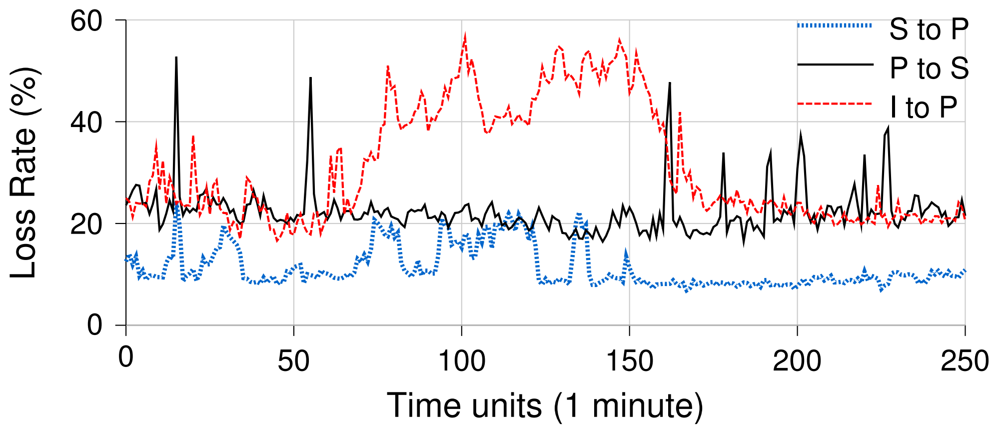

Figure 3.1: Packet loss variation on 2 links over a period of about 4 hours. Traffic was 1 Mbps CBR UDP packets of 1440 bytes each at a PHY datarate of 11 Mbps in 802.11b.

Our study is based on a real-world WiLD network deployment consisting of 6 links with lengths varying from 1–45 km. Unlike existing WiLD deployments [30], our testbed includes both rural and urban links. In addition to the real deployment, we also perform detailed experiments using a wireless channel emulator, which enables repeatable controlled experiments. These key contributions are this chapter are described below. These conclusions were also used to drive the design and implementation of the TDMA-based MAC protocol for our WiLDNet network architecture [104].

- **Protocol-induced losses:** The stock 802.11 MAC protocol is ill-suited for WiLD links due to the breakdown of CSMA over long distances and propagation delays (Section 3.4). Here, we pinpoint the fundamental shortcomings of the 802.11 MAC protocol.

- **Channel loss characterization:** We analyze three well known causes for channel losses in wireless environments, namely, *external WiFi interference*, *non-WiFi* *interference* and *multipath interference*. Among these, we show that external WiFi interference is the most significant source of packet losses in WiLD environments and the effect of multipath and non-WiFi interference is not significant. This is in contrast to the results of the Roofnet mesh network [4] where the authors observed multipath to be the most significant source of packet loss.

The focus of our packet loss characterization study is significantly different from other wireless-based loss measurement studies [4, 120]. The work done by Raman et al. [30] is the only other measurement-based study of WiLD deployments of which we are aware. However, the two studies are orthogonal: we focus on loss variability characterization, determining the impact of different sources of losses and remedies for loss alleviation, while their work focused more on performance analysis of 802.11 network at various layers in the network stack and the effect of other parameters (weather, SNR, payload, datarate) on loss. Our work also differs from mesh networks like Roofnet [4] in that WiLD networks, as we show, have very different loss characteristics, with loss much more due to external interference than multipath effects.

3.3 Experimental Setup

We use three different experimental setups to conduct measurements.

Campus testbed: Figure 3.2 is our real-world campus testbed on which we have currently deployed WiLDNet. The campus testbed consists of links ranging from 1 to 45 km, with end points located in areas with varying levels of external WiFi interference. We also use one of the links in our Ghana network (65 km).

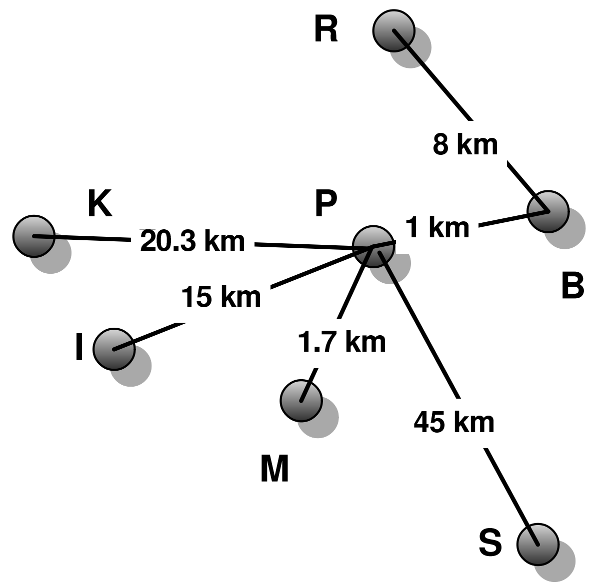

Figure 3.2: Overview of the WiLD campus testbed (not to scale).

Wireless Channel Emulator: The channel emulator (Spirent 5500 [132]) enables repeatable experiments by keeping the link conditions stable for the duration of the experiment. Moreover, by introducing specific propagation delays we can emulate very long links and hence study the effect of long propagation delays. We can also study this in isolation of external interference by placing the end host radios in RF isolation boxes.

Indoor multi-hop testbed: We perform controlled multi-hop experiments on an indoor multi-hop testbed consisting of 4 nodes placed in RF isolated boxes. The setup was designed to recreate conditions similar to long outdoor links where transmissions from local radios interfere with each other but simultaneous reception on multiple local radio interfaces is possible. We can also control the amount of external interference by placing an additional wireless node in each isolation box just to transmit packets mimicking a real interferer. The amount of interference is controlled by the rate of the CBR traffic sent by this node. The indoor setup features very small propagation delay on the links; we use it only to perform experiments evaluating TDMA scheduling and loss recovery from interference.

We use Atheros 802.11 a/b/g radios for all our experiments. The wireless nodes are 266 MHz x86 Geode single board computers running Linux 2.4.26. The choice of this hardware platform is motivated by the low cost ($140) and the low power consumption (< 5W). We use *iperf* to measure UDP and TCP throughput. The madwifi Atheros driver was modified to collect relevant PHY and MAC layer information.

3.4 802.11 Protocol Shortcomings

In this section, we study the three main limitations of the 802.11 protocol: the inefficient link-layer recovery mechanism, collisions in long-distance links, and inter-link interference. These limitations make 802.11 ill-suited even in the case of a single WiLD link. Based on extensive experiments, we also show that modifying the driver-level parameters of 802.11 is insufficient to achieve good performance.

3.4.1 Inefficient Link-Layer Recovery

The 802.11 MAC uses a simple stop-and-wait protocol, with each packet independently acknowledged. Upon successfully receiving a packet, the receiver node is required to send an acknowledgment within a tight time bound ($ACKTimeout$), or the sender has to retransmit. This mechanism has two drawbacks:

• As the link distance increases, propagation delay increases as well, and the sender waits for a longer time for the ACK to return. This decreases channel utilization.

• If the time it takes for the ACK to return exceeds the $ACKTimeout$ parameter, the

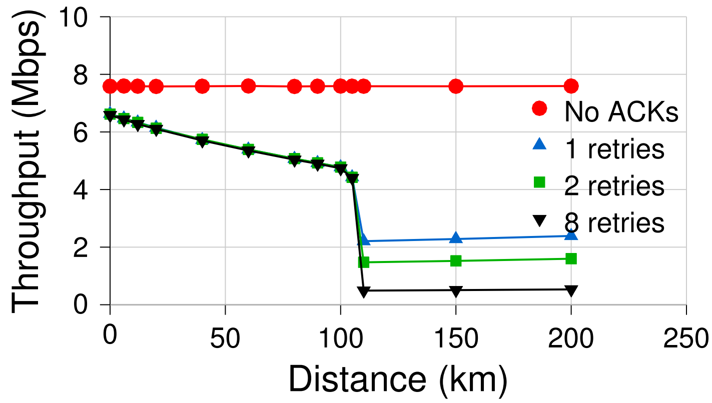

Figure 3.3: Unidirectional UDP throughput for standard 802.11 CSMA on single emulated link ACK timeouts were adjusted with increasing distance (on Atheros cards). Traffic is 1440 byte CBR UDP packets in 802.11b at PHY-layer datarate of 11 Mbps.

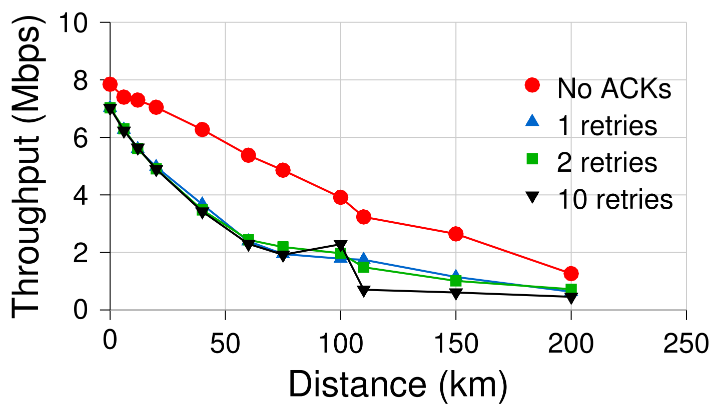

Figure 3.4: Bidirectional UDP throughput for standard 802.11 CSMA on single emulated link ACK timeouts were adjusted with increasing distance (on Atheros cards). Traffic is 1440 byte CBR UDP packets in 802.11b at PHY-layer datarate of 11 Mbps.

sender will retransmit unnecessarily and waste bandwidth.

We illustrate these problems by performing experiments using the wireless channel emulator. To emulate long distances, we configure the emulator to introduce a delay to emulate links ranging from 0–200 km. Figure 3.3 shows the performance of the 802.11 stop-and-wait link recovery mechanism over increasing link distances. With the MAC-layer ACKs

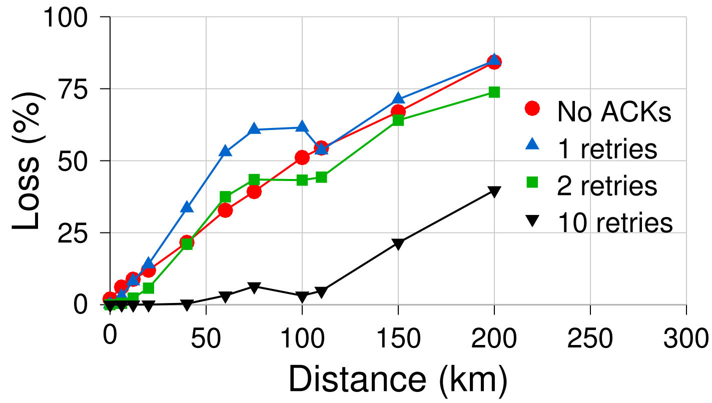

Figure 3.5: Bidirectional UDP loss for standard 802.11 CSMA on single emulated link ACK timeouts were adjusted with increasing distance (on Atheros cards). Traffic is 1440 byte CBR UDP packets in 802.11b at PHY-layer datarate of 11 Mbps.

turned off, we achieve a throughput of 7.6 Mbps at the PHY-layer data rate of 11 Mbps. When MAC ACKs are enabled, we adjust the ACK timeout as the distance increases. In this case, the sender waits for an ACK after each transmission, and we observe decreasing channel utilization as the propagation delay increases. At 110 km, the propagation delay exceeds the maximum ACK timeout (746 μs for Atheros, but smaller and fixed for Prism 2.5 chipsets) and the sender always times out before the ACKs can arrive. We notice a sharp decrease in received throughput, as the sender retries to send the packet repeatedly (even though the packets were most likely received), until the maximum number of retries is reached. This happens because, if an ACK is late, it is ignored. This causes the received throughput to stabilize at $BW_{110km}/(no\_of\_retries + 1)$.

3.4.2 Collisions on long-distance links

The 802.11 protocol uses a CSMA/CA channel-access mechanism, in which nodes listen to the medium for a specified time period ($DIFS$) before transmitting a packet, thus ensuring that the channel is idle before transmission. This translates to a maximum allowable distance at which collisions can be avoided of about 15 km for 802.11b ($DIFS$ is 50 μs), 10.2 km for 802.11a and 8.4 km for 802.11g. For longer links it is possible for a node to start transmitting a packet unaware of another packet transmission at the other end. As the propagation delay increases, this probability of loss due to collisions increases.

We illustrate the above-mentioned effect by using a simple experiment: we send bidirectional UDP traffic at the maximum possible sending rate on the emulated link and measure the percentage of packets successfully received at each end. Figure 3.5 shows how the packet loss rate increases with distance. Figure 3.4 shows the sum of the throughputs achieved at both ends for bidirectional UDP traffic as we increase the distance for a link. Note that there are no losses due to attenuation or outside interference in this controlled experiment; all of the losses are due to collisions.

A possible solution to this issue would be to increase the $DIFS$ time interval in order to permit longer propagation delays. However, just as in the case of the ACK timeout, this approach would decrease channel utilization substantially for longer links. Furthermore, we are not aware of any 802.11 chipsets that allow the $DIFS$ interval to be configured.

3.4.3 Multiple Link Interference

Another important source of errors is the interference between adjacent 802.11 links operating in the same channel or in overlapping channels. Although interference between adjacent links can be avoided by using non-overlapping channels, there are numerous reasons that make it advantageous to operate adjacent links on the same frequency channel, as described by Raman et al. [112]. Moreover, there are WiLD topologies such as the Akshaya network [8] where different channels cannot be allocated to all the pairs of adjacent links, given the high connectivity degree of several nodes.

Inter-link interference occurs because the high-power radios create a strong RF field in the vicinity of the radio, enough to interfere with the receptions at nearby radios. Directional antennas also have sufficiently high gain (4–8 dBi) side lobes [17] in addition to the main lobes.

The first type of problem occurs when multiple radios attached to the same node attempt to transmit at the same time. As soon as one radio starts transmitting after sensing the carrier to be idle, all other radios in the vicinity find the carrier to be busy and backoff. This is desirable in a broadcast network to avoid collisions between two senders at any receiver node. However, in our network where each of these radios transmits over point-to-point long-distance links to independent receivers, this backoff leads to suboptimal throughput. A second problem occurs when packets being received at one link collide with packets simultaneously transmitted on some other link on the same node. The signal strength of packets transmitted locally on a node overwhelms any packet reception on other local radios.

In order to illustrate these effects, we perform experiments on the real-world setup presented in Figure 3.2. First, we attempt to simultaneously transmit UDP packets to both K and M from node P. The total send throughput on both links is 14.20 Mbps when they are on non-overlapping channels (separation ≥ 4) but drops to only 7.88 Mbps when on the same channel. Next we send UDP packets from node M to node K, relayed through node P at different transmitting rates. We then measure received throughput and packet loss rate

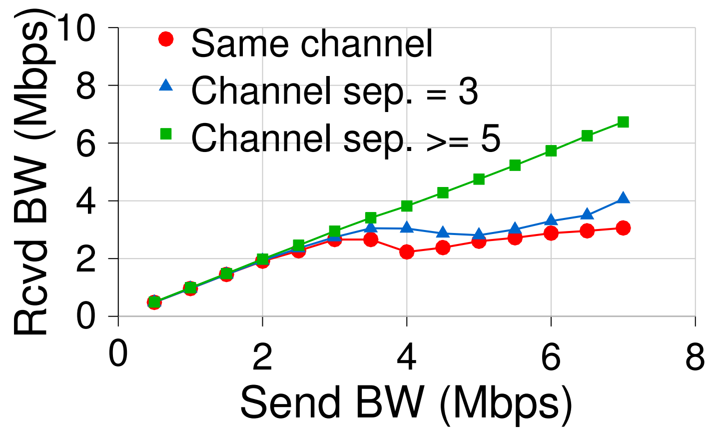

Figure 3.6: Effect of interference on received UDP throughput when sending from M to K through a relay node, P. Channel separation is no. of channels in 802.11b. Traffic is 1440 byte CBR UDP packets in 802.11b at PHY-layer datarate of 11 Mbps.

for various channel spacing between the two adjacent links, as presented in Figures 3.6 and 3.7. We observe that interference does reduce the utilization of the individual links and significantly increases the link loss rate (even in the case of partially overlapping channels).

Therefore, the maximum channel diversity that one can simultaneously use at a single node in the case of 802.11(b) is restricted to 3 (channels 1,6,11) which may not be sufficient for many WiLD networks. This motivates the need for a scheme that allows the efficient operation of same-channel adjacent links. This can be achieved by using a mechanism similar to the one used in 2P [112], that synchronizes both packet transmission and reception across adjacent links to avoid interference and improve throughput.

3.5 External WiFi Interference

In this section, we investigate external WiFi interference as a potential source of packet loss in WiLD links. Any WiFi traffic that is not a part of the primary WiLD link

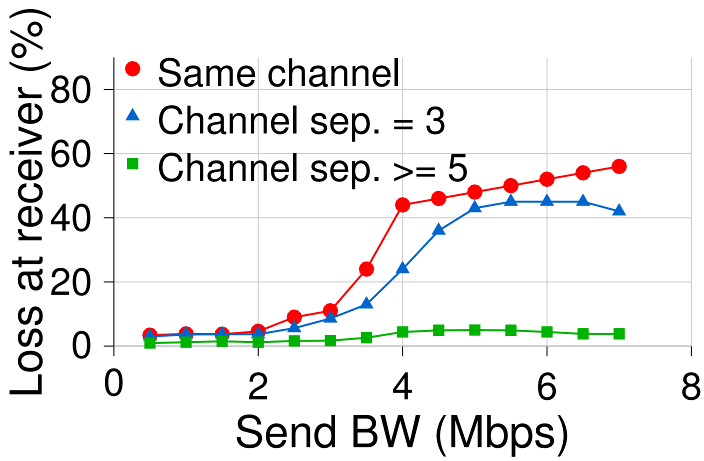

Figure 3.7: Effect of interference on received UDP error rate when sending from M to K through a relay node, P. Channel separation is no. of channels in 802.11b. Traffic is 1440 byte CBR UDP packets in 802.11b at PHY-layer datarate of 11 Mbps.

is categorized as external WiFi interference. Based on the measurements performed on our WiLD testbed and the wireless channel emulator, we show three key results:

- In the presence of external WiFi interference, the loss rate is strongly correlated with the amount of external traffic received on the same and adjacent channels. In contrast, due to the omni-directional antennas used in the Roofnet deployment [20], no such strong correlation was observed.

- Packet loss due to external WiFi interference is far more significant in WiLD deployments than local mesh networks.

- The loss due to external WiFi interference depends on the relative power level between the primary and external traffic, their channel separation, and the rate of external interference.

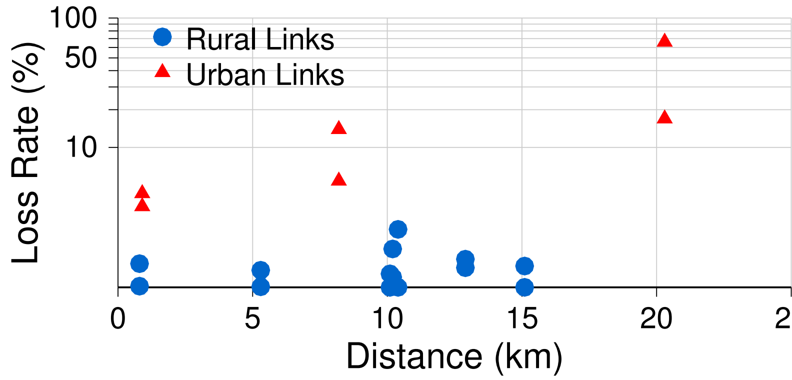

Figure 3.8: Scatter plot of loss rates observed in links deployed in urban and rural areas (note: loss rate is plotted in logscale).

3.5.1 Correlation of loss rate and external WiFi traffic

Figure 3.8 shows the loss rate across all (rural and urban) our WiLD links. We observe that the loss rate of the urban links vary across a wide range (4–70%). In contrast, all the rural WiLD links have a very small loss rate. The maximum loss rate observed in all our rural WiLD links was 2%.

To study this contrast between the rural and urban links, we collected detailed packet level MAC traces. By parsing the MAC header source and destination fields, we are able to count the number of frames received from external WiFi sources. In the traces collected over all our rural links we see negligible external WiFi traffic. However, significant amount of external WiFi traffic was captured from the traces collected in the urban WiLD deployment.

Figure 3.9 shows a scatter plot between the loss rate and the absolute number of external WiFi traffic frames received on an urban link (K $\rightarrow$ P) for a period of 6 hours. The figure shows that a subset of the loss rate samples are strongly correlated with the external

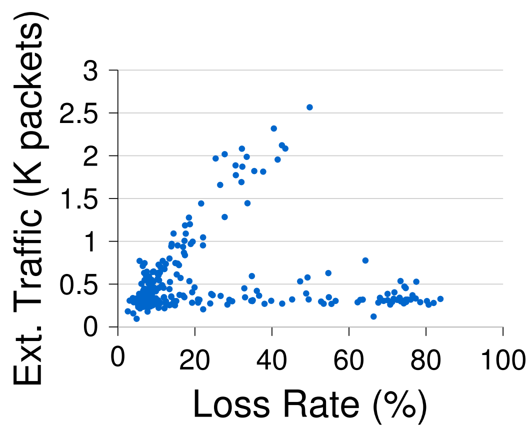

Figure 3.9: Correlation between loss rate and external traffic observed on a WiLD link (K $\rightarrow$ P). Traffic is 5 Mbps UDP CBR packets of 1440 bytes each at 802.11b PHY datarate of 11 Mbps.

traffic. For the other subset of the samples, the loss rate increases even when there is no significant increase in WiFi traffic on the same channel.

To investigate this further, we perform a controlled experiment using the wireless channel emulator. To model interference from an external traffic source, along with the primary link traffic we introduce a controlled interference source at the receiver. The traffic rate of the interference source was varied from 0.1 to 1 Mbps and the traffic rate on the primary link was kept fixed at 5 Mbps. Figure 3.10 shows a scatter plot of the loss rate and the total number of frames received from the external interference source. From the graph, we observe that for a given loss rate, the amount of external traffic captured by the monitor device depends on the channel separation of the primary and interference source.

The above observed trend is the same as that in Figure 3.9. At a channel separation of 0 and 1, the receiver can receive both the primary link traffic as well as the frames from the interference source. Hence, the loss rate is directly correlated with the amount of

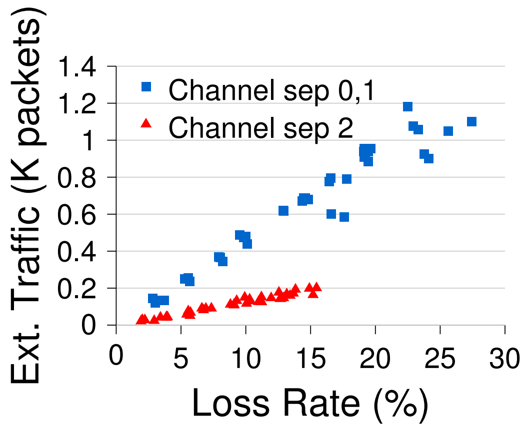

Figure 3.10: Correlation between loss rate and external traffic observed in the wireless emulator. Traffic is 5 Mbps UDP CBR packets of 1440 bytes each at 802.11b PHY datarate of 11 Mbps.

external WiFi traffic captured by the monitor interface. At a channel separation of 2, the receiver is not able to receive the frames from the external interference source. However, the signal spillage of the interference source in the primary channel is sufficient to cause frame corruption. From the traces we observed that almost 100% of the lost frames contained CRC errors.

3.5.2 Effect of hidden terminals in WiLD networks

Unlike WiLD deployments, where we have observed significant correlation between loss rate and external interference, it has been observed that there is no significant correlation in outdoor mesh-network deployments (Roofnet [20]). In a mesh-network deployment, an external interference source (I) that is within range of the omni-directional transmitter (Tx) would be able to sense the medium to be free and backoff its transmission. However in WiLD links, the transmissions are highly directional and the propagation delays are higher.

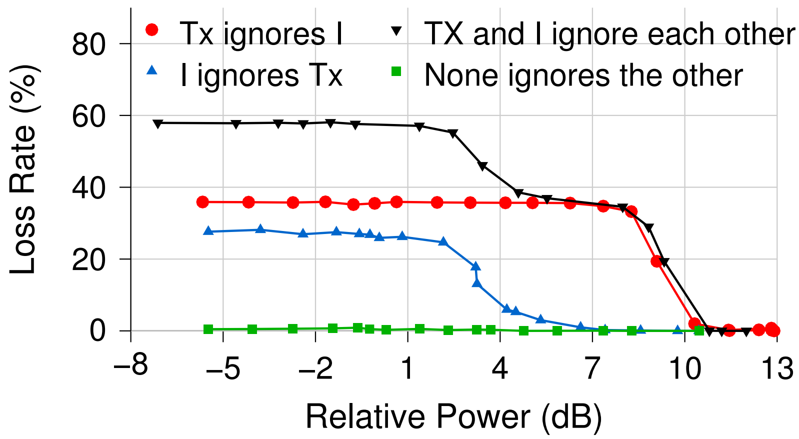

Figure 3.11: Losses due to different hidden terminal effects. Both main and interfering traffic is 1440 byte UDP CBR packets at 11 Mbps PHY datarate of 802.11b.

These factors in combination exacerbate the *hidden terminal* problem in WiLD networks. The transmitter and the interference source can erroneously sense the medium to be free leading to collisions whenever they are out of range of each other (because of the directional nature of transmission) or when they cannot sense the medium to be busy in time to backoff (because of the longer propagation delays).

Collisions at the receiver can manifest in two different situations: a) When I doesn't hear Tx, and initiates a transmission when the medium is busy with an ongoing packet transmission from Tx, and b) When Tx doesn't hear I, and causes a collision by interrupting an ongoing packet transmission from I.

To isolate the above two cases and measure the performance degradation due to each case, we perform controlled experiments using two WiFi links. We simultaneously send packets from both Tx (512 Kbps traffic) and I (3 Mbps), and measure the packet loss rate on the primary link ($Tx \rightarrow Rx$) with MAC-layer ACKs disabled.

To create the situation where Tx cannot hear I, we disable the Clear Channel Assessment (CCA) at Tx, which simply causes Tx to ignore I. We also eliminate propagation delay between Tx and I so that I's CCA works perfectly. We reverse the operations to create the situation in which I cannot hear Tx, but Tx hears I perfectly.

We then run four experiments, reflecting the losses in four situations: when Tx can't hear I, when I can't hear Tx, when neither can hear each other (representative of cases in WiLD networks), and when both Tx and I hear each other (representative of most cases in urban mesh networks).

Figure 3.11 shows the loss rate for each of the above four cases. In the case where I ignores Tx, to overcome the interferer completely (achieve 0% loss), packet transmissions from the Tx have to be 7 dB stronger than the interfering transmissions. This threshold, at which the primary link is loss free, is much higher (12 dB) in the case where Tx ignores I. When neither of Tx and I can hear each other, both the above two types of collisions are possible. Hence the loss rate is the sum of the losses generated by the above two types of collisions. However, when both Tx and I are in range of each other, resembling a mesh-network, losses due to collisions are close to zero. In this case, CSMA ensures that the two transmitters, Tx and I, share the medium properly.

From the above experiment we conclude that the effect of hidden terminals, causing collisions at the receiver, are greatly exacerbated in WiLD networks compared to urban mesh networks.

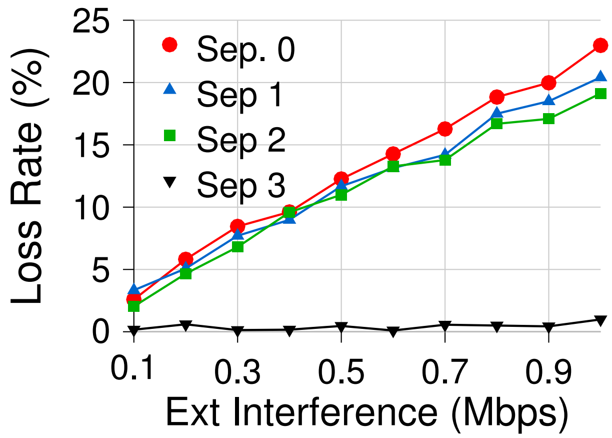

Figure 3.12: Loss rate at different channel separations: Varying interference rate.

3.5.3 Effect of relative power and rate of external interference

To study the effect of relative power and rate of the external WiFi traffic on the loss of the primary link, we perform two experiments using the wireless channel emulator.

In the first experiment, we fix the relative power between the interference source and primary WiLD link, and vary the rate of the external interference source. The received signal strength of the interfering source was approximately 6 dB higher than the primary link traffic. From Figure 3.12 we observe that for channel separations of 0, 1 and 2, the loss rate increases as the rate of the external interference increases. However, beyond a channel separation of 2, there is no significant interference from the external WiFi traffic source and the loss rate is almost zero.

Figure 3.13 shows the variation in loss rate for different relative power levels of the interference source and WiLD link. In this experiment, we fix the power level of the primary WiLD link traffic and vary the relative power of the primary link to the power of

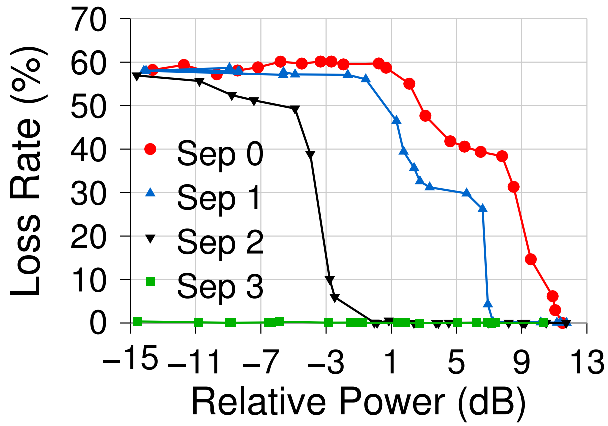

Figure 3.13: Loss rate at different channel separations: Varying interference power.

the interferer from -15 dBm to +13 dBm. The primary link CBR traffic rate is fixed at 512 Kbps, while the interferer transmits at a rate of 3 Mbps.

We observe that when the interference source is on the same channel, even a 12 dB lower signal could lead to packet loss on the primary WiLD link. When the interference source is significantly higher than the WiLD link (6 dB and beyond), the loss rate is very high ($\geq$50%) for channel separations 0, 1 and 2. This corresponds to the situation where any collision results in the capture of the packet on the primary link. Beyond a channel separation of 2, we do not observe any loss on the primary link.

Implications:

- We conclude that external WiFi interference is a significant source of packet loss in WiLD networks. Any deployment of WiLD networks in dense urban deployments has to take into account external WiFi interference.
- When calculating the link budget for urban links, it is beneficial to over-provision the

received power. A high signal strength could potentially immunize the WiLD link from external WiFi traffic.

- MAC-layer adaptation algorithms like adaptive channel switching, rate adaptation, and adaptive FEC could significantly reduce the loss due to external WiFi interference.

3.6 Non-WiFi Interference

The 802.11b communication protocol operates in the 2.4 GHz shared ISM band. This frequency band is shared with a host of other non-802.11 devices, such as microwave ovens, cordless phones, baby monitors, etc. Most of these non-802.11 devices do not follow a channel-access protocol that could lead to a significant amount of interference caused by these devices.

Sheth et al. [127, 128] were able to detect and measure non-WiFi interference by sampling the noise floor of the Atheros chipset. The authors observed that in presence of external non-WiFi noise, the noise floor linearly increases with increasing noise. We performed the same experiment on our WiLD testbed, where we sample the noise floor for every packet received. In presence of external noise causing high loss, we would expect the noise floor to be correlated with the loss rate. However, based on extensive measurements carried out on the urban links we do not see any correlation between noise floor and loss rate. In fact, the noise floor remains mostly constant with minor 1–2 dB variations.

In addition to the above test, we also check for wide-band non-WiFi noise. A wide-band noise source would cause interference across the entire 802.11 spectrum. Ideally, this can be measured using a spectrum analyzer and detecting a rise in power across the entire spectrum. However, using a spectrum analyzer is infeasible on the outdoor WiLD links. Thus, to detect wide band noise in our WiLD deployment we synchronize the two ends of a link to rotate across channel 1, 6 and 11 periodically. The sender generates 1 Mbps UDP CBR traffic on each channel and the receiver measures the loss rate on each channel. In presence of any wide-band noise, we would expect to observe a correlation among loss rates across all three channels. However, based on long-term experiments performed on three urban links, we determined that there was no statistically significant correlation, and thus no significant broadband noise.

3.7 Multipath Interference

Multipath interference is a well known source of packet loss in WiFi networks [4, 33]. It occurs when a RF signal takes different paths from a source to a destination node. Hence, along with the primary line-of-sight signal, the receiver also receives multiple secondary reflections that causes inter-symbol interference (ISI) which might lead to packet corruption. Based on the experiments performed on our WiLD deployments, we conclude that unlike urban mesh deployments, the order-of-magnitude lower delay spreads in WiLD deployments significantly reduces the interference due to multipath.

The two factors contributing to lower delay spreads in WiLD networks are the long distances between the two end hosts and the line-of-sight deployment of the nodes. The strong line-of-sight component in WiLD deployments ensures that the attenuation of the primary signal is only due to path loss, and most of the secondary paths are due to reflections from the ground. In comparison to our WiLD deployment, an urban mesh-

Dist. (km)	Delay spread (μsec)
0.5	(4.75, 3.59)
1.0	(2.4, 1.80)
8.0	(0.3, 0.22)
16.0	(0.15, 0.11)
100.0	(0.02, 0.01)

Table 3.1: Delays between a primary and secondary reflection at midway and quarter-way point.

network deployment (like Roofnet) has shorter and many non-line-of-sight links.

Table 3.1 shows the delay between the primary path and secondary path assuming the antenna is mounted at a height of 30 meters and reflection is only from the ground. The two delays are computed for a secondary path reflecting at the midway point and at the quarter point respectively between the transmitter and the receiver. As the length of the link increases, the primary and the secondary path travel almost the same distance, and hence the delay between the primary and secondary reflection reduces. This reduces the probability of inter-symbol interference. As seen from the table, there is an order-of-magnitude difference between the delay in WiLD links and medium range mesh-network style links. Aguayo et al. [4] also observed that the RAKE receiver is able to tolerate delay spreads upto 0.3–0.4 μs.

Our hypothesis was that most of the loss in our urban deployment was due to external WiFi interference. Hence, in absence of external interference the WiLD links deployed in the rural areas should not have any loss. Figure 3.8 validates our hypothesis, where rural links have a very low loss as compared to urban links.

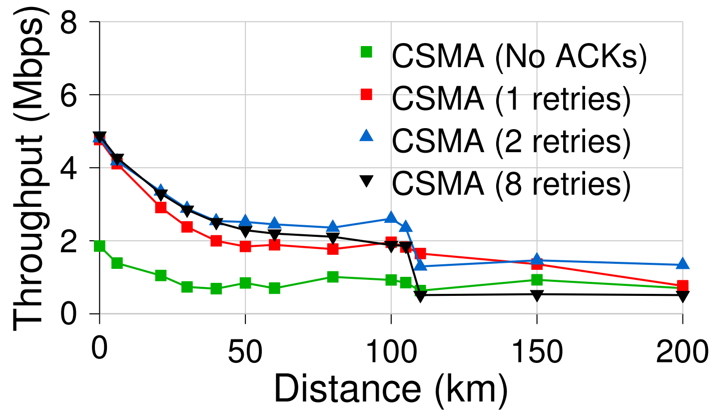

Figure 3.14: Cumulative throughput for TCP in both directions simultaneously over standard CSMA with 10% channel loss on emulated link. Traffic is 802.11b at PHY layer datarate of 11 Mbps.

3.8 Impact on TCP

Taken together, the protocol shortcomings of 802.11 and channel induced losses significantly lower end-to-end TCP performance. The use of stop-and-wait over long distances reduces channel utilization. In addition, we see correlated bursty collision losses due to interference from unsynchronized transmissions (over both single-link and multi-hop scenarios) as well as from external WiFi sources. Under these conditions, TCP flows often timeout resulting in very poor performance. The only configurable parameter in the driver is the number of packet retries. Setting a higher value on the number of retries decreases the loss rate, but at the cost of lower throughput resulting from lower channel utilization.

To better understand this trade-off, we measure the aggregate throughput of TCP flows in both directions on an emulated link while varying distance and introducing a channel packet loss rate of 10%. Figure 3.14 presents the aggregate TCP throughput with various number of MAC retries of the standard 802.11 MAC. Due to increased collisions and larger

ACK turnaround times, throughput degrades gradually with increasing distances.

3.9 Related Work

While there have been several research works on packet loss characterization and methodologies, here, we only focus on those works which are closely related to our work.

Other WiLD deployments: Raman et al. [17] were among the first to deploy a WiLD network consisting of approximately 10 links and lengths ranging from 1–16 km. They also studied the behavior of WiLD links for varying packet sizes, data rates, link lengths, SNRs and weather conditions [30]. This is the only other detailed performance study of WiLD links of which we are aware. Based on their study the authors also experienced high loss due to external interference. In this chapter, we present a comprehensive study of the most common sources of packet loss by the wireless channel and the stock 802.11 protocol. Raman el al. subsequently present modifications to the stock 802.11 MAC protocol to enable point-to-multipoint synchronous transmission and reception in WiLD networks [112]. In our own work, presented in the next chapter, we use lessons from these measurement studies to design a TDMA-based MAC protocol with a synchronization mechanism that is more robust in lossy conditions and with adaptive link loss recovery using bulk ACKs and FEC [104].

Other measurement based studies: Aguayo et al. [4] present a detailed link layer measurement for a outdoor 802.11 mesh deployment, in which they identify the sources of packet loss. Our study indicates that WiLD deployments are faced with a different set of problems as compared to an outdoor 802.11 mesh deployment.

A large number of measurement based studies have also been carried out to study the source of packet loss in indoor large scale 802.11 deployments [55, 69, 70, 120]. Jardosh et al. [70] and Rodrig et al. [120] study the performance of 802.11 in conference settings, where a large number of clients are using the wireless network. The authors observed both short- and long-term variability in link quality and performance degradation under heavy usage of the wireless network. The authors also point out that rate fallback exacerbates the link quality, leading to a higher number of retransmissions and dropped frames.

3.10 Summary

We perform a detailed study of channel induced (WiFi, non-Wifi, and multipath interference) and protocol induced (timeouts, breakdown of CSMA) losses in WiLD settings. Our main result is that most of the losses arise due to external WiFi interference on same and adjacent channels. This result is in contrast to loss studies of urban mesh networks, where multipath is reported to be the most significant source of loss. We also show that 802.11b protocol limitations make it unsuitable not just for point-to-multipoint links, as claimed in prior work, but also unsuitable for simple point-to-point links.

Chapter 4

High Performance Point-to-Point Long Distance Wireless Networks

4.1 Background

As we saw in the previous chapter, the real-world performance of long-distance WiFi links is often below expectations. We found out that there are two main reasons for this poor performance. First, the stock 802.11 protocol has fundamental *protocol shortcomings* that make it ill-suited for WiLD environments. Three specific shortcomings include: (a) the *802.11 link-level recovery* mechanism results in low utilization; (b) at *long distances frequent collisions occur because of the failure of CSMA/CA*; (c) WiLD networks experience *inter-link interference* which introduces the need for synchronizing packet transmissions at each node [112]. The second problem is that the links in our WiLD network deployments (in US, India, Ghana) experienced very *high and variable packet loss rates* induced by external factors (primarily external WiFi interference in our deployment); under such high loss conditions, TCP flows hardly progress and continuously experience timeouts.

The above problems will be conspicous for any link of medium to long distance whether it is part of the point-to-point backbone or the point-to-multipoint access network. In this chapter, we describe the design and implementation of WiLDMAC, system that addresses all the aforementioned problems and provides enhanced end-to-end performance for the point-to-point backbone network.

Prior to our study, the only work addressing this problem was 2P [112], a MAC protocol proposed by Raman et al. The 2P design primarily addresses inter-link interference, and proposes a TDMA-style protocol with synchronous node transmissions. The design of WiLDMAC leverages and builds on top of 2P, making additional changes to further improve link utilization and to make the system robust to packet loss.

Apart from the design and implementation of WiLDMAC, we have had two years experience in deploying and maintaining two production WiLD networks in India and Ghana that support real users. Our network at the Aravind Eye Hospital, India, provides interactive patient-doctor video-conferencing services between the hospital and nine surrounding villages (10–25 km away from the hospital). It is currently being used for about 2000 remote patient examinations per month. The design of WiLDMAC that is presented here has continuously evolved in the past two years to solve many of the performance problems that we faced in our deployments.

4.2 Contributions

Our aim is to build MAC-layer mechanisms to achieve high throughput in multihop point-to-point wireless networks that operate on the same channel. The key contributions

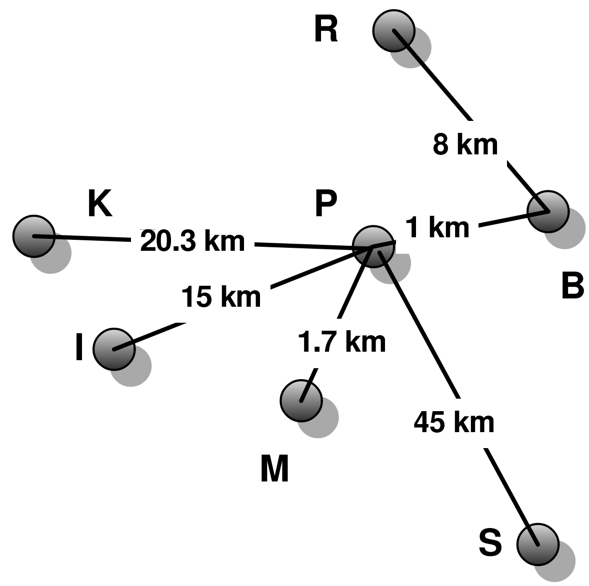

Figure 4.1: Overview of the WiLD campus testbed (not to scale)

of this chapter can be summarized as below.

- **Improving link utilization using bulk acknowledgments:** The current 802.11 protocol uses a stop-and-wait link recovery mechanism, which when used over long distances with high round-trip times leads to under-utilization of the channel. To improve link utilization, WiLDMAC uses a bulk packet acknowledgment protocol.

- **Synchronous operation:** With multiple adjacent directional links at a node that operate on the same channel, WiLDMAC imposes synchronization among them. The nodes use a largely *interference-free* mode of operation termed as Simultaneous Synchronized Operation (*SynOp*) (proposed by Raman et al. [112]), where they either transmit simultaneously (*SynTx*), or receive simultaneously (*SynRx*). We show that the throughput achieved by WiLDMAC is almost the same, whether multiple directional links at a node operate on the same channel or on non-overlapping channels.

- **Designing TDMA in lossy environments:** The stock 802.11 CSMA/CA mechanism is inappropriate for WiLD settings since it cannot assess the state of the channel at the receiver. 2P proposed a basic TDMA mechanism (instead of CSMA/CA) that explicitly synchronized transmissions at each node to prevent inter-link interference. However, with high packet loss rates, explicit synchronization can lead to deadlock scenarios due to loss of synchronization marker packets. In WiLDMAC, we use an implicit approach, using loose time synchronization among nodes to determine a TDMA schedule that is not affected by packet loss.

- **Handling high packet loss rates:** In our WiLD network deployments, we found that external WiFi interference is the primary source of packet loss. The emergence of many WiFi deployments, even in developing regions, will exacerbate this problem. In WiLDMAC, we use an adaptive loss-recovery mechanism that uses a combination of FEC and bulk acknowledgments to significantly reduce the perceived loss rate and to increase the end-to-end throughput. We show that WiLDMAC's link-layer recovery mechanism is much more efficient than a higher-layer recovery mechanisms such as Snoop [14].

- **Application-based parameter configuration:** Different applications have varying requirements in terms of bandwidth, loss, delay and jitter. In WiLDMAC, configuring the TDMA and recovery parameters (time slot period, FEC, number of retries) provides a tradeoff spectrum across different end-to-end properties. We explore these tradeoffs and show that WiLDMAC can be configured to suit a wide range of goals.

We have implemented all our modifications as a *shim layer* above the driver using the Click modular router [78]. We have also deployed WiLDMAC in our campus testbed of 6 long-distance wireless links. Figure 4.1 shows the topology of our campus testbed.

Using a detailed performance evaluation, we roughly observe a 2–5 fold improvement in the TCP throughput over WiLDMAC in comparison to the best achievable TCP throughput obtained by making minor driver changes to the standard 802.11 MAC across a wide variety of settings. On our outdoor testbed, we get upto 5 Mbps of TCP throughput over 3 hops under lossy channel conditions, which is 2.5 times more than that of standard 802.11b. The bandwidth overhead of our loss-recovery mechanisms is minimal.

4.3 WiLDMAC Design

In this section, we describe the design of WiLDMAC and elaborate on how it addresses the 802.11 protocol shortcomings as well as achieves good performance in high-loss environments. In the previous section, we identified three basic problems with 802.11; (a) low utilization, (b) collisions at long distances, and (c) inter-link interference. To address the problem of low utilization, we propose the use of bulk packet acknowledgments (Section 4.3.1). To mitigate loss from collisions at long distances as well as inter-link interference, we replace the standard CSMA MAC with a TDMA-based MAC protocol. We build upon 2P [112] to adapt it to high-loss environments (Section 4.3.2). Additionally, to handle the challenge of high and variable packet losses, we design adaptive loss recovery mechanisms that use a combination of FEC and retransmissions with bulk acknowledgments (Section 4.3.3).

WiLDMAC follows three main design principles. First, the system should not be narrowly focused to a single set of application types. It should be configurable to provide a broad tradeoff spectrum across different end-to-end properties including delay, bandwidth, loss, reliability and jitter. Second, all mechanisms proposed should be implementable on commodity off-the-shelf 802.11 cards. Third, the design should be lightweight, such that it can be implemented on the resource-constrained single-board computers (266-MHz CPU and 128 MB memory) used in our testbed.

4.3.1 Bulk Acknowledgments

We begin with the simple case of a single WiLD link, with each node having a half-duplex radio. As shown earlier, when propagation delays become longer, the default CSMA mechanism cannot determine whether the remote peer is sending a packet in time to back-off its own transmission and avoid collisions. Moreover, such a contention-based mechanism is overkill when precisely two hosts share the channel for a directional link.

Thus, a simple and efficient solution to avoid these collisions is to use an echo protocol between the sender and the receiver, which allows the two end-points to take turns sending and receiving packets. Hence, from a node's perspective, we divide time into send and receive time slots, with a burst of several packets being sent from one host to its peer in each slot.

Consequently, to improve link utilization, we replace the stock 802.11 stop-and-wait protocol with a sliding-window based flow-control approach in which we transmit a *bulk acknowledgment* (bulk ACK) from the receiver for a window of packets. We generate a bulk ACK as an aggregated acknowledgment for all the packets received within the previous slot. In this way, a sender can rapidly transmit a burst of packets rather than wait for an ACK after each packet.

The bulk ACK can be either piggybacked on data packets sent in the reverse direction, or sent as one or more stand-alone packets if no data packets are ready. Each bulk ACK contains the sequence number of the last packet received in order and a variable-length bit vector ACK for all packets following the in-order sequence. Here, the sequence number of a packet is locally defined between the pair of end-points of a WiLD link.

Like 802.11, the bulk ACK mechanism is not designed to guarantee perfect reliability. 802.11 has a maximum number of retries for every packet. Similarly, upon receiving a bulk ACK, the sender can choose to advance the sliding window skipping unacknowledged packets if the retry limit is exceeded. In practice, we support different retry limits for packets of different flows. The bulk ACK mechanism introduces packet reordering at the link layer, which may not be acceptable for TCP traffic. To handle this, we provide in-order packet delivery at the link layer either for the entire link or at a per-flow basis.

4.3.2 Designing TDMA on Lossy Channels

To address the inappropriateness of CSMA for WiLD networks, 2P [112] proposes a contention-free TDMA-based channel access mechanism. 2P eliminates inter-link interference by synchronizing all the packet transmissions at a given node (along all links which operate on the same channel channel). In 2P, a node in transmission mode simultaneously transmits on all its links for a globally known specific period, and then explicitly notifies the end of its transmission period to each of its neighbors using marker packets. A receiving node waits for the marker packets from all its neighbors before switching over to transmis-

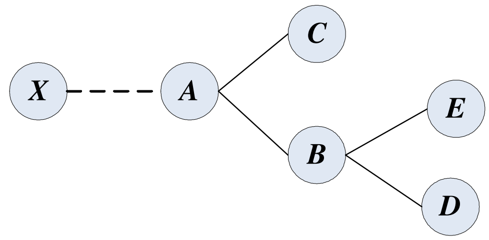

Figure 4.2: Example topology to compare synchronization of 2P and WiLDMAC.

sion mode. In the event of a loss of a marker packet, a receiving node uses a timeout to switch into the transmission mode.

The design of 2P, while functional, is not well suited for lossy environments. Consider the simple example illustrated in Figure 4.2, where all links operate on the same channel. Consider the case where (X, A) is the link experiencing high packet loss-rate. Let T denote the value of the time-slot. Whenever a marker packet transmitted by X is lost, A begins transmission only after a timeout period T_0 $(\geq T)$. This, in turn, delays the next set of transmissions from nodes B and C to their other neighbors by a time period that equals $T_0 - T$. Unfortunately, this propagation of delay does not end here. In the time slot that follows, D's transmission to its neighbors is delayed by $T_0 - T$. Hence, what we observe is that the loss of marker packets has a "ripple effect" in the entire network creating an idle period of $T_0 - T$ along every link. When markers along different links are dropped, the ripples from multiple links can interact with each other and cause more complex behavior.

Ideally, one would want $T_0 - T$ to be very small. If all nodes are perfectly time synchronized, we can set $T_0 = T$. However, in the absence of global time synchronization, one needs to set a conservative value for T_0. 2P chooses $T_0 = 1.25 \times T$. The loss of a marker packet leads to an idle period of $0.25 \times T$ (in 2P, this is 5 ms for $T = 20$ ms). In bursty losses, transmitting multiple marker packets may not suffice.

Given that many of the links in our network experience sustained loss-rates over 5–40%, in WiLDMAC, we use an implicit synchronization approach that aims to reduce the value of $T_0 - T$. In WiLDMAC, we use a simple loose time synchronization mechanism similar to the basic linear time synchronization protocol NTP [85], where during each time slot along each link, the sender acts as the master and the receiver as the slave. Consider a link (A, B) where A is the sender and B is the receiver at a given time. Let t_{send_A} and t_{recv_B} denote the start times of the slot as maintained by A and B. All the packets sent by A are timestamped with the time difference (δ) between the time the packet has been sent (t_1) and the beginning of the send slot(t_{send_A}). When a packet is received by B at time t_2, the beginning of B's receiving slot is adjusted accordingly: $t_{recv\_B} = t_2 - \delta$. As soon as B's receive slot is over, and $t_{send\_B} = t_{recv\_B} + T$ is reached, B starts sending for a period T.

Due to the propagation delay between A and B, the send and corresponding receive slots are slightly skewed. The end-effect of this loose synchronization is that the value of T_0-T is limited by the propagation delay across the link even with packet losses (assuming clock speeds are roughly comparable). Hence, an implicit synchronization approach significantly reduces the value of $T_0 - T$ thereby reducing the overall number of idle periods in the network.

4.3.3 Adaptive Loss Recovery

To achieve predictable end-to-end performance, it is essential to have a loss recovery mechanism that can hide the loss variability in the underlying channel. Achieving such an upper bound (q) on the loss-rate perceived by higher level applications is not easy in our settings. First, it is hard to predict the arrival and duration of bursts. Second, the loss distribution that we observed on our links is non-stationary even on long time scales (hourly and daily basis). Hence, a simple model cannot capture the channel loss characteristics.

In WiLDMAC, we can either use retransmissions or FEC to deal with losses (or a combination of both). A retransmission based approach can achieve the loss-bound q with minimal throughput overhead but at the expense of increased delay. An FEC based approach incurs additional throughput overhead but does not incur a delay penalty especially since it is used in combination with TDMA on a per-slot basis. However, an FEC approach cannot achieve arbitrarily low loss-bounds mainly due to the unpredictability of the channel.

Tuning the Number of Retransmissions: To achieve a loss bound q independent of underlying channel loss rate $p(t)$, we need to tune the number of retransmissions. One can adjust the number of retransmissions $n(t)$ for a channel loss-rate $p(t)$ such that $(1 - p(t))^{n(t)} = q$. Given that our WiLD links support in-order delivery (on a per-flow or on whole link basis), a larger $n(t)$ also means a larger maximum delay, equal to $n(t) * T$ for a slot period T. One can set different values of $n(t)$ for different flows. We found that estimating $p(t)$ using an exponentially weighted average is sufficient in our links to achieve the target loss estimate q. A purely retransmission based recovery mechanism has minimal throughput overhead as only the lost packets are retransmitted but this comes at a cost of high delay due to the long round-trip times over WiLD links.

Adaptive FEC-Based Recovery: Designing a good FEC mechanism in highly variable

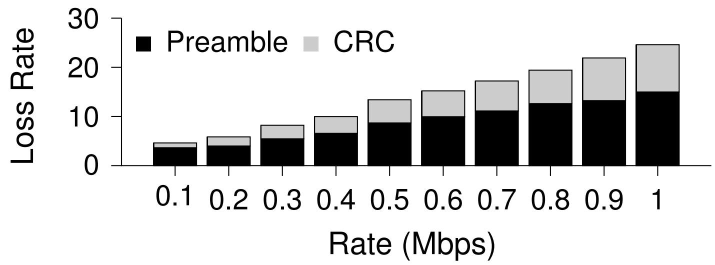

Figure 4.3: Proportion of CRC and preamble errors in channel loss. Traffic is at UDP CBR packets of 1440 bytes each at 802.11b PHY datarate of 11 Mbps. Main link is sending at 2 Mbps. The sending rate of the interferer increases from 0.1 Mbps to 1 Mbps.

lossy conditions requires accurate estimation of the underlying channel loss. When the loss is underestimated, the redundant packets cannot be decoded at all making them useless, but overestimating the loss rate leads to unnecessary overhead.

Motivating inter-packet FEC:: We can perform two types of FEC: inter-packet FEC (coding across packets) or intra-packet FEC (coding redundant blocks within a packet). Based on extensive measurements on a wireless channel emulator we observe that in presence of external WiFi interference, lost packets can be categorized into either CRC errors or preamble errors. A CRC error packet is received by the driver with a check sum error. However, an error in the preamble leads to the entire packet being dropped completely. Figure 4.3 shows the breakup of the loss rate with increasing external interference. We observe although the proportion of preamble errors decreases as external interference increases, it still causes at least 50% of all errors. Moreover a substantial number of the CRC error packets were truncated. We choose not to perform intra-packet FEC because it can only help recover packets that have CRC errors. Hence, we chose to perform inter-packet FEC.

Estimating redundancy:: We apply FEC in combination with TDMA. For every time slot of N packets, we add $N - K$ redundant packets to K original packets. To estimate the redundancy factor, $r = (N - K)/K$, we choose a simple but not perfect estimation policy based on a weighted average of the losses observed in the previous M time slots. Here, we specifically chose a small value of $M = 10$ because it is hard to predict the start of a burst. Secondly, a small value of M, can quickly adapt to both the start and end of a loss burst saving unnecessary redundant FEC packets. For a time slot of $T = 10ms$, $M = 10$ corresponds to $200ms$ (with symmetric slot allocation in both directions) to adapt to a change in the loss behavior. Also due to non-stationary loss distributions, the benefit of using more complicated distribution based estimation approaches [134] is marginal. This type of FEC is best suited for handling residual losses and bursts that are longer than the time required for loss estimation mechanism to adapt.

4.4 Implementation

In this section, we describe the implementation details of WiLDMAC. Our implementation comprises two parts: (a) driver-level modifications to control or disable features implemented in hardware (Section 4.4.1); (b) a *shim* layer that sits above the 802.11 MAC (Section 4.4.2) and uses the Click [78] modular router software to implement the functionalities described in Section 4.3.

4.4.1 Driver Modifications

The wireless cards we use in our implementation are the high-power (200-400 mW) Atheros-based chipsets. To implement WiLDMAC, we have to disable the following 802.11 MAC mechanisms:

- We disable **link-layer association** in Atheros chipsets using the *AdHoc-demo* mode.

- We disable **link layer retransmissions and automatic ACKs** by using 802.11 QoS frames with WMM extensions set to the no-ACK policy.

- We disable **CSMA** by turning off the Clear Channel Assessment (CCA) in Atheros chipsets. With CCA turned off, the radio card can transmit packets right away without waiting for a clear channel.

4.4.2 Software Architecture Modifications

In order to implement single-link and inter-link synchronization using TDMA, the various loss recovery mechanisms, sliding-window flow control, and packet reordering for in-order delivery, we use the Click modular router [78] framework. We use Click because it enables us to prototype quickly a modular MAC layer by composing different Click elements together. It is also reasonably efficient for packet processing especially if loaded as a kernel module. Using kernel taps, Click creates fake network interfaces, such as *fake0* in Figure 4.4 and the kernel communicates with these virtual interfaces. Click allows us to intercept packets sent to this virtual interface and modify them before sending them on the real wireless interface and vice versa.

Figure 4.4 presents the structure of the Click elements of our layered system, with different functionality (and corresponding packet header processing) at various layers:

Incoming/Outgoing Queues: The mechanisms for sliding window packet flow, bulk ACKs, selective retransmission and reordering for in-order delivery are implemented by the incoming/outgoing queue pair. Packet buffering at the sender is necessary for retransmissions,

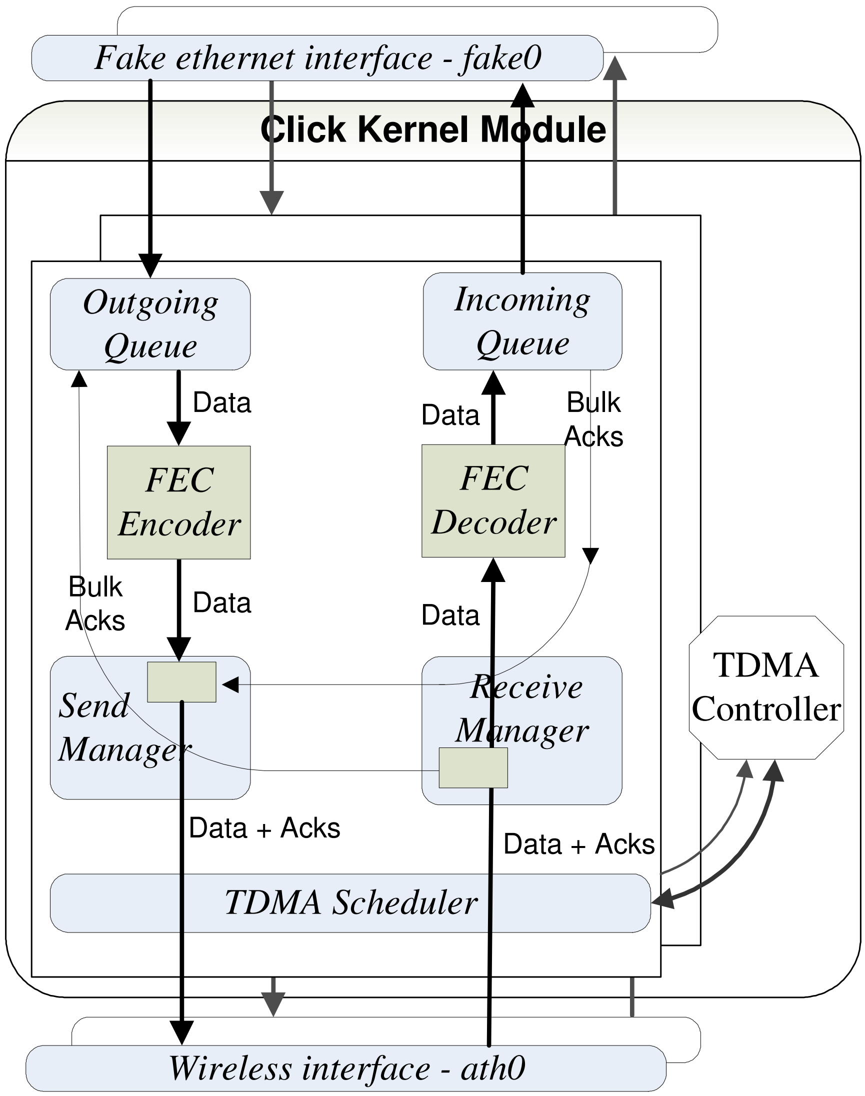

Figure 4.4: Click Module Data Flow

and buffering at the receiver enables reordering. In-order delivery and packet retransmission are optional, and the number of retries can be set on a per-packet basis.

FEC Encoder/Decoder: An optional layer is responsible for inter-packet forward error correction encoding and decoding. For our implementation we modify a FEC library [118] that uses erasure codes based on Vandermonde matrices computed over $GF(2^m)$. This

FEC method uses a (K, N) scheme, where the first K packets are sent in their original form, and $N - K$ redundant packets are generated, for a total of N packets sent. At the receiver, the reception of any K out of the N packets enables the recovery of the original packets. We choose this scheme because, in loss-less situations, it introduces very low latency: the original K packets can be immediately sent by the encoder (without undergoing encoding), and immediately delivered to the application by the decoder (without undergoing decoding).

TDMA Scheduler and Controller: The Scheduler ensures that packets are being sent only during the designated send slots, and manages packet timestamps as part of the synchronization mechanism. The Controller implements synchronization among the wireless radios, by enforcing synchronous transmit and receive operation (all the radios on the same channel have a common send slot, followed by a common receive slot).

4.4.3 Timing issues

We do not use Click timers to implement time synchronization because the underlying kernel timers are not precise at the granularity of our time slots (10 ms-40 ms) on our hardware platform (266 MHz CPU). Also packet queuing in the wireless interface causes variability in the time between the moment Click emits a packet and the time the packet is actually sent on the air interface. Thus, the propagation delay between the sending and the receiving click modules on the two hosts is not constant, affecting time slot calculations. Fortunately, this propagation delay is predictable for the first packet in the send slot, when the hardware interface queues are empty. Thus, in our current implementation, we only timestamp the first packet in a slot, and use it for adjusting the receive slot at the peer.

If this packet is lost, the receiver's slot is not adjusted in the current slot, but since the drift is slow this does not have a significant impact. In the future we intend to perform this timestamping in the firmware - that would allow us to accurately timestamp every packet just before packet transmission.

Another timing complication is related to estimating whether we have time to send a new packet in the current send slot. Since the packets are queued in the wireless interface, the time when the packet leaves Click cannot be used to estimate this. To overcome this aspect, we use the notion of *virtual time*. At the beginning of a send slot, the virtual time t_v is same as current (system) time t_c. Every time we send a packet, we estimate the transmission time of the packet on the channel and recompute the virtual time: $t_v = max(t_c, t_v) + duration(packet)$. A packet is sent only after checking that the virtual time after sending this packet will not exceed the end of the send slot. Otherwise, we postpone the packet until the next slot.

4.5 Evaluation

The main goals of WiLDMAC are to increase link utilization and to eliminate the various sources of packet loss observed in a typical multi-hop WiLD deployment, while simultaneously providing flexibility to meet different end-to-end application requirements. We believe these are the first actual implementation results over an outdoor multi-hop WiLD network deployment.

Raman et al. [112] show the improvements gained by the 2P-MAC protocol in simulation and in an indoor environment. However, a multi-hop outdoor deployment also

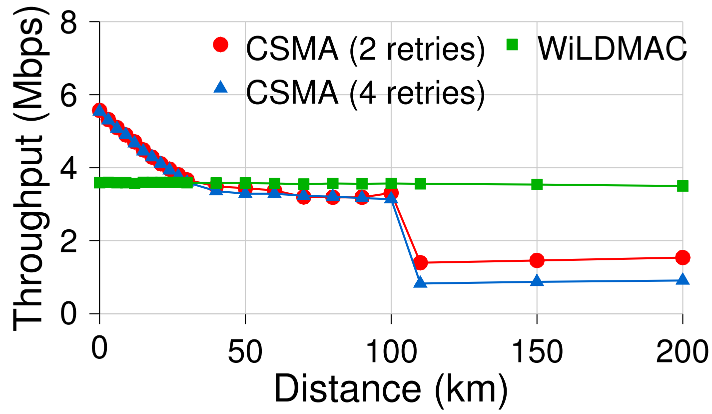

Figure 4.5: TCP throughput for WiLDMAC vs 802.11 CSMA. with flow in one directions 0% channel loss. Each measurement is for a TCP flow of 60 s, 802.11b PHY, 11 Mbps.

has to deal with high losses from external interference. 2P in its current form does not have any built-in recovery mechanism and it is not clear how any recovery mechanism can be combined with the marker-based synchronization protocol. Hence, we do not have any direct comparison results with 2P on our outdoor wireless links. Also, the proof-of-concept implementation of 2P was for the Prism 2.5 wireless chipset and it would be non-trivial to implement the same in WiLDMAC using features of the Atheros chipset.

Our evaluation has three main parts:

- We analyze the ability of WiLDMAC to maintain high performance (high link utilization) over long-distance WiLD links. At long distances, we demonstrate 2–5x improvements in cumulative throughput for TCP flows in both directions simultaneously.

- Next, we evaluate the ability of WiLDMAC to scale to multiple hops and eliminate inter-link interference. WiLDMAC yields a 2.5x improvement in TCP throughput on

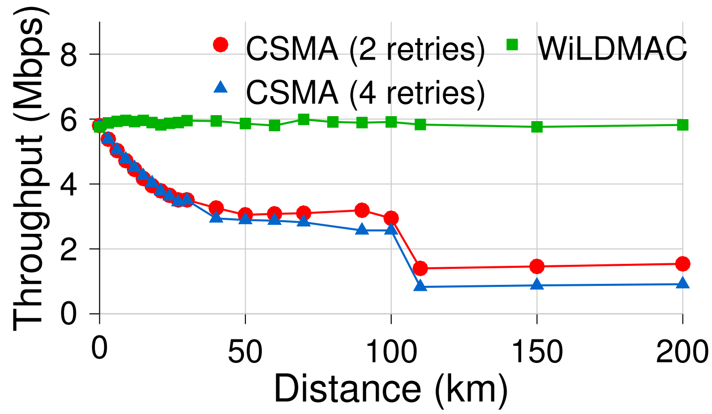

Figure 4.6: TCP throughput for WiLDMAC vs 802.11 CSMA. with flow in both directions, 0% channel loss. Each measurement is for a TCP flow of 60 s, 802.11b PHY, 11 Mbps.

our real-world multi-hop setup.

- Finally, we evaluate the effectiveness of the two link recovery mechanisms of WiLDMAC: Bulk Acks and FEC.

4.5.1 Single Link

In this section we demonstrate the ability of WiLDMAC to eliminate link under-utilization and packet collisions over a single WiLD link. We compare the performance of WiLDMAC (slot size of 20 ms) with the standard 802.11 CSMA (2 retries) base case.

The first set of results show the improvement of WiLDMAC on a single emulator link with increasing distance. Figure 4.5 compares the performance of TCP flowing only in one direction. The lower throughput of WiLDMAC, approximately 50% of channel capacity, is due to symmetric slot allocation between the two end points of the link. However, over longer links (>50 km), the TDMA-based channel allocation avoids the under-utilization of the link as experienced by CSMA. Also, beyond 110 km (the maximum possible ACK

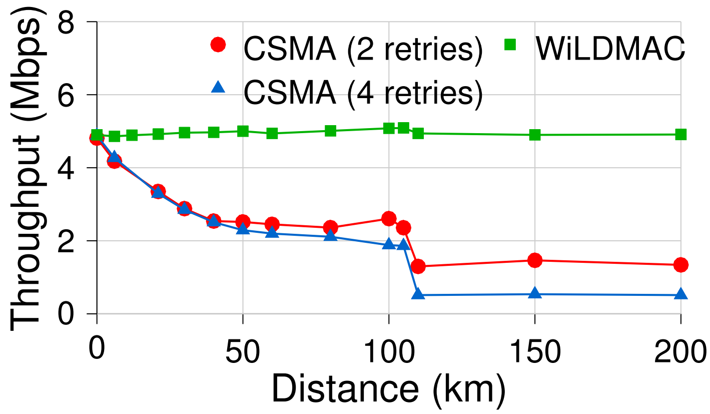

Figure 4.7: TCP throughput for WiLDMAC vs 802.11 CSMA, with flow in both directions, 10% channel loss. Each measurement is for a TCP flow of 60 s, 802.11b PHY, 11 Mbps.

timeout), the throughput with CSMA drops rapidly because of unnecessary retransmits. Figure 4.6 shows the cumulative throughput of TCP flowing simultaneously in both directions. In this case, WiLDMAC effectively eliminates all collisions occurring in presence of bidirectional traffic. TCP throughput of 6 Mbps is maintained for all distances.

Table 4.1 compares WiLDMAC and CSMA for some of our outdoor wireless links. We show TCP throughput in one direction and the cumulative throughput for TCP simultaneously flowing in both directions. Since these are outdoor measurements, there is significant variation over time and we show both the mean and standard deviation for the measurements. We can see that as the link distance increases, the improvement of WiLDMAC is more substantial. Infact, for the 65 km link in Ghana, WiLDMAC's throughput at 5.5 Mbps is about 8x better than standard CSMA.

Link	Distance	Loss rates	802.11 CSMA (Mbps)		WiLDMAC (Mbps)	
	(km)	(%)	One dir	Both dir	One dir	Both dir
B-R	8	3.4	5.03 (0.02)	**4.95** (0.03)	3.65 (0.01)	**5.86** (0.05)
P-S	45	2.6	3.62 (0.20)	**3.52** (0.17)	3.10 (0.05)	**4.91** (0.05)
Ghana	65	1.0	2.80 (0.20)	**0.68** (0.39)	2.98 (0.19)	**5.51** (0.07)

Table 4.1: Mean TCP throughput (flow in one direction and cumulative for both directions simultaneously) for WiLDMAC and CSMA for various outdoor links (distance and loss rates). The standard deviation is shown in parenthesis for 10 measurements. Each measurement is for TCP flow of 30s at a 802.11b PHY-layer datarate of 11 Mbps.

4.5.2 Multiple Hops

This section validates that WiLDMAC eliminates inter-link interference by synchronizing receive and transmit slots in TDMA resulting in up to 2x TCP throughput improvements over standard 802.11 CSMA in multi-hop settings.

The first set of measurements were performed on our indoor setup where we recreated the conditions of a linear outdoor multi-hop topology using the RF isolation boxes. Thus transmissions from local radios interfere with each other but multiple local radio interfaces can receive simultaneously. We then measure TCP throughput of flows in the one direction and then both directions simultaneously for both standard 802.11 CSMA and WiLDMAC (with slot size of 20 ms). All the links were operating on the same channel. As we see in Table 4.2, as the number of hops increases, standard 802.11's TCP throughput drops substantially when transmissions from a radio collide with packet reception on a nearby local radio on the same node. WiLDMAC avoids these collisions and maintains a much higher cumulative TCP throughput (up to 2x for the 3-hop setup) by proper syn-

Linear setup	802.11 CSMA (Mbps)			WiLDMAC (Mbps)		
	Dir 1	Dir 2	Both	Dir 1	Dir 2	Both
2 nodes	5.74 (0.01)	5.74 (0.01)	**6.00** (0.01)	3.56 (0.03)	3.53 (0.02)	**5.85** (0.07)
3 nodes	2.60 (0.01)	2.48 (0.01)	**2.62** (0.01)	3.12 (0.01)	3.12 (0.01)	**5.12** (0.03)
4 nodes	2.23 (0.01)	2.10 (0.01)	**1.99** (0.02)	2.95 (0.05)	2.98 (0.04)	**4.64** (0.24)

Table 4.2: Mean TCP throughput (flow in each direction and cumulative for both directions simultaneously) for WiLDMAC and standard 802.11 CSMA. Measurements are for linear 2,3 and 4 node indoor setups recreating outdoor links running on the same channel. The standard deviation is shown in parenthesis for 10 measurements of flow of 60 s each at 802.11b PHY-layer datarate of 11 Mbps.

chronization of send and receive slots.

We can also see that although WiLDMAC has more than 2x improvement over standard 802.11, the final throughput (4.6 Mbps) is still much smaller than the raw throughput of the link (6-7 Mbps). This can be attributed to the overhead of synchronization and packet processing in Click running on our low-power (266 MHz) single board routers. A more efficient synchronization mechanism implemented in the firmware (rather than Click) would deliver much better improvement.

We also measure this improvement on our outdoor testbed between the nodes K and M relayed through node P. We again compare the TCP throughput for WiLDMAC and standard 802.11 CSMA with links operating on the same channel. In order to quantify the effect of inter-link interference, we also perform the same experiments with the links operating on different, non-overlapping channels, in which case the inter-link interference is almost zero, as previously shown in Figures 3.6 and 3.7.

We can see that, for same channel operation, the cumulative TCP throughput in

Description (Mbps)	One direction	Both directions
Standard TCP: same channel	**2.17**	**2.11**
Standard TCP: diff channels	3.95	4.50
WiLD TCP: same channel	**3.12**	**4.86**
WiLD TCP: diff channels	3.14	4.90

Table 4.3: Mean TCP throughput (flow in single direction and cumulative for both directions simultaneously) comparison for WiLDMAC and standard 802.11 CSMA over a 3-hop outdoor setup ($K \leftrightarrow P \leftrightarrow M$). Averaged over 10 measurements of TCP flow for 60 s at 802.11b PHY-layer datarate of 11 Mbps.

both directions with WiLDMAC (4.86 Mbps) is more than twice the throughput observed over standard 802.11 (2.11 Mbps). The improvement is substantially lower for the unidirectional case (3.14 Mbps versus 2.17 Mbps), because the WiLD links are constrained to send in one direction only roughly half of the time.

Another key observation is that WiLDMAC is capable of eliminating almost all inter-link interference. This is shown by the fact that the throughput achieved by WiLDMAC is almost the same, whether the links operate on the same channel or on non-overlapping channels.

4.5.3 Link-Recovery Mechanisms

Our next set of experiments evaluate WiLDMAC's adaptive link recovery mechanisms in conditions closer to the real world, where errors are generated by a combination of collisions and external interference. We evaluate both the bulk ACK and FEC recovery mechanisms.

Bulk ACK Recovery Mechanism: For our first experiment, presented in Figure 4.5.1, we vary the link length on the emulator, and we introduce a 10% error rate through external

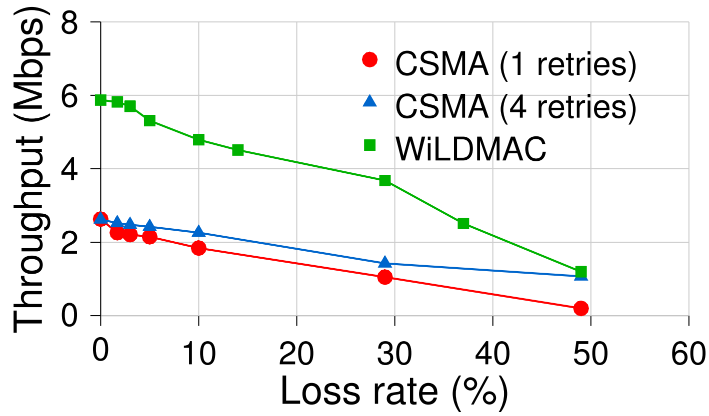

Figure 4.8: Comparison of cumulative throughput for TCP in both directions simultaneously for WiLDMAC and standard 802.11 CSMA with increasing loss on 80 km emulated link. Each measurement was for 60 s TCP flows of 802.11b at 11 Mbps PHY datarate.

interference. We again measure the cumulative throughput of TCP flows in both directions for WiLDMAC and standard 802.11 CSMA. As can be seen, WiLDMAC maintains a constant throughput with increasing distance as opposed to the 802.11 CSMA. Due to the 10% error, WiLD incurs a constant throughput penalty of approximately 1 Mbps compared to the no-loss case in Figure 4.6.

In our second experiment we fix the distance in the emulator setup to 80 km, and vary channel loss rates. The results in Figure 4.8 show that WiLDMAC maintains roughly a 2x improvement over standard CSMA's recovery mechanism for packet loss rates up to 30%.

Forward Error Correction (FEC):

To measure the jitter introduced by the FEC mechanism, we performed a simple experiment where we measured the jitter of a flow under two conditions: in the absence of any loss and in the presence of a 25% loss. Figure 4.9 illustrates the jitter introduced by

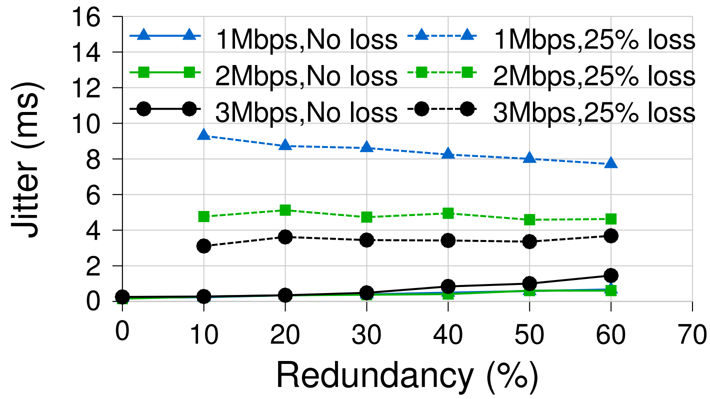

Figure 4.9: Jitter overhead of encoding and decoding for WiLDMAC on single indoor link. Traffic is 1440 byte UDP CBR packets at PHY datarate of 11 Mbps in 802.11b.

WiLDMAC's FEC implementation. We can see that in the absence of any loss, when only encoding occurs, the jitter is minimal. However, in the presence of loss, when decoding also takes place, the measured jitter increases. However, the magnitude of the jitter is very small and well within the acceptable limits of many interactive applications (voice or video), and decreases with higher throughputs (since the decoder waits less for redundant packets to arrive).

Moreover, considering the combination of FEC with TDMA, the delay overheads introduced by these methods overlap, since the slots when the host is not actively sending can be used to perform encoding without incurring any additional delay penalties.

4.6 Tradeoffs

One of the main design principles of WiLDMAC is to build a system that can be configured to adapt to different application requirements. In this section we explore the tradeoff space of throughput, delay and delivered error rates by varying the slot size, number

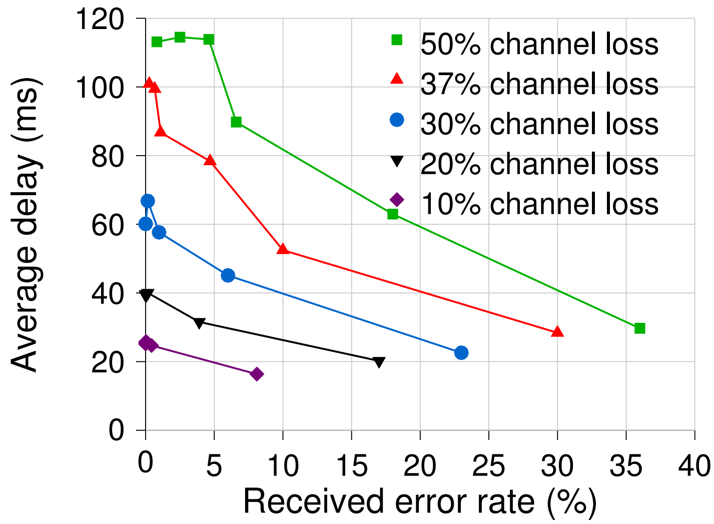

Figure 4.10: Average delay with decreasing target loss rate (X-axis) for various loss rates in WiLDMAC on single emulated 60 km link (slot size=20 ms).

of bulk retransmissions and FEC redundancy parameters. We observe that WiLDMAC can perform in a wide spectrum of the parameter space, and can easily be configured to meet specific application requirements.

4.6.1 Choosing number of retransmissions

The first tradeoff that we explore is choosing the number of retries to get a desired level of final error rate on a WiLD link. Although retransmission based loss recovery achieves optimal throughput utilization, it comes at a cost of increased delay; the loss rate can be reduced to zero by arbitrarily increasing the number of retransmissions at the cost of increased delay. This tradeoff is illustrated in Figure 4.10 which shows a plot of delay versus error rate for varying channel loss rates (10% to 50%). Retries are decreased from 10 to 0 from left to right for a given series in the figure. All the tests are with unidirectional UDP at 1 Mbps for a fixed slot size of 20 ms on a single emulator 60 km link. We can see

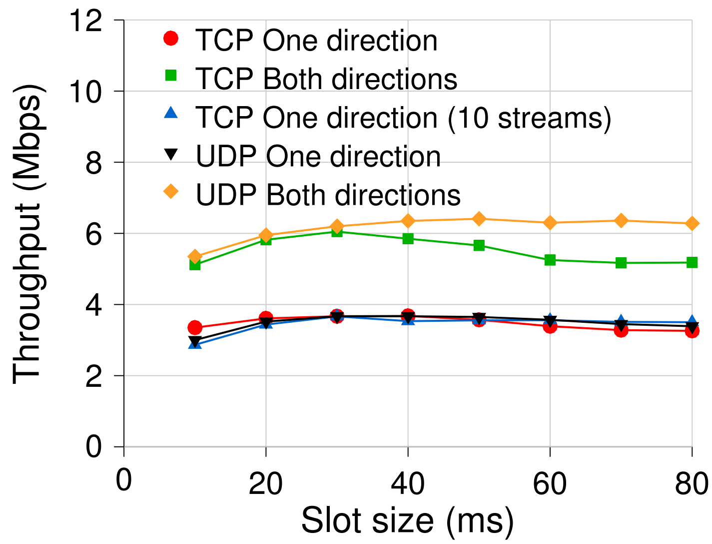

Figure 4.11: Throughput for increasing slot sizes (X-axis) in WiLDMAC for various types of traffic on single emulated 60 km link.

that as we try to reduce the final error rate at the receiver, we have to use more retries and this increases the average delay. In addition, we also observe that larger the number of retries, larger the end-to-end jitter (especially at higher loss rates).

This tradeoff has important implications for applications that are more sensitive to delay and jitter (such as real time audio and video) as compared to applications which require high reliability. For such applications, we can achieve a balance between the final error rate and the average delay by choosing an appropriate retry limit. For applications that require improved loss characteristics without incurring a delay penalty, we need to use FEC for loss recovery.

4.6.2 Choosing slot size

The second tradeoff that we explore is the effect of slot size on TCP and UDP throughput. Our experiments are performed on a 60-km emulated link (Figure 4.11). As discussed in Section 4.3.2, switching between send and receive slots incurs a non-negligible overhead for the Click based WiLDMAC implementation. This overhead although constant for all slot sizes, occupies a higher fraction of the slot for smaller slots sizes. As a consequence, at small slot sizes the achieved throughput is lower. However, the UDP throughput levels off beyond a slot size of 20 ms. We also observe the TCP throughput reducing slightly at higher slot size. This is because the bandwidth-delay product of the link increases with slot size, but the send TCP window sizes are fixed. UDP throughput does not decrease at higher slot sizes.

In the next experiment, we measure the average UDP packet transmission delay while varying the slot size, for several channel error rates. The results are presented in Figure 4.12; each series represents a unidirectional UDP test (1 Mbps CBR) at a particular channel loss rate with WiLDMAC using maximum number of retries. Figure 4.12 shows the increase in delay with increasing slot size. It is clear that slot sizes beyond 20 ms do not result in substantially higher throughputs, but they do result in much larger delay. However, if lower delay is required, smaller slots can be used at the expense of some throughput overhead consumed by the switching between the transmit and receive modes.

4.6.3 Choosing FEC parameters

The primary tunable FEC parameter is the redundancy factor $r = (N-K)/K$, also referred to as throughput overhead. Although FEC incurs a higher throughput overhead

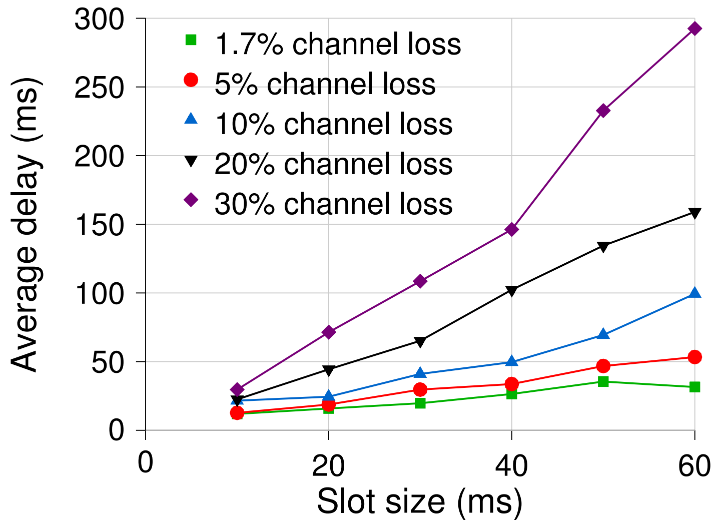

Figure 4.12: Average delay at increasing slot sizes (X-axis) for various loss rates in WiLD-MAC on single emulated 60 km link.

than retransmissions, it incurs a smaller delay penalty as illustrated earlier in Section 4.5.3. To analyze the tradeoff between FEC throughput overhead and the target loss-rate, we consider the case of a single WiLD link (in our emulator environment) with a simple Bernoulli loss-model (every packet is dropped with probability p). Figure 4.13 shows the amount of redundancy required to meet three different target loss-rates of 10%, 5% or 1% as the raw channel error rates (namely p) increase. We see that in order to achieve very low target loss-rates, a lot of redundancy is required (for example, FEC incurs a 100% overhead to reduce the loss-rate from 30% to 1%). Also, when a channel is very bursty and has an unpredictable burst arrival pattern, it is very hard for FEC to achieve arbitrarily low target loss-rates.

For applications that can tolerate one round of retransmissions, we can use a combination of FEC and retransmissions to provide a tradeoff between overall throughput

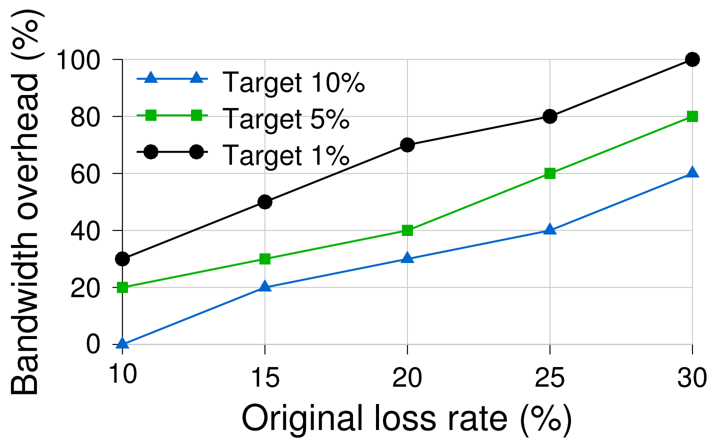

Figure 4.13: Throughput overhead vs channel loss rate for FEC on single emulated 20 km link. Traffic is 1 Mbps CBR UDP.

overhead, delay and target loss-rate. In the case of a channel with a stationary loss distribution, OverQoS [134] shows that the optimal policy to minimize overhead is to not use FEC in the first round but use it in the second round to pad retransmission packets. With unpredictable and highly varying channel loss conditions, an alternative promising strategy is to use FEC in the first round during bursty periods to reduce the perceived loss-rate.

4.7 Related Work

Long Distance WiFi: The use of 802.11 for long-distance networking with directional links and multiple radios per node, raises a new set of technical issues [17]. Raman et al. built upon this work in [111, 112] and proposed the 2P MAC protocol. WiLDMAC builds upon 2P to make it robust in high loss environments. Specifically we modify 2P's implicit synchronization mechanism as well as build in adaptive bulk ACK based and FEC based link recovery mechanisms.

Other wireless loss recovery mechanisms: There is a large body of research literature in wireless and wireline networks that have studied the tradeoffs between different forms of loss recovery mechanisms. Many of the classic error control mechanisms are summarized in the book by Lin and Costello [80]. OverQoS [134] performs recovery by analyzing the FEC/ARQ tradeoff in variable channel conditions and the Vandermonde codes are used for reliable multicast in wireless environments [118].

Of particular interest for this work are the Berkeley Snoop protocol [14, 15] which provides transport-aware link-layer recovery mechanisms in wireless environments. To compare the WiLDMAC bulk ACK recovery mechanism with recovery at a higher layer, we experimented with a version of the original Snoop protocol [14, 15] that we modified to run on WiLD links. Basically, each WiLD router ran one half of Snoop, the fixed host to mobile host part, for each outgoing link and integrated all the Snoops on different links into one module.

We measured the performance of modified Snoop as a recovery mechanism over both standard 802.11 (CSMA) and over WiLDMAC with no retries. We found that WiLDMAC was still 2x better than Snoop. We also saw that Snoop was better than vanilla CSMA only at lower error rates (less than 10%). Thus, this indicates that higher layer recovery mechanisms might be better than stock 802.11 protocol, but only at lower error rates.

Other WiFi-based MAC protocols: Several recent efforts have focused on leveraging off-the-shelf 802.11 hardware to design new MAC protocols. Overlay MAC Layer (OML) [114] provides a deployable approach towards implementing a TDMA-style MAC on top of the 802.11 MAC using loosely-synchronized clocks to provide applications and competing nodes better control over the allocation of time-slots. SoftMAC [96] and FreeMAC [126] are other similar platforms to build experimental MAC protocols. MultiMAC [36] builds on SoftMac to provide a platform where multiple MAC layers co-exist in the network stack and any one can be chosen on a per-packet basis.

WiMax: An alternative to WiLD networks is WiMax [2]. WiMax does present many strengths over a WiFi: configurable channel spectrum width, better modulation (especially for non-line of sight scenarios), operation in licensed spectrum with higher transmit power, and thus longer distances. On the other hand, WiMax currently is primarily intended for carriers (like cellular) and does not support point-to-point operation. In addition, WiMax base-stations are expensive ($10,000) and the high spectrum license costs in most countries dissuades grassroots style deployments. Currently it is also very difficult to obtain licenses for experimental deployment and we are not aware of open-source drivers for WiMax base-stations and clients. However, most of our work in loss recovery and adaptive FEC would be equally valid for any PHY-layer (WiFi or WiMax). With appropriate modifications and cost reductions, WiMax can serve as a more suitable PHY-layer for WiLD networks.

4.8 Summary

The commoditization of WiFi (802.11 MAC) hardware has made WiLD networks an extremely cost-effective option for providing network connectivity, especially in rural regions in developing countries. However providing coverage at high performance in real-world WiLD network deployments raises many research challenges: optimal planning and placement of long-distance links, design of appropriate MAC and network protocols to provide quality of service to a wide variety of applications, remote management and fault tolerance to handle unpredictable node and link failures [135].

One of the most important challenges in this space is the sub-optimal performance of the standard 802.11 MAC protocol. In this chapter, we identify the set of link- and MAC-layer modifications essential for achieving high throughput in multihop point-to-point WiFi based backbone networks. Specifically, using a detailed performance evaluation, we show that the conventional 802.11 protocol is ill-suited for WiLD settings. Our proposed solution provides a 2-5x improvement in TCP throughput over the conventional 802.11 MAC.

Although this constitutes a substantial improvement, designing decentralized TDMA slot scheduling schemes for multi-hop and multi-channel networks to achieve optimal bandwidth and delay characteristics for realistic real-world asymmetric traffic demands is a significant future research direction. Our current solution builds the basic link mechanisms to provide quality of service. End-to-end QoS solutions that leverage these mechanisms and adapt to a realistic traffic mix can be built on top of our mechanisms.

Encouraged by our initial results on our long-distance outdoor testbed, we implemented these modifications in our live rural deployments in India and Ghana. Our network at the Aravind Eye Hospital, India, that provides interactive patient-doctor video-conferencing services between the hospital and nine surrounding villages (10–25 km away from the hospital) is being expanded to reach 50 village eye centers from five main hospitals. We will describe other aspects of these deployments such as operational issues, training of local manpower and financial sustainability in Chapter 6.

Chapter 5

Highly Scalable Point-to-Multipoint Long-Distance Wireless Networks

In this dissertation, we have proposed WiLDNet, a new multi-tier deployment model for wireless networks in rural areas. The WiLDNet architecture consists of a combination of *a)* long-distance (5–150 km) point-to-point high-bandwidth backhaul WiFi links connecting towns and villages, and *b)* medium-range point-to-multipoint access links distributing the connectivity to schools, hospitals, kiosks and individual users. Both types of links use inexpensive, high-gain directional or sector antennas in order to increase the range of outdoor WiFi communication effectively (example network illustrated in Figure 5.1). This deployment model has been explored in many real networks including the Akshaya network [8] in southern India and the AirJaldi network [5], among others [10, 29].

In Chapter 4, we focused on designing protocols to build high-throughput long-

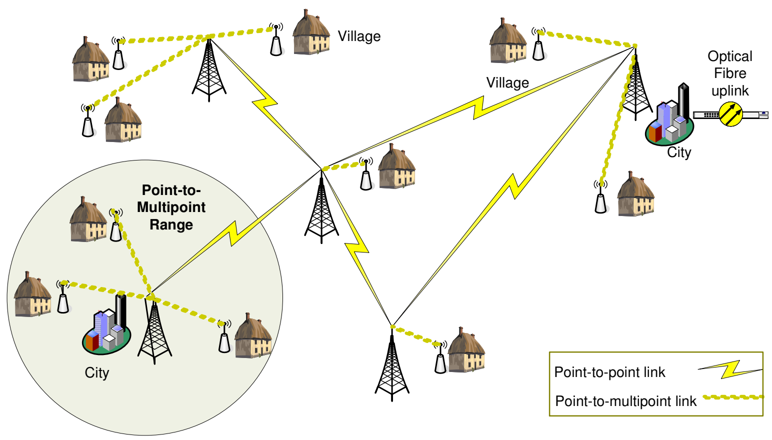

Figure 5.1: Example of a WiFi rural network. Showing a combination of long-distance point-to-point backhaul links and medium-range point-to-multipoint access links.

distance point-to-point backhaul links by using a TDMA protocol in place of WiFi's normal CSMA/CA protocol [104].

In this chapter, we explore architectural design issues pertaining to the point-to-multipoint component of WiLDNet, where each base-station provides high bandwidth access to a large number of client-stations. A few other proposals such as SRAWAN [115] and WiFiRe [106] also aim to achieve higher aggregate throughput in point-to-multipoint settings.

5.1 Background

It is imperative to design point-to-multipoint access networks that afford high throughput to a large number of client-stations, even if the external Internet connectivity to such networks is modest. The reason is that the client-stations often further share their connectivity among several users in tele-centers and schools, where individual users can rarely afford individual network connections and personal devices. Moreover, we have recently seen an increase in interactive applications such as tele-medicine [138] and remote education, which require high-quality video feeds and low-loss VoIP connectivity.

With tall communication towers being a dominant cost factor [87], serving many client-stations with few communication towers is essential. For example, in Akshaya [8], one of the largest rural wireless networks in the world, the ratio of client-stations to each base-station is between 10 and 40. Although this model decreases network deployment costs, it also decreases the throughput delivered to each client-station (each client-station in Akshaya only enjoys a typical throughput of less than 1 Mbps) and bounds the maximum number of client sites supported per tower to the number of antennas that can be mounted on a given tower.

In order to increase both capacity and tower coverage, the standard approach taken by network operators is to increase the base-station capacity by *sectorization*: the co-location of multiple base-station radios on the same tower, and the use of multiple sector antennas, one for each radio. Assuming the availability of as many as K non-overlapping wireless channels, and assuming that S sector antennas are deployed for each channel, in theory, the per-tower capacity scales with $K \times S$. This implies that for each channel, we need to operate S base-station radios, one each for all the S sectors. However, in practice, such scaling is limited due to interference (because the S sector antennas operating on the same channel can cause interference at both the client-station and the base-station) and the cost of installing additional antennas.

Thus, the key challenges that motivate our architectural framework for improving capacity and throughput in point-to-multipoint networks are performance, interference, cost, and flexibility.

- **Performance:** We saw in Chapter 3 that the link utilization and throughput of long-distance 802.11 links decreases drastically if we just use the standard WiFi MAC protocol. As the link distances in point-to-multipoint access networks can be as long as 20 km, we need to use directional antennas at the client-stations pointing towards the base-station. This rules out using the standard carrier sensing MAC protocol (CSMA) since the client-stations will be hidden from each other. In fact, WiMAX, which is usually deployed in similar scenarios uses a TDMA-based MAC protocol with the base-station as the master node. Another option is to use 802.11's PCF (Point Coordination Function) MAC that uses polling based packet delivery. Unfortunately, PCF is not implemented in any commercial WiFi products.

- **Interference:** Using several sector antennas on the same channel can lead to both client-side and base-station-side interference, given the imperfect nature of the directional signal amplification obtained by using sector antennas. On the downlink direction (i.e., from base-station to client-stations), simultaneous transmission from different sector antennas on the same channel can interfere at a client-station. Similarly, on the uplink direction (i.e., from client-stations to the base-station), client-stations in different sectors operating on the same channel can cause mutual interference at their respective antennas. Although this problem can be mitigated to some extent

Figure 5.2: Picture of a radio tower in the Akshaya deployment featuring multiple antennas for point-to-multipoint connectivity. The antennas need to be aligned on the vertical axis and aimed carefully.

through careful aiming and positioning of sector antennas, the large number of antennas required by today's designs decreases the amount of interference that is mitigated by careful antenna alignment and positioning.

- **Cost:** Deploying a large number of sector antennas at each tower is both expensive and physically challenging. Let us illustrate this point by examining Figure 5.2, which shows the top part of a real communication tower in the Akshaya networks [8], featuring 9 directional or sector antennas. Although the setup shown here seems easy and feasible to the uninformed eye, the Akshaya network managers were very unhappy with the large deployment costs and even larger maintenance costs demanded by such an installation. In order to ensure enough spatial isolation among several sector antennas, the communication tower was built much taller and heavier than it would have been required to ensure line of sight conditions, making it by far the most expensive tower in the entire network. Installing and aligning the antennas proved a challenging task as well.

- **Flexibility:** In addition, since today's architectures drive each antenna through only a single radio, they're unable to adapt to *dynamically changing traffic patterns*. The most challenging part is maintaining and adapting the installation to accommodate an increasing number of users. In Akshaya, as more client-stations were connected, the antennas were realigned several times to accommodate the changes in workload and spatial distribution of the client-stations. Recurring interference problems also demanded countless rounds of antenna shifting and realignment, making it a nightmare for maintenance technicians involved. After this experience, the tower remained the only example when the installation of that many antennas was ever attempted in the network.

5.2 Contributions

In this chapter, we build a TDMA-based MAC protocol for point-to-multipoint networks building on top of the WiLDMAC protocol, that we had earlier built for point-to-point links (Chapter 4), and explore three novel techniques to deliver high throughput in point-to-multipoint access networks.

TDMA-based point-to-multipoint MAC protocol: In the basic scheme, the base-station acts as the master node and allocates downlink and uplink slots to different client-stations. The framing is controlled by the base-station to ensure that the client-stations are synchronized to the base-station at the start of every time slot. Instead of implementing tight time synchronization between the base-station and the client-stations, we have a mechanism where each client-station records the time at which it receives a beacon from its base-station and uses that to schedule its uplink slot.

When the base-station has multiple sectorized antennas, we use our insight from point-to-point networks where we found that we can synchronize multiple radios to achieve simultaneous transmit and receive during a time slot. Thus, we can get the same throughput while running all the radios on the same channel (with synchronization) as we would get while running the radios on different channels without any synchronization. This means that in a transmit slot, we simultaneously transmit from all the radios (after disabling CCA) at a base-station and in a receive slot, we simultaneously receive at all the radios at a base-station.

In the ideal situation, if we have S base-station radios each with its own sector antenna and a throughput of M for each radio, then we can achieve a total throughput of $S \times M$ for each channel at the base-station.

Dynamic power control: We design adaptive power control mechanisms to address the problem of interference for all the base-station radios that operate on the same channel.

Our objective is to maximize the total network bandwidth, while satisfying all the per-client minimum bandwidth constraints. To achieve this, we frame our problem as an LP optimization problem, and compute an optimal combination of *a) client-station to base-station allocation*, *b) link transmission schedule* and *c) radio transmit powers*. We also quantify the importance of performing power control by comparing the maximum bandwidth achievable with or without power control.

We show that techniques performing smart transmission scheduling and power control have the potential to substantially increase the number of client-stations that can be accommodated, and also the total network throughput.

Decrease cost by combining antennas: Given the difficulty and cost of installing and aligning many sector antennas on one tower, we propose architectures aimed at *maximizing network capacity* while using *the minimum number of antennas per tower*. To this end, we leverage a set of techniques allowing multiple radios to connect to the same physical antenna, or allowing radios to dynamically switch the antennas they are connected to, adapting to the dynamic client traffic demand.

We show that static allocation using inexpensive off-the-shelf splitters is indeed feasible in practice, pointing to the need for further evaluation of these architectures.

Adaptive channel width: A key limitation to increasing the capacity of wireless mesh networks is the limited number of orthogonal channels available at the base-station. If we recall, when we have K non-overlapping wireless channels, and S sector antennas are deployed for each channel, in theory, the per-tower capacity scales with $K \times S$.

We show that when we use more number of narrower width channels instead of just using a single wider fixed wireless channel (as defined by 802.11), we can increase the number of simultaneous transmissions and thereby, increase the cumulative capacity of the whole network by almost 30-110%.

5.3 Assumptions

In this section we describe our assumptions of the point-to-multipoint setup and MAC-layer requirements. A base-station has multiple co-located radios, each connected to a sector antenna. In the ideal scenario, we have a radio for each orthogonal channel in every sector. Also, to maximize spectrum usage, we would like all the radios on a particular channel to operate simultaneously.

Fixed client-stations are situated upto 20 km away and have directional antennas pointing to the base-station. Time is divided into downlink and uplink slots for base-station and client-station transmissions respectively. Base stations can synchronize transmissions and receptions among all their radios. All client-station traffic flows via the base-station to the outside world. Each client-station has a simple interference model where the difference in received signal strengths must be greater than the isolation threshold (Th_{cap}), which we set to be 20 dB.

5.4 Adaptive Power Control

In this section, we present our first technique, namely *adaptive power control* to improve capacity of point-to-multipoint access networks. The main objective is to improve spatial reuse for single channel settings of point-to-multipoint networks. Note that we will also use other non-overlapping channels, to scale capacity by adding an extra set of radios for each extra channel.

Using multiple radios with sectorized antennas at the base-station that share single channel introduces interference, of the following types:

Rx-Tx interference at base-station: When a base-station radio, say R_a, receives a packet from a client, while another base-station radio, R_b at the same tower is sending a packet, interference can occur at the base-station radio, R_a which is trying to receive packets. Due to imperfections in the directional characteristics of the sector antennas (side lobes), R_a can easily overhear R_b's transmissions, and this almost always results in the corruption or capture of the packet being received at R_a.

Tx-Tx interference at a subscriber: When client c_i can simultaneously hear transmissions from two different base-station radios (out of which only one packet is intended for c_i), interference occurs if the transmissions have *similar signal strengths*. By similar signal strengths we mean that the signal difference is less than an *isolation threshold*. For Atheros based WiFi radios, we have experimentally determined this isolation threshold to be approximately 12 dB in Chapter 3.

Rx-Rx interference at base-station: This is symmetric to the previous case, but now the interference occurs at base-station radio R_a due to conflicting transmissions from two subscribers c_i and c_j. Again, a large enough isolation between these two transmissions is required to avoid interference.

Of these three types of interference, the first type (Rx-Tx) is difficult to handle when using sector antennas, and most existing TDMA-based solutions [104, 112] disallow the situation itself by synchronizing co-located radios such that they either all *transmit simultaneously*, or they all *receive simultaneously*. We therefore look at the remaining types

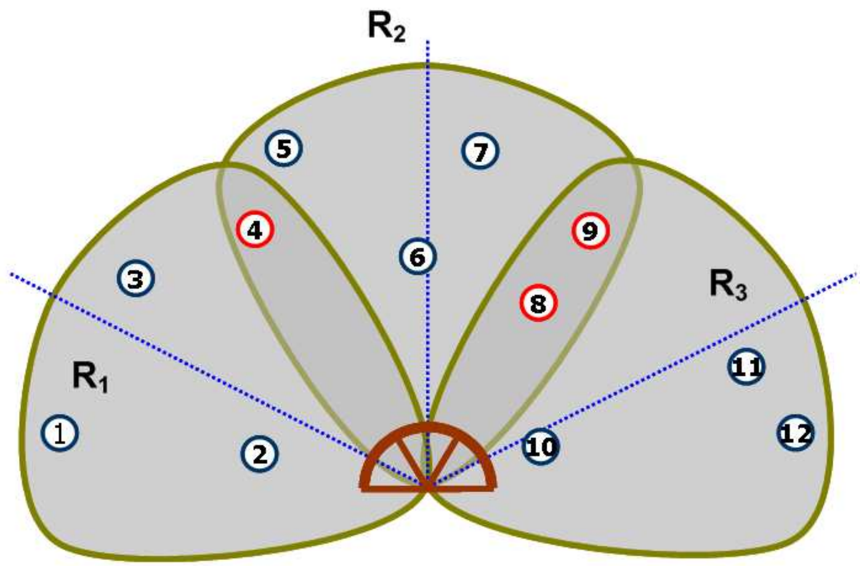

Figure 5.3: Basestation with 3 radios and 12 clients with maximum power All radios are using maximum power. Clients c_4,c_8,c_9 are in a region with potential interference.

of interference, and propose solutions to address them.

Let us begin by analyzing one of these types of interference (Tx-Tx) in an example scenario featuring three base-station radios, each connected to a sector antenna; we assume these radios are synchronized as described above. Figure 5.3 shows the polar propagation plot of the base-station radios with their antennas, in the default case (used in current deployments) when they all transmit at maximum power. The propagation region of each antenna is defined as the region where the received signal strength is higher than the receive signal threshold Th_{rx}, over which a client can successfully receive the packet. For example, clients c_1, c_2, c_3 and c_4 fall in the receive lobe corresponding to radio R_1. However, if all the radios are transmitting at the same time, a few clients like c_4, c_8 and c_9 will fall in the intersection of the antenna receive lobes of two neighboring radios. At these clients, Tx-Tx interference can occur if the signal strength difference between the primary radio and the

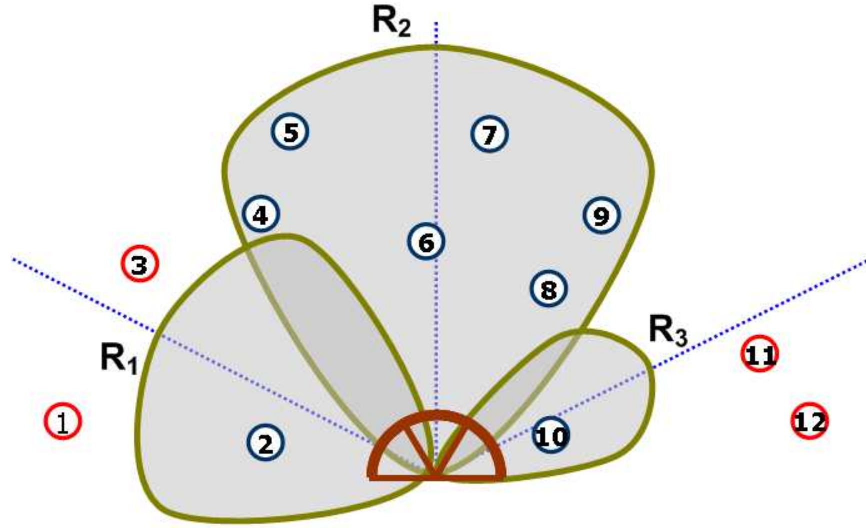

Figure 5.4: Basestation with 3 radios and 12 clients with power control. While radio R_2 is using maximum power, radios R_2 and R_3 are using power control to reduce interference. Clients c_4,c_8,c_9 are now in interference-safe regions.

interfering radio is lower than the *isolation threshold* (Th_{cap}). The set of clients that can receive a clean signal from a radio is referred to as the *safe set* of the radio. Thus, when radio R_1 is transmitting, it's safe set consists of the clients c_1, c_2 and c_3.

One way to eliminate interference and ensure that the minimum isolation threshold is achieved is to change the transmit power at the base-stations. For example, if we want to eliminate interference at clients c_8 and c_9, we can reduce the transmit power of the radio R_3, such that clients c_8 and c_9 now fall only into the receive region of radio R_2. On the other hand, now radio R_3 cannot reach clients c_{11} and c_{12} any longer (Figure 5.4).

Thus the problem of minimizing Tx-Tx interference (or maximizing the number of clients that receive), can be framed as finding the optimal radio power levels, and the optimal allocation of clients to base-station radios, such that it enables the base-station to serve the most clients. The power levels and client allocation could be fixed, or they could change during every time slot, with the latter approach providing additional gains. The solution to this problem can then be used to regulate transmissions during the time slots in which the base-station radios transmit to clients.

The problem of minimizing Rx-Rx interference is very similar, and entails finding optimal transmit power levels for the client radios, together with optimal client allocation, that would enable the most clients to transmit successfully to the base-station. The solution to this optimization problem would be used during the slots when the base-station receives packets from clients.

5.4.1 Power Control Strategies

The different possible strategies for power control are described here as we increase the degree of freedoms in controlling the base-station radios.

Always maximum power (F-MAX): The base-station radios always transmit at the fixed and maximum power setting in all slots. This could be very inefficient as there could be many clients that can never be in the safe set of any of the radios in any slot.

Greedy scheduling with maximum power (G-MAX): The base-stations use a greedy algorithm to coordinate with each other and select a set of M clients for each slot, where M is the number of base-station radios. For each slot in a period, the strategy first tries to allocate a set of clients whose bandwidth bounds have not been satisfied. Other radios that are not allocated any client are switched OFF.

Optimal maximum power allocation (O-MAX): All the base-station radios can only be set to the maximum power setting or are switched OFF during a slot. We find the best possible allocation of clients to radios for each slot in a period so that we satisfy the bandwidth bounds for each client.

Optimal fixed power allocation (O-FIXED): All the base-station radios and clients have a fixed power setting that is determined to be the optimal power setting for the best possible allocation of clients to radios.

Optimal dynamic power allocation (O-DYN): All the base-station radios can change their transmit power in every slot. These power settings and the allocation of clients to radios in each slot is optimally decided every slot.

5.4.2 Optimal LP Formulation

We define RS_j as the set of clients at which the received power from radio j is above the receive power threshold (Th_{rx}). IS_j is the set of clients where the difference between the power from radio j and the other neighboring radios is less than the capture threshold (Th_{cap}), i.e. it is the set of clients which do not receive successfully from radio j due to interference. Then the *safe set* is $RS_j - IS_j$.

The aim of the optimal formulation is to select for each radio, R_j, one client c_i from its safe set at every slot such that we satisfy the minimum bandwidth requirement for each client over the whole period. This problem can be formulated as a mixed integer LP problem.

Definitions:

Set of clients: $C = (c_1, c_2 ... c_N)$

Set of radios: $R = (R_1, R_2 ... R_M)$

Set of slots: $S = (s_1, s_2 ... s_P)$

Path loss between client c_i and radio R_j: L_{ij}

Maximum transmit power for a radio or client: P_{max}

Sum of antenna gain between client c_i and radio R_j: G

Receive power threshold: Th_{rx}

Capture threshold: Th_{cap}

Minimum BW limit for a client: BW_{min}

Large constant:P_{inf}

Variables:

X_{ij}^s, binary variable that is 1 iff c_i is in the *safe set* of radio R_j in the sth downlink slot.

PD_j^s, transmit gain of radio R_j in the sth slot.

PU_i^s, transmit gain of client c_i in the sth slot.

Objective Functions: Maximize the total bandwidth over all the slots such that all the clients satisfy a minimum bandwidth requirement.

$$\text{Maximize} \sum_{i=1}^{N} \sum_{s=1}^{P} \sum_{j=1}^{M} X_{ij}^s \tag{5.1}$$

Subject to:

$$PD_j^s + G - L_{ij} > Th_{rx} + P_{inf}(X_{ij}^s - 1)$$
$$\forall s \in S, \forall c_i \in C, \forall R_j \in R \tag{5.2}$$

$$(PD_j^s - L_{ij}) - (PD_k^s - L_{ik}) >$$
$$Th_{cap} + P_{inf}(X_{ij}^s - 1)$$
$$\forall s \in S, \forall c_i \in C, \forall R_j, R_k \in R \text{ and } j \neq k \tag{5.3}$$

$$\sum_{j=1}^{M} X_{ij}^s \leq 1,\ \forall c_i \in C, \forall s \in S \tag{5.4}$$
$$\sum_{i=1}^{N} X_{ij}^s \leq 1,\ \forall R_j \in R, \forall s \in S \tag{5.5}$$
$$0 \leq PD_j^s \leq P_{max}\ \forall R_j \in R, \forall s \in S \tag{5.6}$$
$$\sum_{s=1}^{P} \sum_{j=1}^{M} X_{ij}^s \geq BW_{min}\ \forall c_i \in C \tag{5.7}$$

Equation 5.2 enforces the condition that the client selected for a slot has a receive power above the receive threshold. Equation 5.3 enforces the condition that every selected client can capture the transmission from the corresponding radio without interference. Equation 5.4 ensures that every client is included in the safe set of exactly one radio. Equation 5.5 ensures that each radio is sending to only one client in each slot. Equation 5.6 enforces the maximum power constraint for each radio. Finally, Equation 5.7 ensures that the bandwidth of each client is at least more than the minimum bound.

If there is no feasible solution that satisfies all the clients, we fold back to solving a different LP optimization, that has as an objective the maximization of the total number of clients for which the minimum bandwidth constraint is satisfied. For this we introduce an additional binary variable Z_i, that is 1 *iff* the total bandwidth of client c_i over all the slots is greater than BW_{min}. Then, our new optimization function becomes:

$$\text{Maximize} \sum_{i=1}^{N} Z_i \tag{5.8}$$

and condition 5.7 (which proved to be unsatisfiable) is now replaced with the following:

$$\sum_{s=1}^{P} \sum_{j=1}^{M} X_{ij}^{s} \geq BW_{min} Z_i \ \forall c_i \in C \tag{5.9}$$

Equation 5.9 ensures that the binary variable Z_i is set to one iff client c_i satisifies its bandwidth requirement over all the slots.

To constrain the power output of the radios according to various power control strategies discussed in section 5.4.1, we add further constraints:

Always Maximum power (F-MAX): Since each radio is *always* set to the maximum power, the additional constraint (Eq. 5.10) is :

$$PD_j^s \quad = \quad P_{max} \ \forall s \in S, \forall R_j \in R \tag{5.10}$$

Optimal maximum power (O-MAX): Here each radio can only be set to the maximum power setting, but we also have the choice of switching off a radio in a slot. We need an additional constraint (Eq. 5.11) to enforce this. This constraint can be transformed into linear constraints by using additional binary variables.

$$
\begin{aligned}
PD_j^s &= P_{max} \text{ iff } R_j \text{ is ON} \\
&= 0 \text{ iff } R_j \text{ is OFF} \\
& \quad \forall s \in S, \forall R_j \in R
\end{aligned}
\tag{5.11}
$$

Optimal fixed power (O-FIXED): Now each radio's power output is fixed for all the slots. However, we still have the choice of switching off a radio in a slot. We add another constraint (Eq. 5.12) that can be transformed into linear constraints by using additional binary variables.

$$
\begin{aligned}
PD_j^s &= PD_j \text{ iff } R_j \text{ is ON} \\
&= 0 \text{ iff } R_j \text{ is OFF} \\
& \quad \forall s \in S, \forall R_j \in R
\end{aligned}
\tag{5.12}
$$

Optimal dynamic power (O-DYN): Each radio can change its power in every slot. Equation 5.6 suffices and no additional constraints are required.

Receive power constraints: In practice, we can also use similar power control strategies for uplink slots from the clients to the radios. In this case, we have to ensure that the signal from the selected client is not drowned out by other clients at the radio. We can add additional constraints to the LP formulation by including an additional set of uplink slots (same as the number of downlink slots). We will need an analogous set of variables (Y_{ij}^s) to specify whether a client uses an uplink slot to a radio. The new capture constraint is shown in equations 5.13 and 5.14. We have not included the full set of equations due to space constraints.

$$PU_i^s + G - L_{ij} > Th_{rx} + P_{inf}(Y_{ij}^s - 1)$$

$$\forall s \in S, \forall c_i \in C, \forall R_j \in R \quad (5.13)$$

$$(PU_i^s - L_{ij}) - (PU_k^s - L_{kj}) >$$

$$Th_{cap} + P_{inf}(Y_{ij}^s - 1)$$

$$\forall s \in S, \forall R_j \in R, \forall c_i, c_k \in C \text{ and } i \neq k \quad (5.14)$$

Steerable Antennas: With electronically steerable antennas, we have more degrees of freedom for various antenna configuration parameters. To simplify analysis, we discretize the steerable antenna configuration: there are D directions (0 to 360 degrees) and L types of lobe patterns (sizes). Now we have to define the path loss for each configuration from each radio to each client. Let the path loss from client c_i from radio R_j in the antenna configuration with direction D_p and lobe size L_q be L_{ij}^{pq}. We also define ${}^sT_j^{pq}$, binary variable that is 1 iff R_j is in configuration with direction D_p and lobe size L_q, in the sth downlink slot. Then the loss on the path can be expressed as an additional linear constraint as equation 5.15. Finally we have to ensure that each antenna has exactly one configuration in every slot (equation 5.16).

$$L_{ij}^s = \sum_{p=1}^{D}\sum_{q=1}^{L} {}^sT_j^{pq} L_{ij}^{pq} \quad (5.15)$$

$$\forall c_i \in C, R_j \in R, D_p \in D, L_q \in L, s \in S$$

$$\sum_{p=1}^{D}\sum_{q=1}^{L} {}^sT_j^{pq} = 1 \quad \forall D_p \in D, L_q \in L, s \in S \quad (5.16)$$

5.4.3 Evaluation

In this section, we present some preliminary evaluation of the various point-to-multipoint strategies presented earlier for selected topologies. The objective is to quantify the gap between the naive strategies and the smarter optimal strategies which will enable us to look for practical implementations of the optimal strategies.

To that end, we first compare the different strategies in terms of the maximum possible bandwidth achievable while satisfying the minimum bandwidth constraints of all the clients in the network. In case if it is not possible to satisfy the constraints of all the clients, we try to maximize the total number of clients whose constraint can be satisfied.

We evaluate the following scheduling strategies: 1) Always maximum power (F-MAX), 2) Greedy scheduling with maximum power (G-MAX), 3) Optimal scheduling with maximum power (O-MAX), 4) Optimal scheduling with fixed power (O-FIXED) and 5) Optimal scheduling with dynamic power (O-DYN).

Methodology:

Topology generation: We randomly place clients at a range of $1-20$ km from the base-station within an angular coverage of 180 degrees.

Radio antenna pattern: We place three radios at the base-station. The radios are connected to fixed antennas that are placed at an equidistant angular separation. The maximum power of each radio is set to $P_{max} = 23$ dBm (200 mW). The antenna gains of the base-station antennas and the client-side antennas are assumed to be 18 dBi at 2.4GHz. The capture threshold is set to $Th_{cap} = 20$ dB which is conservatively set with respect to measurements conducted by us previously 3. The antenna lobe size is selected such that roughly one-third

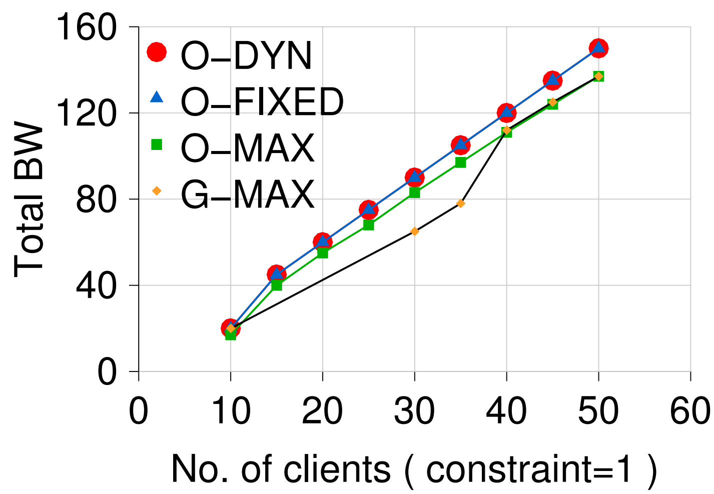

Figure 5.5: Maximum bandwidth with increasing clients with constraint = 1 for 3 base-station radios. The naive maximum power strategy (F-MAX) did not find any feasible solution.

of the clients are reachable by more than one radio.

Mixed Integer Linear Program (MILP) solving: We formulate all the point-to-multipoint scheduling strategies except the greedy scheduling as MILP problems. We solve the optimization for a fixed number of slots (referred to as a period) which gives us a periodic schedule. We set the number of slots in a period to be equal to the number of clients. This is to ensure that it is always possible to find a schedule where each client is served at least once during the entire period. This also bounds the maximum delay of each client to the number of slots in the period.

We then use the CPLEX LP solver [63] to optimize the mixed integer LP formulations. We use a maximum time limit of 600 seconds to find optimal solutions; this was generally sufficient in most cases.

Results:

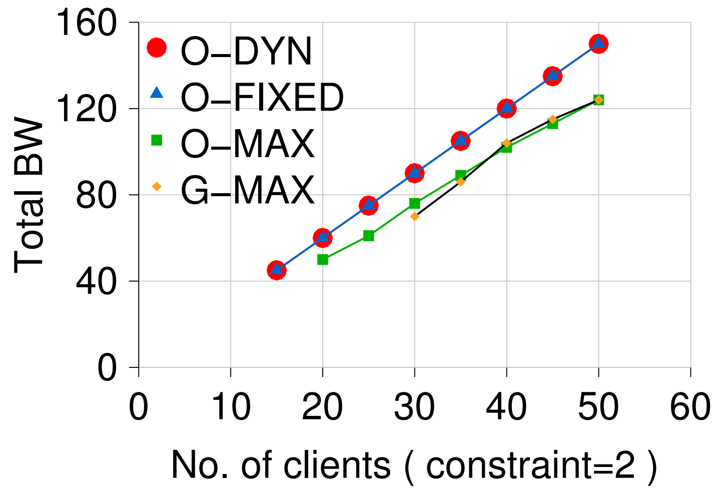

Figure 5.6: Maximum bandwidth with increasing clients with constraint = 2 for 3 base-station radios. The naive maximum power strategy (F-MAX) did not find any feasible solution.

Maximum overall bandwidth: We first compare the strategies in terms of the total bandwidth while satisfying all client constraints. Figures 5.5 and 5.6 show the maximum bandwidth achieved by the different algorithms for bandwidth constraint of 1 slot per period (50 slots) and 2 slots per period (50 slots) respectively as we increase the number of clients in the network. None of the strategies were able to achieve perfect scheduling of 3 slots per period for all the clients (50 slots and 3 radios). The F-MAX strategy was not able to find any feasible solution that satisfied all clients even for the smallest constraint. The greedy strategy (G-MAX), was only able to find feasible solutions for larger topologies (clients > 30) where there is more flexibility to choose slots. The optimal scheduling strategies (O-MAX, O-FIXED, O-DYN) were able to satisfy all clients but there was a gap (upto 17%) between the strategy that can use only the maximum power (O-MAX) and the others. We can see that this gap is higher as we increase the minimum constraint to 2 slots. It can

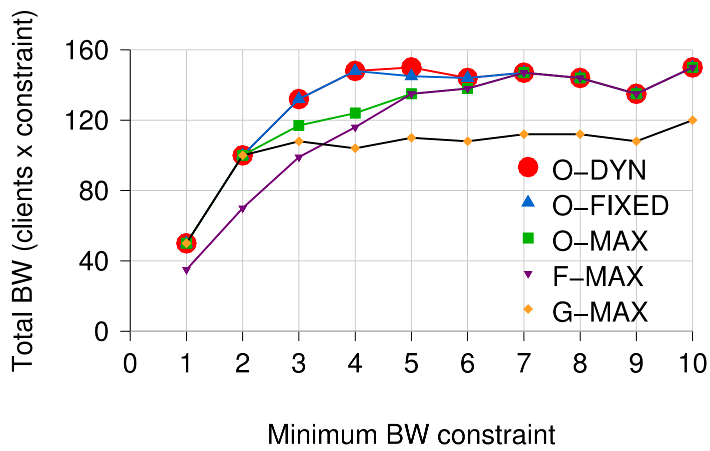

Figure 5.7: Maximum number of clients with increasing the bandwidth demand per client. Y-axis is the total bandwidth utilized in slots i.e. product of the no. of clients and the minimum constraint. For 50 clients and 3 radios.

also be seen that both the Optimal Fixed (O-FIXED) and the Optimal Dynamic (O-DYN) strategies achieve the maximum possible bandwidth.

Maximize clients with bandwidth constraints: When it is not possible to satisfy the constraints for all the clients, we want to maximize of the number of clients whose bandwidth constraint can be actually satisfied. Figure 5.7 shows the number of clients that satisfy the minimum bandwidth constraint as we increase the constraint for a topology with a fixed number of clients (50). On the Y-axis, we plot total bandwidth which is being utilized, i.e., the product of the number of clients and the bandwidth constraint. We can see that there is a substantial gap between the optimal strategies and the fixed power strategy (about 50% at constraint of 3). This gap reduces as we make the constraint tighter. All the strategies can reach upto the maximum capacity of the network (150 for 50 slots and 3 radios) as it becomes easier to just select a small subset of clients that can be scheduled multiple times easily. It is interesting to note that even the strategy with always maximum power

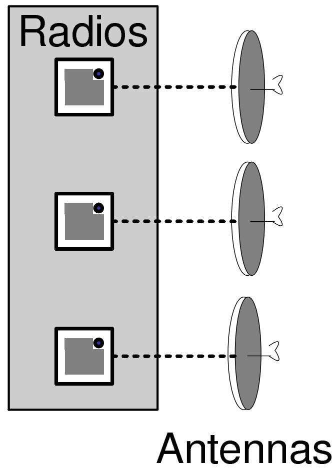

Figure 5.8: Basic design where each radio connects to only one antenna. We achieve a throughput of $K \times S$, where K is number of channels and S is number of sectors, only by using $K \times S$ antennas.

(F-MAX) also manages to reach a subset of the clients. The greedy strategy (G-MAX) does much worse than all the other solutions at higher constraints. We can also see that both the Optimal Fixed (O-FIXED) and the Optimal Dynamic (O-DYN) strategies achieve the maximum possible bandwidth all throughout. To summarize, the part of the graph which is relevant in practice is where the minimum bandwidth constraint is lower and here we see that the optimal power strategies outperform the greedy strategies.

5.5 Antenna Combination

In this section, we propose techniques to design point-to-multipoint networks with a large number of radios that can provide high throughput, while using a smaller number of physical antennas to send and receive signals from these radios. This desire to design configurations that provide high throughput but require fewer antennas to be installed and maintained is driven by a practical constraint on the number of antennas that can be

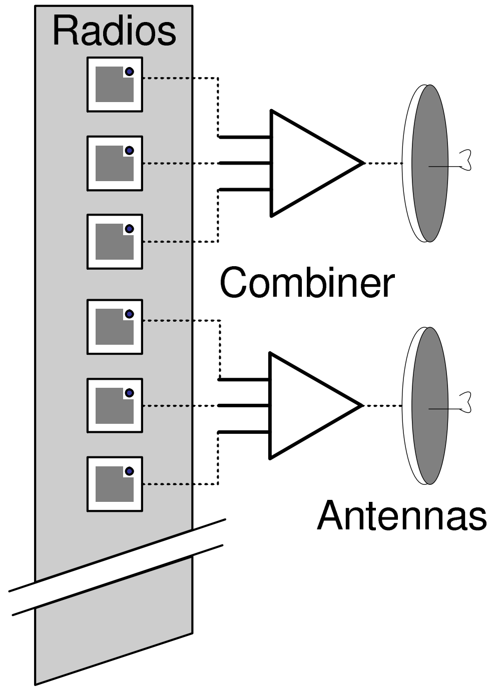

Figure 5.9: Each antenna is connected to $K = 3$ radios using a splitter/combiner. We achieve a throughput of $K \times S$ using only S antennas but $K \times S$ radios.

installed at the base-station in a point-to-multipoint setting.

5.5.1 Antenna Architecture Options

Consider the prevailing status quo. In this architecture, presented in Figure 5.8, we allocate one radio per antenna. This architecture is simple, but has the drawbacks mentioned in Section 5.1: *a)* leads to a large number of antennas – at most $K \times S$ (number of channels times number of sectors per channel); *b)* it reduces the flexibility of capacity allocation. As an example, if all active clients happen to be in the range of a single sector antenna, only one radio can be used to serve all of these clients, while the other radios remain idle. Conversely, if an antenna loses alignment or experiences connectivity failures, the radio connected to it and its corresponding throughput is effectively lost.

In order to overcome these drawbacks, we examine the following alternatives:

Combining several radios for each antenna: We argue that it is possible to design a multiplexing system that allocates multiple radios, operating on different channels, to a single sectorized antenna, by using a single n-port splitter/combiner device (Figure 5.9). This device acts as a combiner when one or more of the radios it is multiplexing are transmitting, and as a splitter when its radios are receiving. The main constraint is that, when radios are transmitting, the splitter/combiner must provide enough inter-port isolation to shield one radio from receiving the high-energy transmission of a neighboring radio, even if the two radios are on different non-overlapping (i.e., orthogonal) channels. We will evaluate this architecture later in more detail in Section 5.5.2.

The main advantage of this approach is the fact that it only uses S sector antennas to achieves the same capacity as the default solution with $K \times S$ antennas. Since the number of orthogonal channels K is three (channels 1, 6, and 11) in the 2.4 GHz band, and up to 12 in the 5 GHz band, we can obtain significant savings in the number of required antennas. In turn, reducing the number of antennas decreases the deployment costs and simplifies installation and maintenance.

The disadvantage of this architecture is the fact that it is not load aware, with radios being statically tied to sectors, regardless of whether there is enough traffic demand for each sector. Another disadvantage is related to the loss introduced by the splitter on the receive path, when the incoming signal from the antenna is split across all radios, thereby reducing the signal received at the relevant radio by $\approx 10\log(n)$ dB, where n is the number

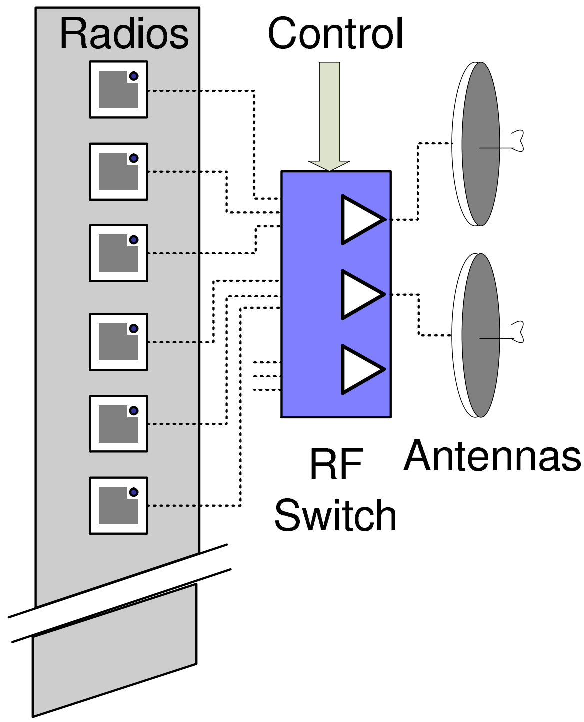

Figure 5.10: All the radios are connected to the antennas through a RF switch that is controlled by the host. The switch can route a radio to any of the antennas. We achieve a throughput of $K \times S$ using only S antennas and radios.

of multiplexed radios.

Multi-port switch connecting radios to antennas: To address the flexibility issue, we propose a second alternative (described in Figure 5.10), that uses dynamic switching in order to connect multiple sectorized antennas and implement a two-stage switching. In Stage 1, each radio uses a RF-multiplexer to switch its power to a particular antenna port. In Stage 2, the output of all the RF multiplexers is combined into an antenna.

This architecture is more flexible, as radios can be steered to antennas based on dynamic demand as well as in order to tolerate remote link failures. Although electronically steerable antenna systems, such as those proposed in [91], also have the potential to increase throughput by taking client traffic demands into account, electronically steerable antenna systems remain expensive (ranging several thousands of dollars upward), which decreases their attractiveness for low-cost wireless networks.

However, this architecture also has two significant disadvantages. First, there is increased cost and complexity due to the use of RF multiplexers. Second, each additional RF component introduces losses ranging from 0.7 dB up to 3 dB. One way to combat these losses is to use power amplifiers and low-noise amplifiers (LNAs) on the transmit and receive side respectively of the antennas. In order to understand how severe these isolation and signal loss problems are in practice, we evaluate the static multiplexing architecture in detail in the next section.

5.5.2 Evaluation

We evaluate the static multiplexing design of combining multiple radios to a single antenna using off the shelf splitter/combiners (henceforth referred to simply as splitters).

In order to recreate outdoor long-distance environments, we use the experimental setup described in Figure 5.11. We use two RF-isolated enclosures, one for the base-station radios (enclosure A), and another for the client-side radios (enclosure B). Each of these enclosures contains three 802.11b/g mini-PCI Atheros-based radios (Ubiquiti SR-2) with a maximum TX power of 25 dBm, driven by WRAP wireless routers. The radios in the first box are combined using a 3-way commercial splitter/combiner, offering a 20 dB port-to-port isolation. As we will see shortly, this isolation between the radios is not enough: even if configured on non-overlapping channels, the transmit energy of one radio is too large to be fed directly to the other radio, with only 20 dB attenuation. We therefore simulate a splitter

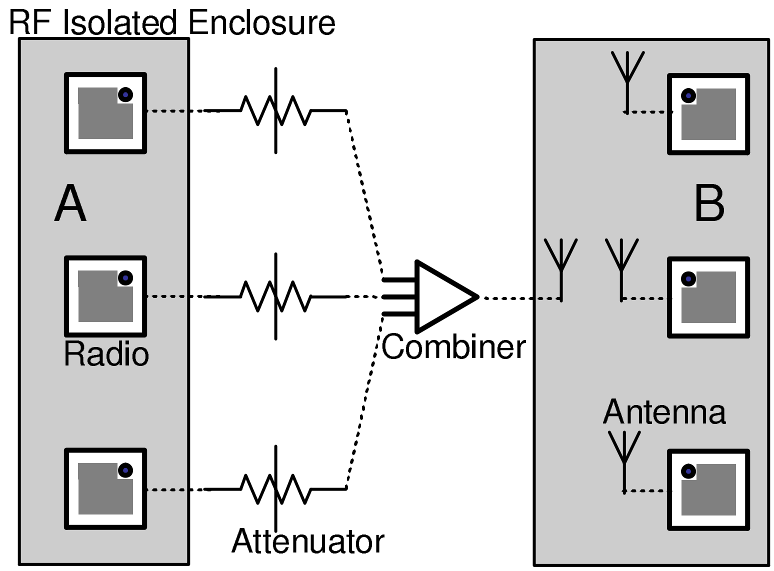

Figure 5.11: Experimental setup for combining multiple radios using the same base-station antenna.

featuring more port-to-port isolation by adding additional attenuators at every base-station radio.

In the other enclosure B, corresponding to client-side radios, each radio is equipped with an independent “rubber duck” antenna.

Using this setup, and with 20 dB attenuators at each base-station radio, we first test that combining two or three radios on a single antenna is feasible. Therefore, we measure the UDP throughput achievable by combining the transmissions between the three pairs of radios under three scenarios: *a) simultaneous TX* from the base-station, *b) simultaneous RX* at the base-station, and *c) mix of TX and RX*. We investigate several scenarios with either two or three pairs of radios used simultaneously on non-overlapping channels.

The results of this experiment are presented in table 5.5.2, and show that the sum of the throughput achieved by combining N radios to one antenna is indeed roughly equal

No. of radios	Channels	All radios TX	All radios RX	Mix TX/RX
2	1,6	14.51	14.89	14.77
2	6,11	14.74	13.98	14.16
2	1,11	13.90	13.79	13.70
3	1,6,11	21.14	20.34	

Table 5.1: UDP throughput (sum on all radios) achieved for various scenarios when we combine multiple radios to a single antenna.

to the throughput achievable by using N independent antennas, one for each radio (the individual UDP throughput is ≈ 7.5 Mbps). This total throughput is sustained when using either 2 or 3 radios connected to one antenna.

This experiment shows that combining multiple radios is feasible, but our experiment required additional attenuation at each radio in order to insure the necessary inter-radio isolation. In the next experiment, we vary this inter-radio isolation, and measure its effect on the total achieved throughput. We begin with testing the scenario where two combined radios transmit simultaneously, and we measure their combined UDP throughput.

Figure 5.12 shows the aggregate throughput achieved by the two radios when they transmit simultaneously, as we vary the isolation between the two radios by changing the attenuation of each radio from 10 dB to 25 dB. We are using 10 dB attenuators at both the base-station radios. The x-axis is the total port-port isolation between the two simultaneous transmitters, and the y-axis is the aggregate throughput achieved by both the radios.

To get maximum isolation, we set the transmitting radios to minimum TX power (at 10 dBm). This corresponds to a total isolation of 55 dB between the two transmitters (the rightmost point on the x-axis) and our system achieves the maximum possible system

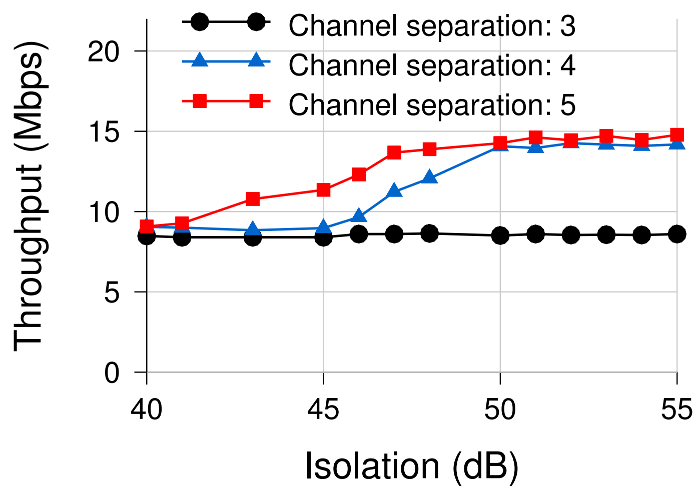

Figure 5.12: Aggregate throughput under simultaneous transmission for different isolation levels.

throughput of around 15 Mbps for a channel separation of at least 4. As we decrease the effective attenuation inserted between each transmitter and the splitter/combiner by increasing the transmit power of each radio, the aggregate throughput begins to suffer. For a total isolation of 40 dB (leftmost point on the x-axis), which is achievable by commercial splitter/combiners, we find that the achievable aggregate throughput drops by a factor of 2/3, to around 8–10 Mbps.

We attribute this decreased throughput to transmission back-offs caused by CCA (Clear Channel Assessment) as well as lost acknowledgments due to lower isolation, even if the two transmitters are separated by several channels. We have two potential ways of dealing with this problem. One relies on using lower-powered transmitters, thereby providing good isolation (i.e., 50 dB or more), and using a power amplifier closer to the antenna in order to boost the signals to their normal power levels. The other option is to set transmission parameters to disable CCA and ACKs and extend the current TDMA protocols to work even across channels. We are currently exploring both possibilities.

5.6 Channel width adaptation

In this section, we present our third technique to increase capacity of point-to-multipoint network. We have seen earlier that in a point-to-multipoint network, when we have K non-overlapping wireless channels and S sector antennas are deployed for each channel at the base-station, the per-tower capacity ideally scales with $K \times S$.

Although dynamic transmit power control (Section 5.4), can optimize the simultaneous operation of a set of sector antennas operating on the same channel, we can also increase capacity by increasing the number of orthogonal channels.

We will first show that the cumulative throughput of multiple concurrent transmissions on narrower channels is higher than the throughput of a single wider fixed width channel. This basic insight is that because we use multiple narrower channels, the total transmitted and received powers are increased, while the noise on the each channel is reduced.

5.6.1 Optimal capacity with variable width channels

We analyze the throughput improvement produced by encouraging multiple concurrent transmissions using orthogonal variable-width channels compared to TDMA schemes such as CSMA and Time-Based Fairness (TBF) that use fixed-width channels.

We consider a single cell with n clients and an AP. The AP has a single radio and antenna. Assume that the transmissions between the clients and the AP are in the uplink,

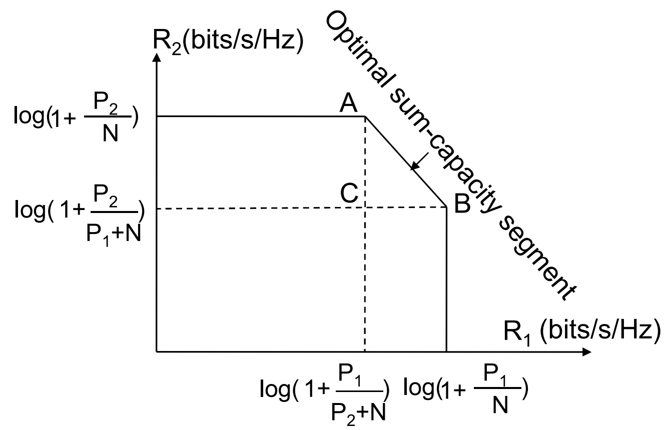

Figure 5.13: Achievable throughputs and the optimal capacity pentagon.

and that there is demand on all n links. Consider two backlogged transmitters 1 and 2 whose signals are received with powers P_1 and P_2. The receiver noise power is N per Hz. If transmitter 1 alone is active, the capacity C_1 of 1, assuming a Gaussian channel, is given by the Shannon-Hartley theorem: $C_1 = \log_2(1 + \frac{P_1}{N})$ bits/s/Hz [44, 140]. Transmitter 1 can achieve any throughput rate R_1 that is less than C_1 [44, 140].

The line segment A–B with slope -1 represents the optimal sum-capacity and is given by $R_1 + R_2 = \log_2(1 + \frac{P_1+P_2}{N})$ (shown in Figure 5.13). The reason is that, no matter how the two users code their transmissions, independently or cooperatively, it is not possible for them to exceed the capacity limit that occurs when there is a single user with total received power $P_1 + P_2$.

If we use variable width channels for 1 and 2 such that the total width is equal to the spectrum available to the receiver, we achieve non-interfering throughput rates for 1 and 2 that are given by:

$$R_1 < \alpha \log_2(1 + \frac{P_1}{\alpha N}) \text{ bits/s/Hz,}$$
$$R_2 < (1-\alpha) \log_2(1 + \frac{P_2}{(1-\alpha)N}) \text{ bits/s/Hz.} \quad (5.17)$$

where α is the fraction of the spectrum allocated to 1 ($0 \leq \alpha \leq 1$). The noise term for R_1 in Equation 5.17 is reduced by a factor α because the signal is now confined to a narrower band, while noise still occupies the entire band with power N per Hz.

We can also show that we achieve the optimal throughput when transmitters are assigned channel widths proportional to their received power at the AP [47]. It can also be shown that no TDMA scheme, such as CSMA or TBF, is optimal (i.e., its throughput does not lie on the A–B segment) because TDMA only keeps one transmitter active at a time, thereby reducing the total transmitted and received powers.

5.6.2 Evaluation

We present a preliminary evaluation of variable width channels by quantifying the potential improvements using experiments on our local Berkeley testbed. The objective of the evaluation is to verify whether we can increase total capacity by using simultaneously a larger number of narrower width but orthogonal channels instead of a single wider channel.

Testbed setup: We ran our experiments on our campus testbed, which consists of 6 wireless nodes and 10 links, 8 of which ranged from 1 km to 4 km, and 2 of which are co-located between different radios at P (Figure 5.14). Subsets of these links interfere with one another at either end-point, and each link interferes with at least one other link. The node at P has three wireless radios, the one at B has two radios and all the other nodes

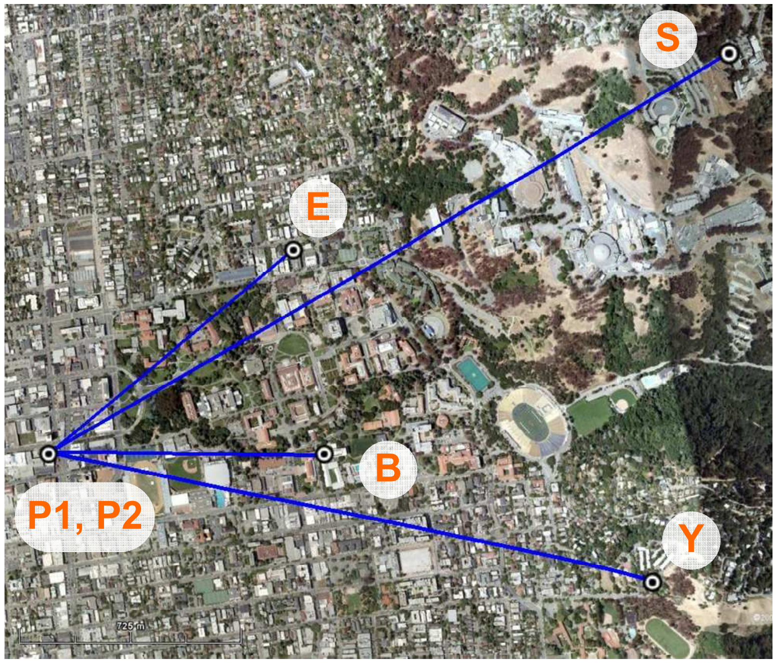

Figure 5.14: Point-to-multipoint topology configured on the Berkeley outdoor testbed. The nodes are the PowerBar building (P1,P2), Etcheverry hall (E), Space Sciences Lab (SSL), Barrows hall (B) and Yahel's house (Y).

(S, B, E and Y) have one radio each. The nodes have directional antennas of 25 dBi gain. However, because of the relatively short distances involved, we were able to configure the links into various topologies such as point-to-point and point-to-multipoint by assigning the right transmit powers to the links. We selected a fixed bit-rate for each radio based on the maximum sustainable throughput (i.e., without getting disconnected after a while) across all its links.

Channel width adaptation: We used the Ubiquiti radio driver that allows variable channel widths (5, 10 and 20 MHz). The channel width can be changed by setting appropriate register values that change the frequency of the reference clock that drives the PLL on the radio. Slowing or increasing the clock rate affects 802.11 timing parameters. For example, a

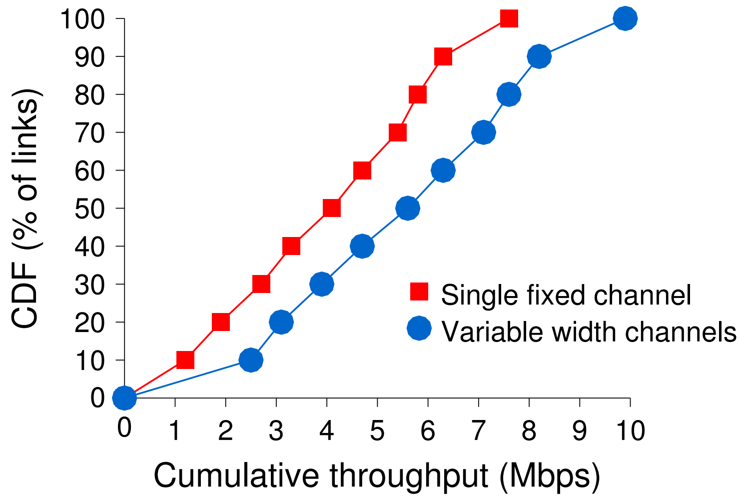

Figure 5.15: Comparing variable width channels with fixed channel for point-to-multipoint scenario

4 μs OFDM symbol in 20 MHz channel width gives symbols of length 2 μs in 40 MHz, and 16 μs in 5 MHz. While the normal 20 MHz-wide channel supports a maximum bit-rate of 54 Mbps according to the 802.11a standard, the half-width (i.e., 10 MHz) channel supports up to 27 Mbps, while the quarter-width (i.e., 5 MHz) channel supports up to 13.5 Mbps.

For our evaluation, we assigned variable-width channels to links within a single 20 MHz channel. In addition to selecting the channel width, we must also select the channel positions for 5 and 10 MHz channels within the 20 MHz channel. For example, assigning 5 MHz channels at 5.185 GHz and 5.195 GHz for two interfering links might be better in practice than assigning channels at 5.185 GHz and 5.19 GHz, because the former provides more channel separation even if neither provides perfect orthogonality. So, for every link, we consider 4 choices for placing 5 MHz channels and 2 choices for placing 10 MHz channels, in addition to retaining the 20 MHz channel option. Thus, we have seven channel choices for each link.

Results: We measured the throughput of one-way UDP flows across ten links of the testbed while they were configured in point-to-multipoint topologies, for both the scenarios - a) single fixed 20 MHz channel and b) narrower channels within the 20 MHz channel. Figure 5.15 shows the comparison between the bi-directionally averaged percentile throughputs of the links as a CDF. We can see that the using narrower variable width channels gives us improvements ranging from 30%-110%.

5.7 Implementation

In this section, we describe the details of our prototype implementation for evaluating the techniques presented in this chapter. In order to implement a time division based channel access mechanism for point-to-multipoint networks with the base-station as a master node, we build up on the WiLDMAC implementation [104] that used the Click modular router [78] framework.

Our basic platform is based on wireless nodes are 266 MHz x86 Geode single board computers running Linux 2.4.26, and is the same as the platform for our point-to-point WiLDMAC implementation. We use high-power Atheros 802.11a Ubiquiti XR5 radios (600 mW) that work in the 5 GHz spectrum.

Similar to WiLDMAC, we start from the standard Atheros MadWiFi driver, and disable link-layer association, link layer retransmissions and automatic ACKs, and instead use the *adhoc-demo* mode of operation. By using Click, we replace the stock 802.11 stop-and-wait protocol with a sliding-window based flow-control approach in which the receiver transmits a *bulk acknowledgment* (bulk ACK) for a whole window of packets. Click allows us to create virtual network interfaces, through which we can intercept packets, and perform packet re-ordering them before sending them on the real wireless interface and vice versa. To synchronize slots between the base-station and the client-stations, we use the implicit synchronization mechanism from WiLDMAC. The first packet of a slot from the base-station is timestamped, and includes the slot schedule for the client-stations. It is then used by the client-station to adjust its receive slot and schedule the subsequent transmit slot according to the received slot schedule.

To enable dynamic configuration of slots at each node, we also export two interfaces. The first interface allows us to add client-stations for each radio interface. Each client-station is specified by its IP and MAC addresses. The second interface allows us to add the dynamic slot schedule for each radio interface. The slot schedule consists of a list of slots where each slot is specified by the client-station allocated for the slot and the type of slot (receive or transmit).

5.8 Related Work

Together with the increasing use of 802.11 for long- and medium-distance outdoor networks in developing countries, point-to-multipoint 802.11 networks have also become increasingly popular.

Point-to-multipoint MAC implementation: The SRAWAN project in IIT, Kanpur [115] has implemented a TDMA-based point-to-multipoint MAC protocol on Atheros WiFi chipsets. It uses beacons from the base-station to synchronize the clients and a mix of round-robin and WFQ scheduling for quality of service. However the design focuses on single radio base-stations and does not explore the use for multiple radios to increase capacity.

The proposed WiFiRe standard from IISc, Bangalore [106] also describes a TDMA-based mechanism for implementing point-to-multipoint MAC on WiFi radios. WiFiRe also seeks to increase spatial usage by synchronizing TX and RX from multiple radios but does not address the issue of optimal allocation of clients using steerable antennas or transmit power control.

The WiMAX [2] standard also proposes a TDMA-based MAC for supporting multiple clients at distances going up to 70 km, smart antennas and other enhancements. However WiMAX does not support any sort of synchronization between different radios at the same base-station to increase throughput.

Steerable antennas [91] have also been proposed to increase network capacity, but such approaches are still in their infancy.

Channel width adaptation: Current networks use either interference suppression on a packet-by-packet basis using MAC protocols [48, 93, 104, 105, 139, 146], or cope with interference using interference cancellation and related techniques such as interference subtraction, interference alignment and ZigZag decoding [24, 44, 45, 50, 133, 140, 154]. We have introduced the idea of variable-width channels as an alternative, in which multiple transmitters operate nd ensure their orthogonality to avoid interference further.

While commodity hardware has supported variable-width channels out of necessity of narrow-width operation outside the unlicensed bands, this potential seems to have been recognized only recently. Moscibroda et al. [89] have used them for adjusting an AP's channel width based on load, while Chandra et al. [28] have examined their properties in detail for the single-link case. As newer standards such as 802.11-2007 mandate narrow-width channels even in unlicensed bands, we can expect more commodity hardware to offer variable-width channel support.

5.9 Summary

In this chapter, we have proposed a TDMA-based MAC layer for operating point-to-multipoint networks. The base-station that is equipped with multiple radios and multiple antennas is the master node of the TDMA schedule and is responsible for allocating slots to clients according to traffic demand.

However, scaling up the capacity these point-to-multipoint networks raises many design challenges. These include handling interference between the different base-station radios, time synchronization between base-station radios and clients, providing bandwidth and loss guarantees to clients and reducing the cost (both installation and maintenance) of using large number of antennas on towers. We proposed and evaluated the following three techniques to be built on top of the base TDMA MAC protocol to overcome these challenges.

- *Dynamically adapt transmit power:* To minimize interference (both at the base-station and at the clients) resulting from simultaneous operation of multiple co-located base-station radios on different sectors but on the same channel, we evaluated dynamic transmit power control at the base-station and the clients. We formulated the problem to determine the optimal allocation of clients to radios and a schedule of transmissions as a linear program, and compared various transmit power control strategies. We showed that we can schedule up to 50% more clients compared to naive strategies indicating that there is an opportunity to leverage transmit power control in real deployments.

- *Combine radios to physical antennas:* To minimize the cost of installation of towers, the difficulties in maintenance, and to provide flexibility with dynamically changing demands and numbers of active clients, we investigated various architectures that increase capacity while minimizing the number of antennas needed in a point-to-multipoint network.

  We showed through measurements that even a simple static approach of multiplexing radios to antennas using inexpensive off-the-shelf splitters is indeed feasible in practice. pointing to the need for further evaluation of these architectures.

- *Dynamically adapt channel width:* We have examined the theoretical and practical potential of using narrower variable-width channels. Our evaluation on our campus testbed of outdoor links configured as a point-to-multipoint network showed upto 2x throughput improvements and better loss with narrower-width but orthogonal channels compared to using a single wider channel.

We have implemented a basic prototype of a fully featured point-to-multipoint MAC layer that allows us to the above techniques for capacity scaling and throughput improvement. We are currently involved in integrating our various proposals for adaptive power control and dynamic channel adjustment as part of the base MAC protocol. In addition, we also need to implement other standard point-to-multipoint MAC mechanisms such as discovery of new client-stations using a spare contention slot from the base-station, admission control for joining and leaving of client-stations and dynamic adaptation according to traffic demand from clients. Some of these standard mechanisms are well understood from previous work on WiMAX [2] and proposals such as WiFiRe [106].

Going forward, we are also working on new mechanisms that will best utilize smart beamforming antennas to improve the range and flexibility in adjusting to client demand in point-to-multipoint networks. Finally, we plan to undertake a real world deployment of a point-to-multipoint network using our protocols in collaboration with our partners in AirJaldi, India [5].

Chapter 6

Conclusion

We had started this dissertation with the aim of building a new network architecture for rural areas of developing regions using long-distance WiFi links based on the insight that it is most cost effective to cover only those few places where connectivity is required. These goals have been mostly satisfied - we have designed, implemented and deployed MAC- and network-layer mechanisms that to enable the low cost point-to-point and point-to-multipoint networks using off-the-shelf WiFi equipment.

Towards this goal, we studied channel-induced and protocol-induced losses on WiFi links in long-distance settings [129], and addressed these problems by building a WiLDMAC, new TDMA-based MAC with adaptive loss-recovery mechanisms. We used WiLDMAC to improve end-to-end throughput and spectrum usage in multihop point-to-point backbone links [104]. We also use the same WiLDMAC to scale up the capacity of point-to-multipoint access links by proposing three techniques - dynamic transmit power control [105], variable width channels [47] and antenna combination [48].

An important conclusion of our work is that we are no longer limited by performance over long distances in rural networks. Instead, our experiences in deploying and maintaining networks in rural regions of India and elsewhere for the last three years has shown that operational challenges are also a big factor to successful deployments.

In this chapter, we will first present performance results of our MAC protocols from real world testing in Section 6.1. We then describe a real world deployment of our solutions in the Aravind Eye hospital network in Section 6.2. We summarize the contributions of this dissertation in Section 6.3. Finally, we discuss the various limitations of our work in Section 6.4 and future research directions in Section 6.5.

6.1 Real World Performance of WiLDNet

An aim of our work was to make WiFi-based long-distance point-to-point links feasible by achieveing high performance, typically expressed as high throughput and low packet delay. We designed and built WiLDMAC, a TDMA-based MAC with adaptive loss-recovery mechanisms [104]. We showed a 2–5 fold increase in TCP/UDP throughput (along with significantly reduced loss rates) in comparison to the best throughput achievable by the standard 802.11 MAC. We showed these improvements on real medium-distance links and emulated long-distance links (Chapter 4).

We also confirmed the emulated results with data from several real long-distance links in developing regions. Working with Ermanno Pietrosemoli of Fundación Escuela Latinoamericano de Redes (EsLaRed), we were able to achieve a total of 6 Mbps bidirectional TCP throughput (3 Mbps each way simultaneously) over a single-hop 382 km link between Pico Aguila and Platillon in Venezuela using WiLDMAC over 802.11b. Each site used a 2.4

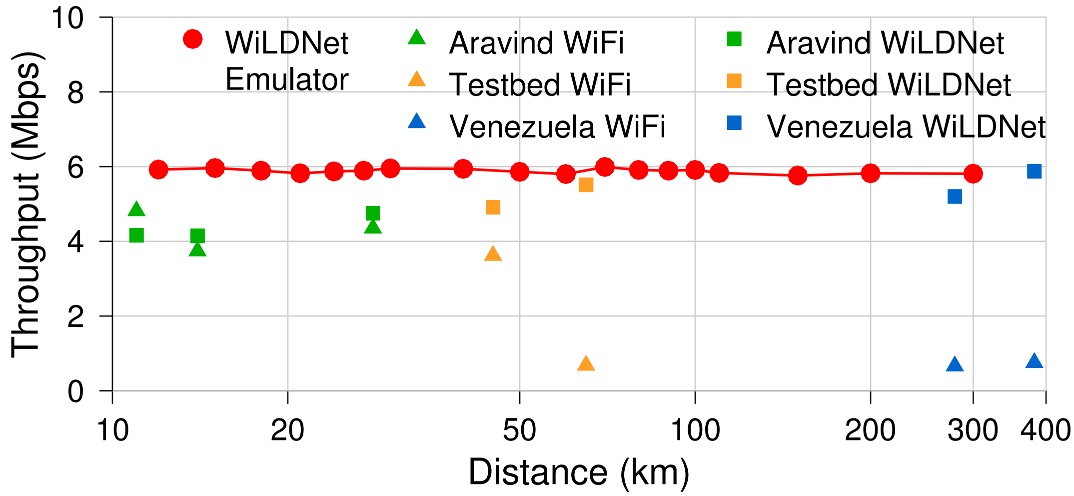

Figure 6.1: Comparison of TCP throughput for WiLDMAC (squares) and standard WiFi MAC (triangles) from links in Aravind, Venezuela, Ghana (the 65 km link), and our local testbed in the Bay Area. Most urban links in Aravind had up to 5–10% loss, and so WiLDMAC did not show substantial improvement over standard WiFi. However, WiLDMAC's advantage increases with distance. Each measurement is for a TCP flow of 60s, 802.11b PHY, 11 Mbps.

GHz 30-dBi reflector grid antenna with 5.3° beam-width and a 400 mW Ubiquiti SR2 radio card with the Atheros AR5213 chipset. To the best of our knowledge, this is currently the longest distance at which a stable high-throughput WiFi link has been achieved without active amplification or custom antenna design.

Figure 6.1 presents results from running WiLDMAC on real links from our various deployments in Aravind (India), Venezuela, Ghana, and our local testbed in the Bay Area. We match the performance of WiLDMAC over emulated links and greatly exceed the performance of the standard WiFi MAC protocol at long distances.

6.2 The Aravind Network

In recent years, our research group, TIER has helped the Aravind Eye Care System [10] to deploy a rural wireless telemedicine system in southern India. Aravind comprises five eye hospitals in Madurai, Theni, Tirunelveli, Coimbatore, and Pondicherry in the state of Tamil Nadu. By volume, it is the largest eye-care provider in the world. In 2006-2007 alone, Aravind saw 2.3 million patients and performed 270,000 surgeries, most of which were for cataracts [11].

6.2.1 Background

Aravind's stated mission is to eradicate needless blindness. The most common causes of preventable blindness are refractive errors and cataracts, which can be treated with prescription glasses and cataract surgery, respectively. With 15 million people needlessly blind, India has the largest share of preventable blindness globally. About 70 percent of India's population is rural, where the risks are higher but access to eye care is lowest [100] The primary limitation to eye care is a severe shortage of trained doctors and nurses in rural areas. India has only 10,000 ophthalmologists serving a population of 1 billion, with 90 percent of doctors based in urban areas [49]. Rural patients must typically travel long distances to clinics or hospitals. Travel expenses, even if not large in absolute terms, can be significant fractions of rural patients incomes, and as a result many are unable or simply decline to get treatment.

To increase utilization of doctors, Aravind has adopted the vision center model, in which doctors remain at the hospital but interact with rural patients over a communication

Figure 6.2: Aravind Village Centers (VCs) are staffed by two people: a technician who operates the ophthalmic equipment and PC, and a counselor who follows up with patients based on the diagnosis provided by a doctor at the base hospital.

network.

A VC, shown in Figure 6.2, is typically a room Aravind rents from a rural family's home in the village. It is equipped with some basic ophthalmic equipment and a PC with a webcam. The center is staffed by two people: a technician who operates the ophthalmic equipment and PC, and a counselor who follows up with patients based on the diagnosis. Center staff generally do not have a degree or a broad technical skill set; Aravind trains them specifically for their duties. At the VC, the technician performs some basic tests for refractive errors and cataracts. The counselor presents the results to the doctor at the base hospital via a videoconference, after which the patient interacts with the doctor. The counselor then follows up on the doctor's advice, for example, by handing out prescriptions, filling out referral forms, or creating glasses. If advised by the doctor, the counselor refers the patient to the base hospital for further examinations or treatments such as cataract surgery. The cost to the patient for a VC consultation visit is 25 cents. Cataract surgery, if required, can cost up to $75 at the hospital surgery of comparable quality in the US costs about $2,000. However, about two-thirds of patients cannot pay and receive surgery for free, which is also true for non-VC patients; these procedures are subsidized by paying patients.

6.2.2 Network description

In 2005, Aravind's eye hospital at Theni created three VCs based on corDECT wireless local-loop technology [13], supplied by a local carrier focusing on rural connectivity. Each site, including the base hospital, had a total bandwidth of 36.5 kilobits per second. Not surprisingly, the video quality was insufficient, although the audio had some value. Going through a carrier limited Aravind's ability to start centers in areas with dire need. Despite being ready with clinical equipment and personnel, it could not start VCs in two locations as the carrier did not consider those areas profitable enough to deploy a base-station. The same year, we established our own long-distance Wi-Fi link as an alternative connection for the Ambasamudram VC in Theni, mostly for operational experience. By early 2006, satisfied with the high performance of 5-6 Mbps per link and the operational freedom of an unlicensed spectrum, Aravind phased out the corDECT links, converting the existing three VCs and completing two others.

The Aravind network at Theni consists of five vision centers connected to the main hospital in Theni (Figure 6.3). The network has total of 11 wireless routers (6 endpoints, 5 relay nodes) and uses 9 point-to-point links. The links range from just 1 km (Theni–Vijerani) to 15 km (Vijerani–Andipatti). Six of the wireless nodes are installed on towers, heights of which range from 24–42 m; the others use short poles on rooftops or existing

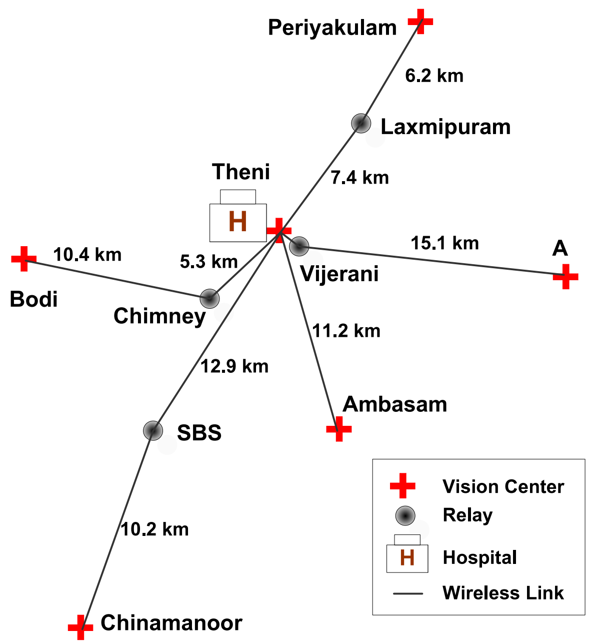

Figure 6.3: Aravind Telemedicine Network. Theni hospital is connected to 5 vision centers. The other nodes are all relays.

tall structures, such as the chimney of a power plant on the premises of a textile factory. Recently, Aravind has expanded this model to their hospitals in Madurai and Tirunelveli where they have added more vision centers. The network is currently financially viable and a further expansion to 50 clinics around 5 hospitals is being planned to provide 500,000 annual eye examinations.

Hardware: The wireless nodes are 266 MHz x86 single-board computers. These routers have up to 3 Atheros 802.11 a/b/g radio cards (200–400 mW). The longer links use 24 dBi directional antennas. The software of the routers include the base Linux operating system, our WiLDNet MAC drivers and tools for monitoring, logging, and remote management. The routers are placed in small and lightweight waterproof enclosures, and are mounted externally, close to the antennas, to minimize signal losses. They are powered via power-over-ethernet (PoE); a single ethernet cable from the ground to the router is sufficient.

Applications: The primary application is video-conferencing. We currently use software from Marratech [84]. Although most sessions are between doctors and patients, we also use the video conferencing for remote training of staff at vision centers. Typical throughput on the links ranges between 5–7 Mbps with channel loss less than 2%. But 256 Kbps in each direction is sufficient for very good quality video conferencing. Our network is thus over provisioned, and we also use the network to transmit 4-5 MB-sized retinal images. The hospital has a VSAT link to the Internet, but most applications require only intranet access within the network (except for remote management).

6.2.3 Operational Issues

To leverage the range of skills and availability of potential support personnel, we have created a three-tiered support system. Tier 1 consists of local staff responsible for basic management and maintenance. Tier 2 includes local network integrators and local vendors trained in installation, configuration, and debugging of networking components. Tier 3 consists of the remote management team, comprising of highly skilled professionals familiar with all the hardware and software. At Aravind, TIER has played the role of the remote management team. We also built a monitoring system, alternative backchannels for remote management during link failures, and automatic-recovery mechanisms that together have improved operational sustainability.

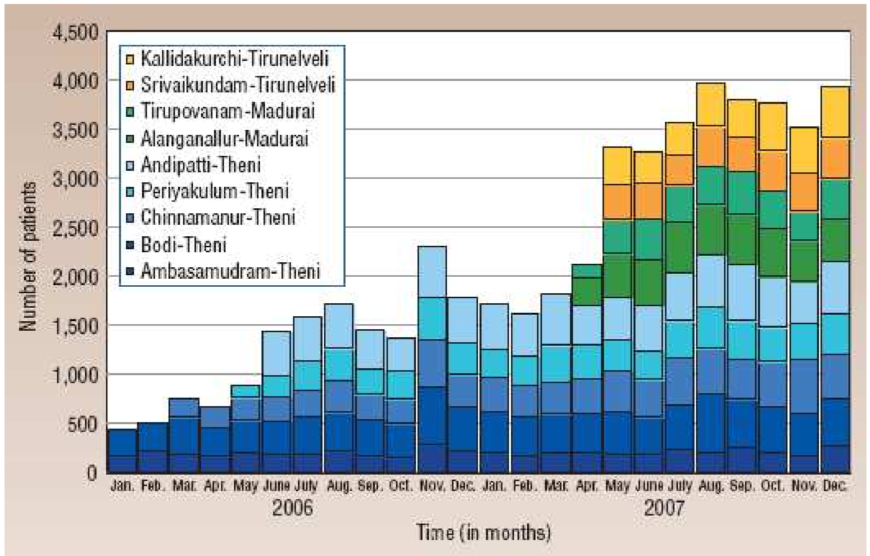

Figure 6.4: The graph shows patients treated per month per center for two years TIER set up the links to all Theni VCs; Aravind set up subsequent links to the new centers in Madurai and Tirunelveli.

6.2.4 Sustainability

Our work at Aravind reflects the optimization principle: The doctors are great; we only needed to improve rural access to them. Figure 6.4 shows the growth in the number of Aravind patients and VCs from January 2006 to December 2007. Ambasamudram, Andipatti, and Bodi were the existing centers that migrated to our high-bandwidth links, while the rest are new VCs enabled by our solution. Overall, the network enabled 51,205 remote eye examinations during that two-year period. From May 2007, when all nine VCs were up, until December 2007, the system served an average of 3,632 patients per month. About 75 percent of patients visiting all VCs were new, while the remaining 25 percent came in for follow-ups. According to Aravind, this implies more extensive eye care than the eye-camp approach because new patients get treatment every month, and higher-quality eye care with increased patient follow-ups. Overall, 9,835 patients were diagnosed with severe cataract or refractive errors and needed significant vision improvement. Of these, 90 percent (8,814) got their sight back through prescription glasses or cataract surgeries as advised by the doctor during the videoconference. None of these villages have any ophthalmologists, except for Bodi, which has one doctor in private practice who visits once a week. We know from interviews that these patients generally would not have received treatment if not for the VCs.

These patients are also likely to return to income generation, the first step out of poverty. A recent study revealed that 96 percent of Aravind patients who get cataracts stop working [74]. Among those who lost their jobs, about 85 percent of men and 58 percent of women who get surgery return to wage-earning activities within a week.

Over time, operational responsibility for the network has migrated from us (Tier 3) to local staff (Tier 1), and we consider this a big success. In the beginnning, we were solely responsible for the entire network. Through training and development of management tools, local staff has learned to maintain and manage the network on their own, while a local vendor (Tier 2) handles tower construction, antenna alignment, and other installation issues. In the past year we have not installed any links ourselves even though Aravind has established four additional VCs. Our role is reduced to supplying equipment for new wireless installations, which Aravind now also pays for as the centers have demonstrated they can recover network costs.

6.3 Contributions

Our proposed WiLDNet network architecture for rural areas consists of a combination of long-distance point-to-point backhaul links, and medium-range point-to-multipoint access links to provide high performance wireless connectivity. We now present a brief summary of the contributions of this dissertation in tackling challenges at different aspects of this architecture.

6.3.1 Understanding poor performance of WiFi

To explain the poor performance of WiFi in outdoor long-distance environments, we conducted a rigorous measurement study to identify the causes of high and variable loss using indoor emulated links and outdoor testbeds [129]. Our findings were as follows:

- We found out that the 802.11 MAC protocol is ill-suited for long-distance links. In particular, the default link-level recovery mechanism results in low channel utilization at long distances and frequent collisions occur because of the failure of CSMA/CA.
- Although adjusting some 802.11 timing parameters (such as $ACKTimeout$ and $SlotTime$) can decrease the number of collisions at long distances, the overall throughput achieved is still much below the potential throughput.
- Outdoor wireless links also suffer from bursty and variable packet loss that can be traced to interference from external WiFi sources, and this variability causes to TCP flows to stall and experience timeouts.
- In multihop settings, co-located radios (same physical location) on the same router op-

erating on the same wireless channel interfere with each other if one of them transmits while the other receives.

6.3.2 Building WiLDMAC: A novel TDMA-based MAC

To overcome the protocol problems with 802.11, we designed and built a new MAC protocol that uses time-division based slots instead of CSMA/CA. We implemented this new MAC in Linux on top of the madwifi driver for Atheros radio cards using the Click modular router framework and tested it on real WiFi hardware. The main features of this MAC protocol are:

- The use of time-division slots eliminates all collisions that happen at long distances from the failure of CSMA to detect transmissions from other end-points in time.
- We used sliding-window flow control instead of the stop-and-wait packet transmission of 802.11 to improve channel utilization irrespective of the propagation delay on the link.
- We used bulk ACKs (aggregated over multiple packets in a slot) instead of per-packet acknowledgments used in 802.11 to implement retransmission based loss recovery (ARQ).

6.3.3 Achieving high throughput on point-to-point backbone network

We used WiLDMAC, the base TDMA-based MAC to improve end-to-end performance for the multihop long-distance point-to-point backbone part of our network architecture [104]. The main results show that:

- We achieved sustained high TCP and UDP throughput (6 Mbps, bidirectional) on a single long-distance 802.11b WiFi link at any distance (tested upto 382 km).

- For delay sensitive traffic, we implemented forward error correction (FEC) as the primary loss recovery mechanism. We showed that depending on application requirements, we can use a combination of ARQ and FEC to tradeoff between average delay and bandwidth overhead of loss recovery.

- We implemented the *SynOp* constraint i.e. transmit simultaneously from all links at a node in the same slot to all its neighbors (similarly for receive slot). This allows us to operate the whole (bipartite) network on the same channel, thereby optimizing spectrum usage.

- We used an implicit synchronization mechanism to align the time slots between neighboring nodes. We showed that this mechanism is resilient to packet loss on the channel.

6.3.4 Capacity scaling in point-to-multipoint access networks

We used WiLDMAC, the same base TDMA-based MAC to implement a point-to-multipoint MAC protocol and investigated three techniques to a) enable scaling up capacity of the network, b) handle dynamic client traffic demand and c) lower cost of deployment.

- We showed that smart transmit power control and slot scheduling at the base-station for simultaneous transmission from different sector antennas on the same channel substantially increases the number of clients that can be accommodated and thereby the total throughput of the network.

- We showed that it is feasible to connect multiple radios to the same physical antenna thus greatly reducing the difficulty and cost of installing and aligning many sector antennas on one tower. We also explored other architectures to allow radios to dynamically switch antennas to adapt to dynamic client traffic demand.

- We showed that using variable width channels at the base-station instead of just using the small number of fixed wireless channels as defined by 802.11 (only three for 11b) can significantly improve cumulative capacity in a point-to-multipoint network by increasing the number of simultaneous transmissions in the network.

6.3.5 Deploying in real world scenarios

Finally, we validated our solutions by undertaking several deployments of our long-distance wireless links in India (for the Aravind telemedicine network), Uganda, Ghana and the Philippines. As we saw earlier, our network in Aravind that targets doctor-patient video-conferencing [10] has been successfully running for more than two years, has expanded to 10 vision centers, has restored vision to about 10,000 patients (by 2007) and has achieved operational sustainability.

6.4 Limitations

Although our proposed solutions address several problems and enable many applications with development impact, there are some limitations to our approach. While some of these limitations are in the process of being addressed, others need more research.

6.4.1 Topology restrictions

We chose WiFi to build rural wireless networks because it could get high performance, good range and low cost. Thus, it is very well suited for connecting places that have line of sight with high throughput links at very long distances. However, to satisfy line of sight requirements, we often need to place the wireless routers and the directional antennas at higher elevation on top of towers or poles. Establishing a link also often involves antenna alignment to maximize the signal strength. All these factors make WiFi less than ideal for scenarios where we don't have line of sight or need more flexibility to move the end-points in the network.

In addition, WiFi does not work with mobile clients very well since it was not designed to handle fading that occurs with high mobility. In general, WiFi chipsets also consume more power than radios designed for cellular networks such as GSM that are highly optimized for battery operated cellphones.

While some of these challenges can be overcome by using smart electronically steerable antennas for dynamic alignment, for other challenges, we need to combine WiFi with other wireless technologies that are more suited for mobile devices.

6.4.2 Bipartite topologies and fixed time slots

While WiLDMAC can run on multihop long-distance wireless networks using only a single channel, it also requires the topology to be bipartite; for non-bipartite topologies, we can run WiLDMAC only on a bipartite subgraph of the overall topology graph.

Figure 6.5 shows an example of such a bipartite network. Using 2P or WiLDMAC, all nodes in partition A first transmit on all of their links (for a time slot of size $t_{A \to B}$).

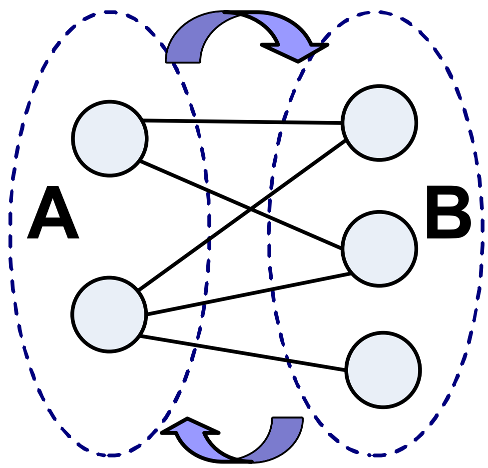

Figure 6.5: SynOp operation in 2P

Following this, all nodes in partition B transmit on all their links (for a time slot of $t_{B \to A}$). The ratio between these slot sizes regulates the bandwidth allocation for every network link between the two partitions. In practice, $t_{A \to B}$ and $t_{B \to A}$ are almost always set to be equal since this maximizes throughput for traffic paths spanning more than two hops [30, 110]. Because of this constraint on having fixed-length slots, WiLDMAC cannot adapt to dynamic traffic variations.

The ideal MAC protocol would run without any topology-related constraints (beyond the usual line-of-sight constraints) and also adapt the length of slots according to changing traffic conditions.

These limitations are addressed in JazzyMAC [92, 93], which is an extension of WiLDMAC. JazzyMAC is a fully distributed, practical MAC-layer that uses local traffic information to adapt the transmission slot sizes. It uses dynamic slot sizing to negotiate the delay-throughput tradeoff in WiLD networks, and exploits asymmetric traffic, time varying traffic, and non-bipartite topologies. It also enables more efficient use of network capacity by allowing more parallel communication among non-interfering links.

Simulation results show that JazzyMAC achieves superior throughput (with up to 100% improvement) over 2P and WiLDMAC across various network sizes, topologies, and traffic workloads. Moreover, this improvement increases dramatically in the case of asymmetric traffic, a commonly-occurring workload in rural Internet access.

6.5 Future Directions

The work from this dissertation leads towards several new research directions. This includes building custom radios that are optimized for long distance operation, and using smart antennas that can substantially increase the range and robustness of long-distance wireless networks leading to better operational sustainability.

6.5.1 Hardening of the WiLDNet platform

Although we have deployed and experimented with our new MAC protocols in a number of real world wireless links, installing and configuring them is still a non-trivial job requiring expertise of networking concepts and Linux. Often for a new deployment, we were required to provide significant handholding to local administrators before they were confident enough to manage and debug the network on their own.

We are currently exploring ways to convert our *research* platform into a more *production ready* platform. This involves work to refine various aspects of the platform – a) the hardware base including computer board, radio cards, connectors and stable power solutions; b) the Linux software operating system; c) the modified MAC protocols and d) management tools (including remote monitoring and administration). We need to standardize this platform, make it robust to power failures, and write documentation so that anybody interested in building such networks can start on their own.

6.5.2 More efficient implementation of protocols

While WiLDMAC currently achieves 6 Mbps bidirectional throughput with both TCP and UDP on long-distance links with 802.11b, this throughput is substantially less than maximum raw throughput of 802.11b (around 7.5 Mbps). This 15% overhead is because of inefficient implementation of the TDMA implementation that is done by scheduling time slots using Click modular router in our current implementation. We plan decrease this overheard by implementing more accurate time slots using hardware timers on the Atheros card.

To support adaptation to dynamic traffic demands and non-bipartite topologies, we plan to implement the JazzyMAC extensions on top of the WiLDMAC base protocol.

For the point-to-multipoint MAC protocol that currently only supports static configuration of clients, we are in the process of building beaconing mechanisms for client discovery, joining and leaving, and admission control.

6.5.3 Custom Radio design for long distances

Although we have used pre-existing off-the-shelf radios (e.g. 802.11a/b/g and 802.11n in the future) for our long-distance links, an important question that is often asked is what is best radio design for long-distance environments?

The main advantage of using pre-existing radios is cost and reliability (equipment have been tested and operated on field). However, there are several drawbacks. The first drawback is related to the limited bandwidth and the second drawback stems from the inadequacy of the OFDM and DSSS modulation schemes for long-distance operation.

In order to get at least 100 Mbps reliably over long-distance, the total bandwidth must be at least 100 MHz. 802.11n (or superG) radios have a maximum bandwidth of 40 MHz. Therefore, in order to achieve the desired data-rates, we have to aggregate multiple radios. However, transmitting on multiple radios may result in packets received out of order, which might impact the higher network layers. Also, if we want to take advantage of the entire available unlicensed spectrum from 5–6GHz, the required number of radios could be large.

The second disadvantage of using 802.11 based standards is that their modulation schemes are not designed or optimized for the long-distance application. 802.11 standards use either DSSS (802.11b), which is not spectrally efficient, or OFDM (802.11agn), which is not power efficient (high peak to average ratio). Although these modulation schemes make sense for the original intent of these radios (indoor and urban outdoor networks), since they provide extra protection against fading and inter-symbol interference (ISI), which are common for these types of networks. However, these issues do not exist in long-distance links where we have line of sight with narrow beams.

A simple constant envelope modulation scheme (e.g. QPSK or differential QPSK) would provide robust high throughput long-distance communication, and be more immune to external interference, ISI, synchronization errors, coupling between antenna elements (especially when multiple data streams are transmitted and received simultaneously). In addition, using custom radios can reduce the MAC-layer overhead and enable more accurate time slots in hardware.

6.5.4 Smart Steerable Antennas

As we mentioned earlier in Chapter 2, steerable antennas not only increase range and coverage but also make installation and operation of long-distance wireless networks much more easier.

A steerable antenna can focus its beam on its desired target, which limits the interference it experiences from other radios that are sufficiently far away. As a result, beamforming can potentially improve the overall system capacity by allowing multiple pairs of radios to have simultaneous conversations (spatial reuse), or improve the capacity of a single network by allowing a single radio to transmit multiple independent data streams to multiple radios with sufficient spatial separation (spatial multiplexing), or even improve the capacity of a single link (in a multipath environment) by allowing a radio to transmit independent data streams on different paths and taking advantage of reflections, often referred to as multi-input/multi-output or MIMO.

Currently, there are promising steerable antenna technologies that could be used in low-cost wireless networks.

- *Phased-array beamforming antennas:* Beamforming is usually implemented using an array of omni-directional or directional antennas. A beam will arrive at the different antenna elements with different phase delays . The delays will in general be a function of the direction of arrival. By applying the appropriate phase-shifts (hence the name phased arrays) to the incoming signals at each antenna, the array can combine signals from the desired direction in-phase while attenuating signals from other directions.

The ability of the array to resolve (or distinguish) beams from different directions depends on the beam-width, which depends on the size of the array and the number of antennas. Antenna elements are usually placed at least half a wavelength apart. Large arrays (more antenna elements) are usually required to achieve narrower beams and higher spatial resolution

Today, are commercial phased-array beamforming solutions that sell for around $3000 [40]. The high cost results from the fact that the phase delays are introduced by converting the signal into a digital domain. However, there are current efforts to build lower cost arrays where the phase delay can be introduced directly in the RF domain. These arrays with could be mass produced for less than $100 [12].

- *Parasitic antennas:* The second approach builds an adaptive parabolic reflector from an array of passive scatterers with tunable reactive loads. The system adjusts the bias voltages on the loads to achieve the best received signal strength and thus adapt the beam. This technology is still under development, but if manufactured in a large scale, antennas with a peak gain of upto 24 dBi can be built for less than $200 [82, 136].

Making full use of steerable antenna capabilities also requires work on new antenna adaptation algorithms, and novel MAC- and network-layer mechanisms.

Adaptation algorithms are important because the array has to adapt and converge on the desired neighbors. For example, a router with a long-distance link has to continuously adapt to account of antenna misalignment. A base-station serving many clients has to adapt and beamform within the packet reception time ($10\mu s$) to be able to receive and send data at the highest SNR.

MAC- and network-layer protocols are important because channel access and route selection in a large wireless network would determine how we form beam patterns for the antenna arrays. In a wireless mesh network, the network layer needs to adapt the beam pattern to choose the best next hop for maximizing throughput. In a wireless LAN, the MAC-layer at the access point has to decide when to change its beam to a new client that is expected to transmit packets to the access point.

6.6 Summary

Finally, from our research over the last five years, we have realized that, to have real sustained impact of any technology to lives of people, we must deal with complex social issues such as underlying gender and ethnic inequalities, as well as existing players that might be negatively impacted by ICT. Therefore, the right strategy is to work closely with social scientists and to partner with strong government or non-governmental organizations (NGOs), who tend to understand local needs and dynamics in a way that is not possible from afar.

ICTs cannot be a panacea for all the complex problems facing nations on the path to economic development. On the contrary, at best ICT can enable new solutions only when applied with a broad understanding and a multi-disciplinary approach.

Bibliography

[1] The Technology and Infrastucture for Emerging Regions (TIER) Research Group. `http://tier.cs.berkeley.edu`.

[2] WiMAX forum. `http://www.wimaxforum.org`.

[3] Abdus Salam International Centre for Theoretical Physics. School on New Perspectives on Wireless Networking. `http://wireless.ictp.it/`.

[4] Daniel Aguayo, John Bicket, Sanjit Biswas, Glenn Judd, and Robert Morris. Link-level Measurements from an 802.11b Mesh Network. In *ACM SIGCOMM*, August 2004.

[5] AirJaldi Wireless Network. `http://summit.airjaldi.com`.

[6] Airspan Incorporated. Airspan Networks WipLL Solution. `http://www.airspan.com/`.

[7] Airvana Networks. `http://www.airvananet.com`.

[8] Akshaya E-Literacy Project. `http://www.akshaya.net`.

[9] Andreas and Olov Schelen. Quality of Service Schemes for IEEE 802.11 Wireless LANs: An Evaluation. *Module Networks and Application, Kluwer*, 2003.

[10] Aravind Eye Care System. `http://www.aravind.org`.

[11] Aravind Eye Care System. *Annual Activities Report*. Aravind Eye Hospitals, 2005.

[12] O Bakr, M Johnson, B Wild, and K Ramchandran. A Multi-Antenna Framework for Spectrum Reuse Based on Primary-Secondary Cooperation. In *New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, 2008.

[13] A. Balaji, T. A. Gonsalves, and R. Kumar. The corDECT Wireless Local Loop: Capacity Predictions. In *Personal Wireless Communication*, 1999.

[14] Hari Balakrishan. *Challenges to Reliable Data Transport over Heterogeneous Wireless Networks*. PhD thesis, University of California at Berkeley, August 1998.

[15] Hari Balakrishnan, Srinivasan Seshan, Elan Amir, and Randy Katz. Improving TCP/IP Performance over Wireless Networks. In *ACM MOBICOM*, November 1995.

[16] Bay Area Research Wireless Network. `http://www.barwn.org`.

[17] Pravin Bhagwat, Bhaskaran Raman, and Dheeraj Sanghi. Turning 802.11 Inside-out. In *Hotnets-III*, 2004.

[18] G Bianchi. Performance Analysis of the IEEE 802.11 Distributed Coordination Function. *IEEE Journal on Selected Areas in Communications*, 18(3):535–547, 2000.

[19] G Bianchi and I Tinnirello. Remarks on IEEE 802.11 DCF Performance Evaluation. *IEEE Communications Letters*, 9(8):765–767, 2005.

[20] Sanjit Biswas and Robert Morris. Opportunistic Routing in Multi-Hop Wireless Networks. In *Hotnets-II*, 2003.

[21] K. Brauchli, D. O'Mahony, L. Banach, and M. Oberholzer. iPath-a Telemedicine Platform to Support Health Providers in Low Resource Settings. *Studies in Health Technology and Informatics*, 114:11–7, 2005.

[22] Eric Brewer, Michael Demmer, Bowei Du, Melissa Ho, Matthew Kam, Sergiu Nedevschi, Joyojeet Pal, Rabin Patra, Sonesh Surana, and Kevin Fall. The Case for Technology in Developing Regions. *IEEE Computer*, 38(6):25–38, 2005.

[23] Michael Buettner, Eric Anderson, Gary Yee, Dola Saha, Anmol Sheth, Douglas Sicker, and Dirk Grunwald. A Phased Array Antenna Testbed for Evaluating Directionality in Wireless Networks. In *International workshop on System Evaluation for Mobile Platforms (MobiEval)*, 2007.

[24] Viveck Cadambe, Syed A. Jafar, and Shlomo Shamai. Interference Alignment on the Deterministic Channel and Application to Fully Connected AWGN Interference Networks. Nov 2007. `http://arxiv.org/abs/0711.2547v1`.

[25] Manuel Castells. *Rise of the Network Society*. Blackwell Publishers, 1996.

[26] Champaign-Urbana Community Wireless Network. `http://www.cuwin.net`.

[27] Y. Chandani and G. Breton. Contraceptive Security, Information Flow and Local Adaptations: Family Planning in the Philippines and Morocco. *EJISDC*, 5(3):1–18, 2001.

[28] Ranveer Chandra, Ratul Mahajan, Thomas Moscibroda, Ramya Raghavendra, and Paramvir Bahl. A Case for Adapting Channel Width in Wireless Networks. In *ACM SIGCOMM*, 2008.

[29] Kameswari Chebrolu and Bhaskaran Raman. FRACTEL: A Fresh Perspective on (Rural) Mesh Networks. In *ACM SIGCOMM Workshop on Networked Systems for Developing Regions (NSDR)*, August 2007.

[30] Kameswari Chebrolu, Bhaskaran Raman, and Sayandeep Sen. Long-Distance 802.11b Links: Performance Measurements and Experience. In *ACM MOBICOM*, 2006.

[31] C.M. Cheng, P.H. Hsiao, HT Kung, and D. Vlah. Adjacent Channel Interference in Dual-radio 802.11a Nodes and Its Impact on Multi-hop Networking. In *IEEE Globecom*, 2006.

[32] Romit Roy Choudhury, Xue Yang, Ram Ramanathan, and Nitin H. Vaidya. On Designing MAC Protocols for Wireless Networks Using Directional Antennas. *IEEE Transactions on Mobile Computing*, 5(5):477–491, 2006.

[33] Martin V. Clark, K.K. Leung, B. McNair, and Zoran Kostic. Outdoor IEEE 802.11 Cellular Networks: Radio Link Performance. In *IEEE ICC*, 2002.

[34] CRCNet. Connecting Rural Communities Using WiFi. `http://www.crc.net.nz`.

[35] Digital Gangetic Plains. `http://www.iitk.ac.in/mladgp/`.

[36] Christian Doerr, Michael Neufeld, Jeff Filfield, Troy Weingart, Douglas C. Sicker, and Dirk Grunwald. MultiMAC - An Adaptive MAC Framework for Dynamic Radio Networking. In *IEEE DySPAN*, November 2005.

[37] Engenius Technologies. `http://www.engeniustech.com/`.

[38] Enlace HispanoAmericano de Salud (EHAS) Foundation. `http://www.ehas.org`.

[39] Ermanno Pietrosemoli. Long Distance WiFi Trial. *International Summit for Community Wireless Networks*, 2007. `http://www.eslared.org.ve/articulos/Long%20Distance%20WiFi%20Trial.pdf`.

[40] Fidelity Comtech. Phocus Phased Array Antennas. `http://www.fidelity-comtech.com/`.

[41] Rob Flickenger. *Wireless Networking in the Developing World. A Practical Guide to Planning and Building Low-cost Telecommunication Infrastructure. Second Edition.* Limehouse Book Sprint Team, 2008. `http://wndw.net`.

[42] R Gaiha. Decentralization and Poverty Alleviation in Developing Asia. *Economics and Development Resource Center, Asian Development Bank, Manila*, 2001.

[43] R. Gandhi, R. Veeraraghavan, K. Toyama, and V. Ramprasad. Digital Green: A Participatory Digital Framework to Deliver Targeted Agricultural Information to Small and Marginal Farmers. In *Information Technologies and Development (ICTD)*, 2007.

[44] Andrea Goldsmith. *Wireless Communications*. Cambridge University Press, 2005.

[45] Shyamnath Gollakota and Dina Katabi. ZigZag decoding: Combating hidden terminals in wireless networks. In *ACM SIGCOMM*, 2008.

[46] Grameen Bank. Grameen Telecom Overview. `http://www.grameen-info.org/grameen/gtelecom/`.

[47] Ramakrishna Gummadi, Rabin Patra, Hari Balakrishnan, and Eric Brewer. Interference Avoidance and Control. In *HotNets-VII*, 2008.

[48] Ramakrishna Gummadi, Rabin K. Patra, Sergiu Nedevschi, Sonesh Surana, and Eric A. Brewer. A Radio Multiplexing Architecture for High Throughput Point to Multipoint Wireless Networks. In *ACM MOBICOM Workshop on Wireless Networks and Systems for Developing Regions (WiNS-DR)*, 2008.

[49] GV Murthy and SK Gupta and D Bachani and HK Tewari and N John. Human Resources and Infrastructure for Eye Care in India: Current Status. *National Medical Journal of India*, 17, 2004.

[50] Daniel Halperin, Josephine Ammer, Thomas Anderson, and David Wetherall. Interference Cancellation: Better receivers for a new Wireless MAC. In *HotNets-VI*, 2007.

[51] N Hanna. Exploiting information technology for development: A case study of India. *World Bank*, 1994.

[52] S Hasan. Introducing E-Government in Bangladesh: Problems and Prospects. *International Social Science Review*, 78(3-4):111–126, 2003.

[53] Richard Heeks. ICT4D 2.0: The Next Phase of Applying ICT for International Development. *IEEE Computer*, pages 26–33, 2008.

[54] C. Heffernan. The Livestock Guru: Demand-led knowledge transfer for poverty alleviation. In *Information and Communication Technologies and Development (ICTD)*, pages 94–100, 2006.

[55] Tristan Henderson, David Kotz, and Ilya Abyzov. The changing usage of a mature campus-wide wireless network. In *ACM MOBICOM*, 2004.

[56] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance Anomaly of 802.11b. In *IEEE INFOCOM*, 2003.

[57] HyperLink. HG2427G 2.4GHz 27dBi High Performance Die Cast Reflector Grid Wireless LAN Antenna. `http://www.hyperlinktech.com/multimedia/datasheets/DS_HG2427G.PDF`.

[58] M. Ibrahim, B. Bellows, M. Ho, and S. Surana. Improving Healthcare Delivery in Developing Regions: Coupling Output-based Aid and Mobile Technologies. In *University of California Big Ideas Conference*, 2007.

[59] IEEE 802.11. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification, 1999.

[60] IEEE 802.11a. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification: High Speed Physical Layer Extension in the 5GHz Band, 1999.

[61] IEEE 802.11b. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification: High Speed Physical Layer Extension in the 2.4GHz Band, 1999.

[62] IEEE 802.11n. 802.11n (D7.0) Draft Standard for 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), 2007.

[63] ILOG. CPLEX LP Solver. `http://www.ilog.com`.

[64] Indian Institute of Technology, Madras. The TENET Research Group. `http://www.tenet.res.in`.

[65] Intel. Classmate PC. `http://www.intel.com/intel/worldahead/classmatepc`.

[66] Inveneo. `http://www.inveneo.org`.

[67] Yousuf M. Islam, Manzur Ashraf, Zillur Rahman, and Mawdudur Rahman. Mobile Telephone Technology as a Distance Learning Tool. In *ICEIS*, 2005.

[68] ITC Limited. The e-Choupal Initiative. `http://www.echoupals.com`.

[69] Kyle Jamieson, Bret Hull, Allen Miu, and Hari Balakrishnan. Understanding the real-world performance of carrier sense. In *ACM SIGCOMM E-WIND Workshop*, 2005.

[70] Amit P. Jardosh, Krishna N. Ramachandran, Kevin C. Almeroth, and Elizabeth M. Belding-Royer. Understanding link-layer behavior in highly congested ieee 802.11b wireless networks. In *ACM SIGCOMM E-WIND Workshop*, 2005.

[71] A. Jhunjhunwala. Towards Enabling India through Telecom and Internet Connections. *Variously published essay available online*, 2000. `http://www.tenet.res.in`.

[72] A. Jhunjhunwala, A. Ramachandran, and A. Bandyopadhyay. n-Logue. *The Journal of Community Informatics*, 1(1), 2004.

[73] M. Kam, D. Ramachandran, A. Raghavan, J. Chiu, U. Sahni, and J. Canny. Practical considerations for participatory design with rural school children in underdeveloped regions: early reflections from the field. *Interaction Design And Children*, pages 25–32, 2006.

[74] Kate Gooding. Poverty and Blindness: A Survey of the Literature. *Sight Savers International*, February 2006.

[75] P. D. Kaushik and N. Singh. Information Technology and Broad-Based Development: Preliminary Lessons from North India. *World Development*, 32(4):591–607, 2004.

[76] G R Kiran. E-governance services through Telecentres - Role of Human Intermediary and issues of Trust. In *Information and Communication Technologies and Development (ICTD)*, pages 131–142, 2006.

[77] Murali Kodialam and Thyagarajan Nandagopal. Characterizing the Achievable Rates in Multihop Networks. In *ACM MOBICOM*, September 2003.

[78] Eddie Kohler, Robert Morris, Benjie Chen, John Jannotti, and Frans M. Kaashoek. The Click Modular Router. *ACM Transactions on Computer Systems*, 18(3):263–297, August 2000.

[79] F Librero. Uses of the Cell Phone for Education in the Philippines and Mongolia. *Distance Education*, 28(2):231–244, 2007.

[80] Shu Lin and Daniel Costello. *Error Control Coding: Fundamentals and Applications*. Prentice Hall, 1983.

[81] Rowena Luk, Melissa Ho, and Paul M. Aoki. Asynchronous remote medical consultation for Ghana. In *ACM SIGCHI*, 2008.

[82] A Mainwaring, A L Umnov, V A Filimonov, I S Shishalov, and M Shuralev. Theoretical and Experimental Investigations of Adaptive Reactance Parasitic Antenna Array. *Mobile and Wireless Communications Summit*, 2007.

[83] S. Mangold, S. Choi, P. May, O. Klein, G. Hiertz, and L. Stibor. IEEE 802.11 e Wireless LAN for Quality of Service. In *European Wireless*, volume 18, pages 32–39, 2002.

[84] Marratech. Videoconferencing Software. `http://www.marratech.com`.

[85] David L. Mills. Internet Time Synchronization: The Network Time Protocol. *Global States and Time in Distributed Systems, IEEE Computer Society Press*, 1994.

[86] A. Mishra, M. Shin, and W. Arbaugh. An empirical analysis of the IEEE 802.11 MAC layer handoff process. *ACM SIGCOMM Computer Communication Review*, 33(2):93–102, 2003.

[87] Shridhar M. Mishra, John Hwang, Dick Filippini, Tom Du, Reza Moazzami, and Lakshminarayanan Subramanian. Economic Analysis of Networking Technologies for Rural Developing Regions. In *WINE*, December 2005.

[88] N Moraveji, U Pawar, and T Kim. A Mouse on Each Desk: An Inexpensive Classroom Interaction Technique for Remote Teaching. *Microsoft Research Working Papers*, 2006.

[89] Thomas Moscibroda, Ranveer Chandra, Yunnan Wu, Sudipta Sengupta, and Paramvir Bahl. Load-Aware Spectrum Distribution in Wireless LANs. In *ICNP*, 2008.

[90] nanoGSM Picocells. `http://www.ipaccess.com/picocells/nanoGSM.php`.

[91] Vishnu Navda, Anand Prabhu Subramanian, Kannan Dhanasekaran, Andreas Timm-Giel, and Samir Das. MobiSteer: Using Steerable Beam Directional Antenna for Vehicular Network Access. In *MOBISYS*, 2007.

[92] Sergiu Nedevschi. *Maximizing Performance in Long Distance Wireless Networks for Developing Regions*. PhD thesis, Berkeley, CA, USA, 2008. Adviser-Brewer,, Eric.

[93] Sergiu Nedevschi, Rabin Patra, Sonesh Surana, Sylvia Ratnasamy, Lakshminarayanan Subramanian, and Eric Brewer. An Adaptive, High Performance MAC for Long-Distance Multihop Wireless Networks. In *ACM MOBICOM*, 2008.

[94] Nicholas Negroponte. *Being Digital*. Random House Inc. New York, 1996.

[95] Nepal Wireless Networking Project. `http://www.nepalwireless.net`.

[96] Michael Neufeld, Jeff Fifield, Christian Doerr, Anmol Sheth, and Dirk Grunwald. SoftMAC - Flexible Wireless Research Platform. In *HotNets-IV*, November 2005.

[97] G Nulens and L Van Audenhove. An Information Society in Africa?: An Analysis of the Information Society Policy of the World Bank, ITU and ECA. *International Communication Gazette*, 1999.

[98] NY Wireless Network. `http://www.nycwireless.net`.

[99] Government of India. Community Information Centers. `http://www.cic.nic.in`.

[100] P K Nirmalan, J Katz, A L Robin, R Krishnadas, R Ramakrishnan, R D Thulasiraj, J Tielsch. Utilisation of Eye Care Services in Rural South India: the Aravind Comprehensive Eye Survey. *British Journal of Ophthalmology*, 88, 2004.

[101] Joyojeet Pal, M Lakshmanan, and Kentaro Toyama. My Child Will be Respected: Parental Perspectives on Computers in Rural India. In *Information Technologies and Development (ICTD)*, volume 2, pages 168–177, Bangalore, 2007.

[102] T Parikh, K Ghosh, and A Chavan. Design Studies for a Financial Management System for Micro-credit groups in Rural India. In *Universal Usability*, pages 15–22, 2003.

[103] Partnership for Higher Education in Africa. Securing the Linchpin: More Bandwidth at Lower Cost. *Carnegie Corporation and the Ford, MacArthur and Rockefeller Foundations*, 2006.

[104] Rabin Patra, Sergiu Nedevschi, Sonesh Surana, Anmol Sheth, Lakshminarayanan Subramanian, and Eric Brewer. WiLDNet: Design and Implementation of High Performance WiFi Based Long Distance Networks. In *NSDI*, 2007.

[105] Rabin Patra, Sonesh Surana, Sergiu Nedevschi, and Eric Brewer. Optimal Scheduling and Power Control for TDMA based Point to Multipoint Wireless Network. In *ACM SIGCOMM Workshop on Networked Systems for Developing Regions (NSDR)*, 2008.

[106] Krishna Paul, Anitha Varghese, Sridhar Iyer, and Bhaskar Ramamurthiand Anurag Kumar. WiFiRe: Rural Area Broadband Access Using the WiFi PHY and a Multisector TDD MAC. *New Directions in Networking Technologies in Emerging Economics, IEEE Communications Magazine*, 2006.

[107] Udai Pawar, Joyojeet Pal, R Gupta, and Kentaro Toyama. Multiple Mice for Retention Tasks in Disadvantaged Schools. In *ACM SIGCHI*, 2007.

[108] C K Prahalad and A Hammond. *Serving the World's Poor, Profitably*. Harvard Business Review, 2002.

[109] One Laptop Per Child Project. `http://www.laptop.org`.

[110] Bhaskaran Raman. Channel Allocation in 802.11-based Mesh Networks. In *IEEE INFOCOM*, April 2006.

[111] Bhaskaran Raman and Kameswari Chebrolu. Revisiting MAC Design for an 802.11-based Mesh Network. In *HotNets-III*, 2004.

[112] Bhaskaran Raman and Kameswari Chebrolu. Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks. In *ACM MOBICOM*, August 2005.

[113] Ram Ramanathan. *Antenna Beamforming and Power Control for Ad Hoc Networks*. Wiley-IEEE Press, 2004.

[114] Ananth Rao and Ion Stoica. An Overlay MAC layer for 802.11 Networks. In *MOBISYS*, June 2005.

[115] Narasimha Puli Reddy. The SRAWAN MAC Protocol to support Real-Time Services in Long Distance 802.11 Networks. Master's thesis, IIT Chennai, 2006.

[116] A Rendon, A Martinez, M Dulcey, J Seoane, R Shoemaker, V Villarroel, D Lopez, and J Simo. Rural Telemedicine Infrastructure and Services in the Department of Cauca, Colombia. *Telemedicine Journal and e-Health*, 2005.

[117] C. G. Riley. The Changing Role of the Citizen in the E-Governance and E-Democracy Equation. *Commonwealth Centre for e-Governance*, 2003.

[118] Luigi Rizzo. Effective Erasure codes for Reliable Computer Communication Protocols. *ACM CCR*, 1997.

[119] J W Robinson and T S Randhawa. Saturation throughput analysis of IEEE 802.11e Enhanced Distributed Coordination Function. *IEEE Journal on Selected Areas in Communications*, 22(5):917–928, 2004.

[120] Maya Rodrig, Charles Reis, Ratul Mahajan, David Wetherall, and John Zahorjan. Measurement-based characterization of 802.11 in a hotspot setting. In *ACM SIGCOMM E-WIND Workshop*, 2005.

[121] Sandra Salmerón, Javier Simó, and Rabin Patra. Comparison of MAC protocols for 802.11 based long distance networks. In *1st Workshop on Wireless Broadband Access for Communities and Rural Developing Regions at International Conference on Mobile Communication Technology for Development (M4D)*, 2008.

[122] L S Satyamurthy and A Bhaskaranarayana. Telemedicine: Indian Space Agency's (ISRO) Initiatives for Speciality Health Care Delivery to Remote and Rural Population. `http://laico.org/v2020resource/files/remote_rural_population.htm`.

[123] Sayandeep Sen and Bhaskaran Raman. Long Distance Wireless Mesh Network Planning: Problem Formulation and Solution. In *WWW*, 2007.

[124] Seattle Wireless Network. `http://www.seattlewireless.net`.

[125] A. Seth, D. Kroeker, M. Zaharia, S. Guo, and S. Keshav. Low-cost Communication for Rural Internet Kiosks using Mechanical Backhaul. In *Mobile Computing and Networking*, pages 334–345, 2006.

[126] Ashish Sharma and Elizabeth M. Belding. FreeMAC: Framework for Multi-channel MAC Development on 802.11 Hardware. In *ACM workshop on Programmable routers for extensible services of tomorrow (PRESTO)*, 2008.

[127] Anmol Sheth. *Improving the performance of wireless systems through distributed fault diagnosis*. PhD thesis, Boulder, CO, USA, 2007. Adviser-Han,, Richard.

[128] Anmol Sheth, Christian Doerr, Dirk Grunwald, Richard Han, and Douglas Sicker. MOJO: A Distributed Physical Layer Anomaly Detection System for 802.11 WLANs. In *MOBISYS*, 2006.

[129] Anmol Sheth, Sergiu Nedevschi, Rabin Patra, Sonesh Surana, Lakshminarayanan Subramanian, and Eric Brewer. Packet Loss Characterization in WiFi-based Long Distance Networks. In *IEEE INFOCOM*, 2007.

[130] J Simó, C Figuera, J Seoane, and A Martínez. Distance Limits in IEEE 802.11 for Rural Networks in Developing Countries. In *IEEE WRECOM*, 2007.

[131] J Simo, P Osuna, D Espinoza, L Camacho, and R Quispe. Application of IEEE 802.11 Technology for Health in Isolated Rural Environments. In *WCIT*, 2006.

[132] Spirent Communications. `http://www.spirentcom.com`.

[133] Peter Stavroulakis. Interference suppression techniques: A Twenty-year Survey. *International Journal of Satellite Communications and Networking*, 2003.

[134] Lakshminarayanan Subramanian, Ion Stoica, Hari Balakrishnan, and Randy Katz. OverQoS: An Overlay Based Architecture for Enhancing Internet QoS. In *NSDI*, March 2004.

[135] Lakshminarayanan Subramanian, Sonesh Surana, Rabin Patra, Melissa Ho, Anmol Sheth, and Eric Brewer. Rethinking Wireless for the Developing World. In *Hotnets-V*, 2006.

[136] Chen Sun, Student Member, Akifumi Hirata, Takashi Ohira, Nemai C. Karmakar, and Senior Member. Fast Beamforming of Electronically Steerable Parasitic Array Radiator Antennas: Theory and Experiment. *IEEE Transactions on Antennas and Propagation*, 52:1819–1832, 2004.

[137] Sonesh Surana, Rabin Patra, Sergiu Nedevschi, and Eric Brewer. Deploying a Rural Wireless Telemedicine System: Experiences in Sustainability. *IEEE Computer*, 41(6):48–56, June 2008.

[138] Sonesh Surana, Rabin Patra, Sergiu Nedevschi, Manuel Ramos, Lakshminarayanan Subramanian, Yahel Ben-David, and Eric Brewer. Beyond Pilots: Keeping Rural Wireless Netorks Alive. In *NSDI*, 2008.

[139] Godfrey Tan and John Guttag. Time-based Fairness Improves Performance in Multi-rate WLANs. In *USENIX*, 2004.

[140] David Tse and Pramod Viswanath. *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.

[141] Teresa Sheausan Tung. *QoS in Wireless Networks*. PhD thesis, University of California at Berkeley, November 2006.

[142] Ubiquiti Wireless. `http://www.ubnt.com/`.

[143] United Nations. UN Millennium Development Goals. `http://www.un.org/millenniumgoals/`.

[144] David M. Upton and Virginia A. Fuller. The ICT e-Choupal Initiative. *Harvard Business School Case N9-604-016*, January 2004.

[145] Vihaan Networks Limited. WorldGSM Rural Deployment. `http://www.vnl.in/solutions/rural/`.

[146] Mythili Vutukuru, Kyle Jamieson, and Hari Balakrishnan. Harnessing Exposed Terminals in Wireless Networks. In *NSDI*, 2008.

[147] R. Wang, U. Sahni, S. Sobti, N. Garg, J. P. Singh, M. Kam, A. Krishnamurthy, and T. Anderson. The Digital StudyHall. Technical Report TR-723-05, Computer Science Department, Princeton University, 2005.

[148] U Wehn and R Mansell. *Knowledge Societies*. Oxford University Press, New York on behalf of the United Nations, 1998.

[149] Wi-LAN Incorporated. Wi-LAN Fixed and Mobile Solutions. `http://www.wi-lan.com`.

[150] Wizzy Digital Courier. `http://wizzy.org.za`.

[151] World Bank. The Information for Develoment Program (infoDev). `http://www.infodev.org/`.

[152] World Bank. World Development Indicators, 2008. `http://www.worldbank.org/data`.

[153] Y Xiao. Performance analysis of priority schemes for IEEE 802.11 and IEEE 802.11 e wireless LANs. *IEEE Transactions on Wireless Communications*, 4(4):1506–1515, 2005.

[154] Liang-Liang Xie and P. R. Kumar. A network information theory for wireless communication: Scaling laws and optimal operation. In *IEEE Transactions on Information Theory, Vol 50, No. 5, 2004*.