© Copyright by Romit Roy Choudhury, 2006

UTILIZING BEAMFORMING ANTENNAS FOR WIRELESS MULTIHOP NETWORKS

BY

ROMIT ROY CHOUDHURY

B.S., Haldia Institute of Technology, 2000
M.S., University of Illinois at Urbana-Champaign, 2003

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Computer Science
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2006

Urbana, Illinois

UTILIZING BEAMFORMING ANTENNAS FOR WIRELESS MULTIHOP NETWORKS

Romit Roy Choudhury, Ph.D.
Department of Computer Science
University of Illinois at Urbana-Champaign, 2006
Nitin H. Vaidya, Adviser

Wireless multihop networks are gaining popularity in applications such as mesh networks, wireless backbones, and sensor networks. Typically, communication protocols designed for such networks have explicitly or implicitly assumed an omnidirectional antenna at the radio layer. With recent progress in signal processing and antenna technologies, smart beamforming antennas have become feasible at cheaper prices, making them an attractive replacement to omnidirectional antennas. With beamforming antennas, the ability to guide RF energy in desired directions can lead to higher spatial reuse of the channel. Also, the higher communication range of beamforming antennas can be exploited to achieve stronger network connectivity, and to obtain fewer hop routes. However, existing protocols considered suitable for omnidirectional antennas are incapable of achieving these benefits. We have identified several weaknesses in these protocols, that together degrade overall network performance. We have addressed these weaknesses, and developed a better understanding of the improvements achievable with beamforming antennas. Based on the insights, we have designed medium access control (MAC) and routing protocols to suitably exploit the benefits of beamforming. We have demonstrated part of our solutions on a small prototype testbed, equipped with electronically steerable antennas.

Abstract

Wireless multihop networks are gaining popularity in applications such as mesh networks, wireless backbones, and sensor networks. Typically, communication protocols designed for such networks have explicitly or implicitly assumed an omnidirectional antenna at the radio layer. With recent progress in signal processing and antenna technologies, smart beamforming antennas have become feasible at cheaper prices, making them an attractive replacement to omnidirectional antennas. With beamforming antennas, the ability to guide RF energy in desired directions can lead to higher spatial reuse of the channel. Also, the higher communication range of beamforming antennas can be exploited to achieve stronger network connectivity, and to obtain fewer hop routes. However, existing protocols considered suitable for omnidirectional antennas are incapable of achieving these benefits. We have identified several weaknesses in these protocols, that together degrade overall network performance. We have addressed these weaknesses, and developed a better understanding of the improvements achievable with beamforming antennas. Based on the insights, we have designed medium access control (MAC) and routing protocols to suitably exploit the benefits of beamforming. We have demonstrated part of our solutions on a small prototype testbed, equipped with electronically steerable antennas.

To Shyamal Roy Chowdhury

Acknowledgments

This thesis would not have been possible without the support and encouragement of my adviser, Professor Nitin H. Vaidya. I am deeply grateful for his guidance all throughout my Ph.D. program, especially in times of distress, dissapointment, and frustration. I am thankful to him for all his time, efforts, and most importantly, for his patience during my formative years as a research student. I am also thankful to the members of my Ph.D thesis committee – Professor P. R. Kumar, Professor Jennifer Hou, and Professor Klara Nahrstedt. Their constructive criticism have been invaluable toward the completion of the thesis. I would also like to thank Dr. Ram Ramanathan at BBN Technologies for all the confidence he has shown in me over the duration of five years, Professor Somprakash Bandyopadhyay at Indian Institute of Management, Kolkata, for mentoring me whenever I was in need, and Professor Indranil Gupta at UIUC for the insightful discussions I have had with him. While the list can grow far longer, I would like to mention all my friends and colleagues, including Vartika Bhandari, Chandrakanth Chereddi, Shravan Gaonkar, Pradeep Kyasanur, Matthew Miller, Jungmin So, Xue Yang for their technical and non-technical support during my studies. Furthermore, I would like to thank National Science Foundation, DARPA, Motorola Center for Communications Fellowship, and Vodafone-US Foundation Graduate Fellowship for providing financial support.

There are many others without whom I would not have come this far. My father, Pradip Roy Choudhuri, my mother, Mala Roy Chowdhury, and my brother, Shamik Roy Chowdhury, have stood beside my through the thick and thin of my life. I wish to thank them for everything. My sincere gratitude goes to my uncle, Shyamal Roy Chowdhury, for kindling in me the academic curiosity, and for motivating me in every possible way. His contribution has been ineffable. I would also take this opportunity to thank all my professors in the B.Tech program in Haldia Institute of technology, especially Professor Niloy Ganguly who continues to be an inspiration to me.

Finally, I am deeply indebted to my fiancee, Abrita Chakravarty, for being beside me through the best and worst of times. I cannot imagine how it might have been without her presence. She has been my partner in every possible way during my long Ph.D journey.

Table of Contents

Chapter 1 INTRODUCTION **1**
1.1 Background 1
1.2 The Case for Beamforming 2
1.3 Research Challenges 3
1.4 Research Contribution and Thesis Outline 4

Chapter 2 ANTENNA MODEL **7**
2.1 Single Beam Antennas 7
2.2 Multi-Beam Antennas 10

Chapter 3 IMPACT OF BEAMFORMING – PROBLEM FORMULATION . . **12**
3.1 Directional Medium Access Control 12
3.1.1 Related Work 13
3.1.2 Basic Directional MAC Protocol (DMAC) 17
3.1.3 Problems With Basic DMAC 19
3.1.4 Evaluating Basic DMAC 26
3.2 Routing with Beamforming Antennas 31
3.2.1 Related Work 32
3.2.2 Basic Directional Dynamic Source Routing (DDSR) 35
3.2.3 Evaluating Basic DDSR 37

Chapter 4 MAC PROTOCOLS USING BEAMFORMING ANTENNAS **45**
4.1 Deafness 46
4.1.1 Related Work 46
4.1.2 Impact of Deafness 51
4.1.3 ToneDMAC – Protocol Description 55
4.2 MAC-Layer Capture 66
4.2.1 Problem Formulation 67
4.2.2 Capture-Aware Directional MAC (CaDMAC) 73

Chapter 5 ROUTING PROTOCOLS USING BEAMFORMING ANTENNAS **84**
5.1 Directional DSR (DDSR) 85
5.1.1 Improving DDSR 85
5.2 Capture-Awareness in Routing 92
5.2.1 Capture-Aware Routing Protocol (CaRP) 94

Chapter 6 EXPERIMENTAL EVALUATION **105**
6.1 System Description 105
6.2 Related Work 106
6.3 Experiments 108
6.3.1 Neighbor Discovery 108
6.3.2 MAC Layer: IEEE 802.11 and ESPAR 110
6.3.3 Link State Routing 115
6.3.4 Discussion and Summary 115

Chapter 7 SUMMARY AND FUTURE WORK **117**

References **120**

Author's Biography **127**

List of Tables

3.1 DMAC affected by deafness on a 2-hop flow 26

4.1 Aggregate TCP multihop throughput 65

List of Figures

1.1 (a) ESPAR antenna prototypes at 2.4 GHz. (b) Comparison of ESPAR antenna prototype at 2.4GHz and 5 GHz, the latter being placed on the right of the laptop. . 5

2.1 A realistic radiation pattern including side lobes. 9
2.2 Switched combining with multi-beam antennas 10
2.3 An abstract model of multi-beam radiation patterns 11

3.1 An example topology scenario 20
3.2 An example where node E wishes to send a packet to H while it has a DNAV entry for direction D^{EB}. The Figure shows a condition where node E can safely transmit to node H, since it would not interfere with node B's transmission 20
3.3 Hidden terminal problems arising when nodes change antenna patterns for communication. 22
3.4 An example scenario depicting the possibility of deafness 24
3.5 End to end throughput on a chain topology 27
3.6 Example scenarios showing the decreasing degree of alignment in the choice of routes. 27
3.7 Aggregate throughput for varying degree of route alignment 28
3.8 An example of a random topology. 30
3.9 Average aggregate throughput and average end to end delay for random topologies with random routes. 31
3.10 Performance improvement for chain and grid topologies, when routing using beamforming antennas. 39
3.11 Aggregate throughput and average delay in random topologies with directional routing. 39
3.12 DDSR's ability to connect holes in the network. 40
3.13 DDSR shows better performance for the topology on the right. 41
3.14 (a) Mobility affects performance due to frequent link breakages. (b) Performance varies with variation in beamwidth. 43

4.1 A scenario illustrating the problem of deafness 47
4.2 An example scenario in multi-hop ad hoc networks. 51
4.3 Omnidirectional transmissions affect spatial reuse. 53
4.4 An example scenario in multi-hop ad hoc networks. 53
4.5 A scenario illustrating channel under-utilization due to deafness 55
4.6 An example scenario and timeline of the ToneDMAC protocol. 57
4.7 Scenarios affected by deafness. 59
4.8 Throughput comparison for scenario (i). 60
4.9 Variation of backoff values for scenario (ii) using DMAC. 61
4.10 Variation of backoff values for scenario (ii) using ToneDMAC. 61

4.11 Comparison of delay for scenario (ii). 62
4.12 Packet drop due to deafness in DMAC and ToneDMAC 62
4.13 Performance of ToneDMAC in chain and grid topologies 63
4.14 Throughput in a random network for increasing traffic. 64
4.15 Throughput comparison for different random networks. 64
4.16 Capture prevents parallel dialog between A-B and C-D 69
4.17 Time divided into cycles of ON and OFF-Durations 73
4.18 CaDMAC enables parallel communication by avoiding mutual capture between nearby nodes. 75
4.19 Communicating with beam β_{off} is inefficient (nodes B and C capture each other because of higher interference). 75
4.20 Maximal spatial reuse available from CaDMAC 76
4.21 A simple scenario with 2 flows; A-B-C and E-F-G. 79
4.22 CaDMAC improves throughput for the simple scenario 80
4.23 CaDMAC improves throughput for grid topologies with zigzag flows. 80
4.24 Throughput benefits from capture-aware MAC protocols in large random networks, with beamwidth of 60° 81
4.25 Reduction in end to end latency with capture-aware MAC protocols in random networks. 82

5.1 Partial sweeping optimization with directional DSR 88
5.2 Performance of DDSR on static random networks. 90
5.3 Optimized DDSR degrades more gracefully in comparison to DDSR. 92
5.4 Communication on link A-B captures link M-F. Routes that avoid using link M-F (and chooses link M-E instead) extracts greater spatial reuse from the channel . . . 93
5.5 Accumulation of route cost in the propagating RREQ 98
5.6 Enumeration of routes when using DSR and our capture-aware routing protocol, on a small topology with 4 ordered flows. 99
5.7 Capture-aware routes, chosen by CaRP, offers higher aggregate throughput. 100
5.8 Capture-aware routing in large random networks. 101
5.9 Impact of ON and OFF durations on performance. 102
5.10 Impact of varying weights on the performance of CaRP. 103

6.1 ESPAR antenna prototype (for 2.4 GHz band). 106
6.2 Block diagram showing interaction between laptop and ESPAR. 106
6.3 Each dot results from one sweep – the angular position of the dot represents the directional beam that was selected in that sweep, and the radial distance is proportional to max received signal strength for that particular sweep. The direction of black dots represents the LoS direction. The graph plots results for 1 hour. 109
6.4 Beam selection improves in the anechoic chamber. 110
6.5 Neighbor directions for a plot for 30 seconds. 110
6.6 Node configurations to exhibit spatial reuse. 111
6.7 Comparing spatial reuse for OO and DD flows. 112
6.8 Validating capture in a toy topology. 112
6.9 Performance when capture is prevented. 113
6.10 Packet drops rates at transmitters due to the impact of deafness 114
6.11 Higher packet delivery ratio with increasing packet sizes 114

Chapter 1

INTRODUCTION

1.1 Background

Recent advances in electronic devices and wireless communication have enabled a variety of networking technologies, including wireless local area networks (WLAN), mesh networks, ad hoc networks, sensor networks, and RFID networks. In all these networks, a set of computing devices are interconnected over the wireless medium to form a distributed environment. The ability to communicate without the need for expensive infrastructure has made wireless technology the key to *anytime, anywhere computing.* Envisioned applications include military coordination, disaster relief operations, wireless backbones for rural connectivity, environment sensing, monitoring, tracking, etc. Motivated by emerging applications, the research community has developed, and continues to develop, various protocols for supporting the increasing wireless need. A large fraction of these protocols have been focused on extracting the highest levels of performance using the available spectrum. Such protocols have investigated various dimensions over which improvements may be possible, including frequency, space, time, code, etc. This dissertation aims to exploit spatial characteristics of wireless multihop networks through the use of beamforming antennas. We intend to design protocols at multiple layers of the network stack that together can harness the potentials of directional communication.

1.2 The Case for Beamforming

As opposed to omnidirectional antennas, beamforming antennas posses the ability to radiate energy in a desired direction. When communication ranges are made equal, transmissions using beamforming antennas interfere over a smaller area around the vicinity of the transmitter. This increases the spatial reuse of the network since more simultaneous communications may now be packed within a given area. Packing more simultaneous communications in space increases channel utilization, leading to greater performance.

Focusing more energy in a desired direction may also be useful in achieving better link stability – a link formed by using beamforming antennas may achieve higher SINR, and therefore, might be less prone to disruption from channel fluctuations and fading. The possibility of higher SINR can increase transmission rates, which in turn enhances network performance.

Beamforming antennas can also be used to attain higher communication range in a desired direction. A higher communication range can be useful in finding fewer-hop routes between a source and a destination in a wireless multi-hop network. Fewer-hop routes can reduce the latency of communication, and in some cases, may improve the throughput of the network.

The reception characteristics of beamforming antennas are reciprocal to transmission. Using electronic beamforming, these antennas can amplify signals from specific directions in space, and suppress signals from other directions. This can be envisioned as a spatial filter – like frequency filters that filter out signals in the frequency domain, beamforming antennas are capable of filtering signals on the spatial domain. Exploiting spatial filtering in wireless networks leads to the possibility of avoiding "capture" with beamforming antennas. Briefly, "capture" is a problem in which a wireless device remains engaged in receiving a packet that is not meant for it. A unicast transmission using IEEE 802.11, although intended for a single destination node, can result in *capturing* several other nodes located within the interference footprint of the transmitter. Antennas capable of beamforming may be able to overcome capture by spatially filtering out directions from which signals lead to capture. As a result, nodes may gain more free time to simultaneously schedule their

own productive communication. We later visit the problem of capture, and show how beamformed reception can mitigate capture, and thereby improve the performance of wireless multihop networks.

Several other benefits motivate the case for exploiting beamforming antennas. However, the benefits are not free of tradeoffs, and therefore the problem of utilizing beamforming antennas is not trivial. Some of the challenges with beamforming communication are discussed next.

1.3 Research Challenges

To be able to communicate directionally, a transmitter first needs to beamform in the (appropriate) direction of its intended receiver. Determining appropriate directions of neighbors is a non-trivial problem, especially if nodes are mobile, or if users rotate on a vertical axis. Moreover, in indoor environments, multipath effects can be pronounced, disqualifying the possibility of line-of-sight beamforming. As a result, neighbor discovery and neighbor tracking needs to be performed continuously – a node may require to handoff from one beam to the other as its neighbor moves, or as the signal environment changes with time. *Directional neighbor discovery*, unlike with omnidirectional antennas, is an important research challenge.

The ability to filter signals from specific directions, although effective for achieving higher SINR and avoiding capture, leads to a serious problem called *deafness*. Briefly, deafness is caused when a node, C, attempts to initiate a dialog with a node, A, while A is engaged in communication with another node, B. Node A, by virtue of being beamformed toward B, fails to receive signals from C, and thus cannot reply to C. Node C fails to classify the reason for communication failure, because both deafness and collision could be the potential cause for such failures. The impact of deafness can significantly affect performance of the network, as discussed in later chapters of this dissertation.

Directional communications also lead to new kind of *hidden terminal problems*. Due to the differences in gains of various beam-patterns, a transmitter may cause collisions at the receiver of an ongoing dialog. Exchange of RTS/CTS control packets, although widely used in literature to

cope with hidden terminal problems, prove to be insufficient for directional communication. We visit these problems in later chapters, and analyze their impact on wireless networking.

From the perspective of a network layer, beamforming antennas face the problem of achieving broadcast communication. Several higher layer protocols require a node to broadcast a packet to all its neighbors. When using beamforming antennas, a *broadcast to all neighbors* has to be emulated by transmitting copies of the same packet in all directions. This incurs a high overhead, and can increase the delay of message propagation among nodes. More subtle problems arise, making the problem of utilizing beamforming antennas an interesting research challenge to the wireless networking community.

1.4 Research Contribution and Thesis Outline

This dissertation focuses on utilizing beamforming antennas for wireless multihop networks, with an emphasis on MAC and routing protocols. At the MAC layer, we confine ourselves to the regime of CSMA/CA protocols, and investigate the issues underlying directional channel access. We begin with a study of a simple directional MAC protocol (a directional version of IEEE 802.11), and use it to identify several problems that appear with beamformed communication. These problems are noticeably different from those that appear with omnidirectional antennas, and thereby do not lend themselves to existing solutions. We address these problems systematically in an attempt to fully exploit the capabilities of beamforming antennas. At the MAC layer, our contributions include extensions to existing protocols (IEEE 802.11 is extended to DMAC), as well as the design of new MAC protocols (ToneDMAC, CaDMAC) for beamforming antennas.

Building on our MAC layer solutions, we evaluate the impact of beamforming antennas on the network layer. For this, we modify an omnidirectional routing protocol, named Dynamic Source Routing (DSR), and study its impact when executed with a directional MAC layer – we refer to this as *directional DSR (DDSR)*. Using insights from these studies, we design new routing solutions that aid our MAC protocols in increasing the spatial reuse of the network. More specifically, we

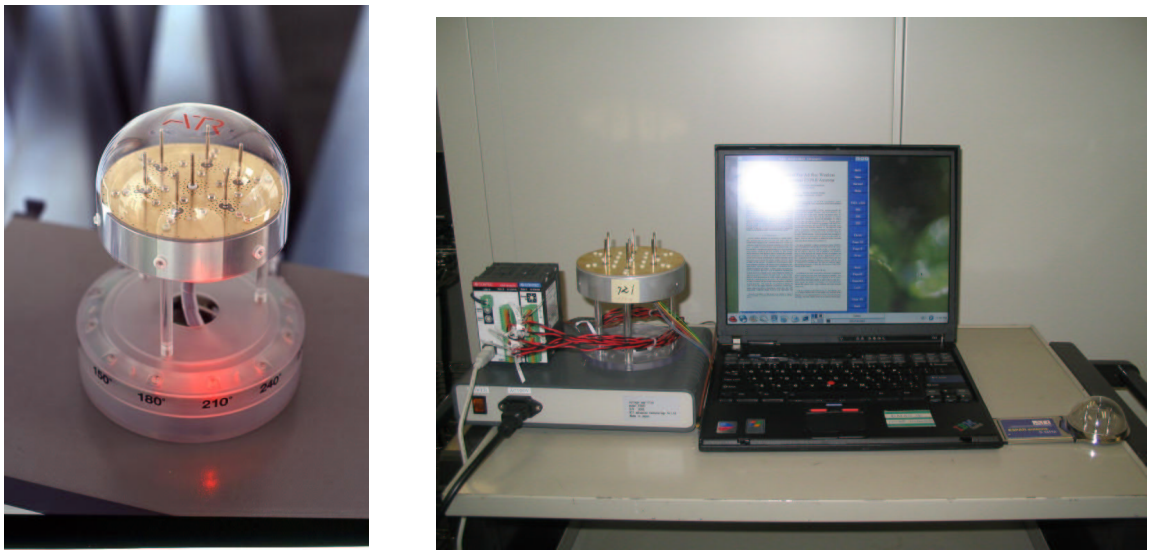

(a) (b)

Figure 1.1: (a) ESPAR antenna prototypes at 2.4 GHz. (b) Comparison of ESPAR antenna prototype at 2.4GHz and 5 GHz, the latter being placed on the right of the laptop.

incorporate the notion of capture-awareness in routing so that chosen routes do not share the channel between them. We show how the combination of our MAC and routing protocols can achieve higher performance than the individual protocols. To attain further benefits, we also investigate cross-layer optimizations available from the inherent diversity of beamforming antennas.

We evaluate our protocol through simulation and a small prototype implementation. For our simulations, we use the Qualnet simulator, version 3.6, and experiment with a varied set of scenarios. As a part of our prototype implementation, we develop a testbed of 6 laptops, each equipped with *Electronically Steerable Passive Array Radiator (ESPAR)* antennas (shown in Figure 1.1). As a proof of concept, we validate some of the key benefits of beamformed communication, and quantify some of the tradeoffs, such as deafness. We also study the negative impacts of multipath and channel fluctuation on the performance of neighbor discovery protocols. Our overall evaluation suggests that when controlled carefully, beamforming antennas can offer encouraging performance benefits in wireless multihop networks.

The rest of this dissertation is organized as follows. Chapter 2 introduces the antenna models that we use over the remainder of the dissertation. Chapter 3 presents the impact of beamforming on higher layers, formulating the need for revision and re-design of MAC and routing protocols. In Chapter 4 we design MAC protocols to address some of the key challenges with beamforming, namely *deafness* and *capture*. In Chapter 5 we present routing protocols that better utilize beamforming antennas – we show that the benefits can be significant when used in conjunction with our proposed MAC protocols. Chapter 6 discusses our practical experiences from a prototype testbed equipped with beamforming antennas, while Chapter 7 discusses other related work in this area of active research. Finally, Chapter 7 concludes the dissertation and identifies future research directions.

Chapter 2

ANTENNA MODEL

We present our assumed antenna models in this chapter. Some precise details of real antenna models are abstracted for the convenience of explanation. However, the actual antenna model used for simulations are obtained from existing beamforming antennas, and are characterized by side lobes and back lobes – one example radiation pattern has been shown in Figure 2.1. We begin with a description of single beam antennas that we use in the initial parts of our dissertation. Then we describe multi-beam antenna models used for our sbsequent proposals on capture-aware protocols.

2.1 Single Beam Antennas

Let us consider a single beam antenna system. The antenna system comprises of N non-overlapping beams such that the union of all the beams covers the entire ($360°$) azimuth plane. The shapes of the beams can be envisioned as sectors, each sector having a beamwidth of $\frac{360}{N}$. A main lobe points in the direction of the sector, while side lobes span over the remaining azimuthal directions. When transmitting with such an antenna beam, more energy is radiated in the direction of the main lobe, while some energy leaks through the side lobes. The transmit and leakage energy are proportional to the main lobe and side lobe gains, respectively. By virtue of reciprocity, reception using such a beam will enable the antenna to amplify signals over the main lobe, and suppress signals that impinge on the side lobes.

The antenna system offers two modes of operation: *Omni* and *Directional*[1]. A node resides in

[1]The *Omni* mode may be imagined as an auxiliary antenna, with an omnidirectional radiation pattern.

the *Omni* mode when it does not know the direction from which a signal might arrive. Once a signal is detected, the antenna system invokes a beamforming algorithm to estimate the beam-pattern on which the impinging signal can be received with highest quality. Various beamforming algorithms have been designed at the physical layer, and depending on their characteristics, they take different time-duration to determine the most appropriate beam-pattern. In the simplest form, a beamforming antenna will (electronically) switch between a finite number of pre-defined beam-patterns, recording the SINR achieved from each of them – called an azimuthal scan. Finally the antenna selects the beam-pattern over which the detected SINR is maximum. This type of antenna systems are referred to as *Switched Beam Antennas*, in view of its ability to switch between a pre-defined set of beams. In contrast to Switched Beam Antennas, *Steerable Beam Antennas* (also called adaptive antennas) employ sophisticated signal-processing algorithms to effectively locate and track the signal of interest. Moreover, they dynamically minimize the interference in the environment by steering nulls toward these directions. While switched and steerable antennas both increase the communication gain, steerable antennas can achieve optimality in a dynamically changing signal environment. For both types of antennas, once a signal is received with a particular beam-pattern, the antenna can cache the direction of arrival (DoA) of this signal, and use it in the near future to communicate back to the transmitter. While we assume switched beam antennas for our simulations, our protocols can be easily used in conjunction with steerable antennas.

When in the omni mode, the gain of the antenna is denoted as G^o. However, when directional beams are used for communication(transmission or reception), the main lobe gain in the direction of the beam is denoted as G^d. Typically, $G^d \geq G^o$. The Friss Equation represents how transmit and receive gains (G_T and G_R) are related to the transmit and receive powers (P_T and P_R) [1]:

$$P_R = \frac{P_T G_T G_R}{K r^\alpha}$$

The term K is a constant that accounts for atmospheric absorption, ohmic losses, etc., r is the distance between the transmitter and the receiver, and α is the path-loss index. Note that G_T and G_R are transmit and receive gains *along the straight line joining the transmitter and receiver.*

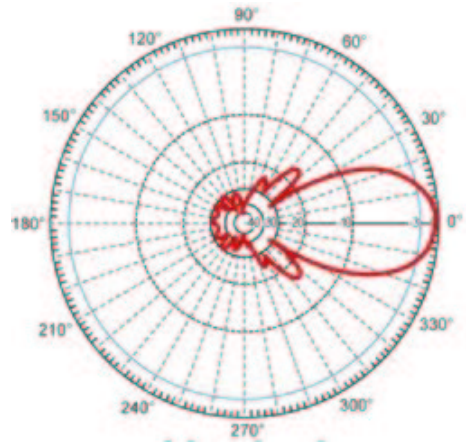

Figure 2.1: A realistic radiation pattern including side lobes.

Observe from the equation that the maximum distance, r, over which two nodes can communicate, increases with increase in transmit and the receive gains. As a result, directional antennas offer the capability of range extension. Put differently, two nodes in omni mode may be out of communication range because the product of their omnidirectional transmit and receive gains, $(G^o \times G^o)$, is not large enough. However, if one of the nodes beamforms in the direction of the other, the new product, $(G^d \times G^o)$, may be sufficiently large to enable direct communication. To distinguish between types of communicating node pairs, we use *OO-neighbors* to refer to node-pairs that can communicate while both nodes are in omni mode. *DO-neighbors* refer to nodes that can communicate only when one or both the nodes have beamformed at its counterpart. *DD-neighbors* refer to nodes that can communicate only when both the nodes have beamformed toward each other.

Relating to IEEE 802.11, a transmitter can transmit the RTS directionally, assuming that it is aware of the direction of its intended receiver. Since the receiver does not know *a priori* who the transmitter might be, it cannot receive the RTS directionally. The receiver receives the RTS omnidirectionally and performs the azimuthal scan (in parallel) to determine the best beam for communicating to this transmitter. Thereafter, the CTS, Data, and ACK packets can be transmitted and received directionally. However, since we assume that the RTS must be received omnidirectionally, the distance between two communicating neighbors is bounded by $G^d \times G^o$.

2.2 Multi-Beam Antennas

We propose to use multi-beam antennas in the latter part of our dissertation. Multi-beam antennas can be abstracted in a similar way as single-beam antennas – we present the abstraction in this section. Similar to the single beam antenna systems, multi-beam antenna systems also comprise of N non-overlapping beams that together cover the entire (360°) azimuth plane. However, each antenna beam has an on-off switch, and signals from each of the beams are combined equally to feed a single transceiver. Using these switches, it is possible to turn on/off any subset of beams, resulting in various multi-beam patterns. A block diagram of the antenna system is presented in Figure 2.2, and example beam-patterns are shown in Figure 2.3. In Figure 2.3(b), a node only has a single beam turned on towards north. In all other directions, it has pointed a null implying that signals arriving in these directions are suppressed. In reality (and in beam-patterns used in our performance evaluation), nulls include low-gain side lobes and backlobes, as illustrated in Figure 2.1.

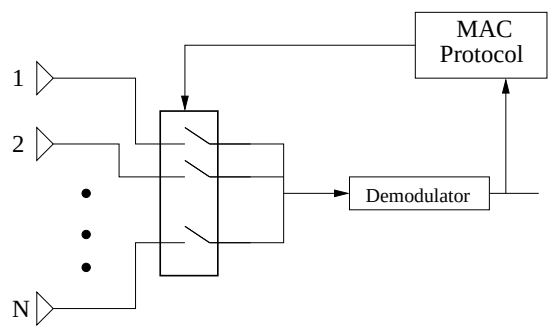

Figure 2.2: Switched combining with multi-beam antennas

When all the beams of the antenna are turned on (Figure 2.3(a)), the beam-pattern is "omnidirectional". The antenna is then said to be in the *omnidirectional mode*. When transmitting in the omnidirectional mode, the antenna uses a fixed power P_{max}. When transmitting using a subset of its beams, the transmit power will be shared only by subset of beams. As a result, the communication range increases. We will discuss the impact of higher communication range on the performance of wireless networks, and regulate the transmission power whenever we desire to maintain an uniform communication range. Equalizing the communication range ensures that the reachable neighbors of a node do not change with change in beam-patterns.

We assume that multi-beam antennas have the ability to detect the beam-of-arrival of a signal. In other words, when a signal arrives, the output from each of the beams can be compared to detect the beam on which the received signal strength was maximum. This information is made available to the higher layers, along with the received packet.

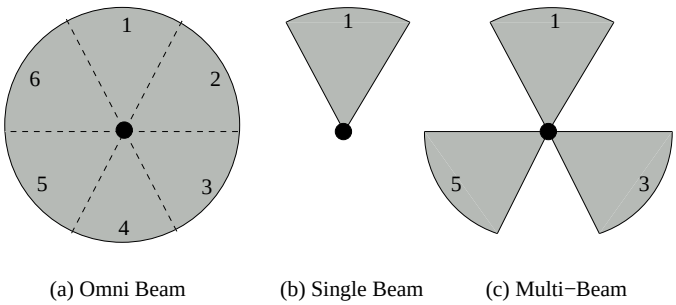

Figure 2.3: An abstract model of multi-beam radiation patterns

As mentioned earlier, signals that impinge on the mainlobe of an active beam, are received with a higher gain, while those that arrive from other directions are suppressed by the lower gain of the antenna's side lobes. One may view such a beam pattern as a spatial filter (a well-understood concept in the area of adaptive beamforming). Like frequency filters that filter out signals on the frequency domain, spatial filters can be made to filter signals in space. The abstract antenna model presented above can be realized in many ways, including diversity antennas and adaptive beamforming antennas [2]. Of course, such realizations will have beamforms with side lobes as appropriately considered in our dissertation.

Chapter 3

IMPACT OF BEAMFORMING – PROBLEM FORMULATION

Advancements in signal processing and antenna design principles have led to the design of beamforming antenna systems. While such antennas have been widely used in cellular networks, their applicability in wireless multi-hop networks have not been adequately studied in the past. We aim to investigate the utility of beamforming antennas in such multi-hop environments. We aim to understand if replacing omnidirectional antennas, with the newer beamforming antennas, is sufficient in achieving the complete benefits from beamforming. To address this question, we begin by designing very simple protocols at the MAC and network layer, and evaluate their performance when used with beamforming antennas. This chapter presents the results of our evaluation, and identifies multiple challenges that underlie beamformed communication in a multi-hop network. These challenges form the basis of our dissertation, and motivates the design of new protocols presented in subsequent chapters. We begin with directional medium access control, and discuss network layer issues thereafter.

3.1 Directional Medium Access Control

This section studies the impact of beamforming antennas on the performance of the MAC layer. We confine ourselves to the regime of carrier sense multiple access/collision avoidance (CSMA/CA) protocols. We consider IEEE 802.11, the *de facto* MAC layer standard for current wireless networks, and extend it to a simple directional version, called directional MAC (DMAC). If the expected benefits from beamforming can be achieved through DMAC, then research on new MAC protocols

can be deemed unnecessary. However, if DMAC is unable to leverage full benefits of beamforming, then MAC protocols need to be designed specifically for beamforming antennas. Motivated from this, we present the details of DMAC, and evaluate it over a wide range of scenarios. We show how DMAC reveals new issues with directional medium access control, and exposes potential problems and tradeoffs. We study these problems and tradeoffs and demonstrate the need for new protocols. We begin with a discussion of related work in this area, followed by a brief description of the IEEE 802.11 standard [3].

3.1.1 Related Work

For several decades in the past, beamforming antennas have been widely utilized in the context of wireless systems, such as cellular networks [4], [5], [6], broadband systems [7], [8], and satellite networks [9]. More recently, some researchers have considered the utility of beamforming antennas in the design of wireless multihop networks [10] [11] [12] [13]. However, the applicability of beamforming antennas in wireless multihop networks involves a wide range of issues at multiple layers of the network stack, much of which has not been fully investigated. We review the contributions of existing work and distinguish our research from them.

The design of IEEE 802.11 implicitly assumes a static omnidirectional antenna at the physical layer. Although 802.11 may operate correctly when using fixed directional antennas, beamforming antennas that change their beams over time can cause 802.11 to operate incorrectly. Moreover, performance can get seriously affected unless the MAC protocols are made antenna-aware. Recently, several MAC protocols have been proposed that suitably adapt 802.11 for beamforming antennas. Ko *et al.* [14] make the simple but useful observation that directional antennas can be used to improve the number of concurrent communications in a network (spatial reuse). The authors propose a simple MAC protocol capable of transmitting directionally, and receiving omnidirectionally. The authors have proposed to transmit an RTS directionally, only if the RTS does not collide with other ongoing communications. However, when the direction of the intended receiver is not known, the MAC protocol reverts to transmitting an omnidirectional RTSs. Results show that improvements are possible when exploiting such concurrency in the channel. Nasipuri *et al.* [15] also propose

to reduce the interference in the wireless channel by communicating directionally. However, their proposals require the transmission of an omnidirectional CTS to inform the receiver's neighborhood about the imminent dialog. This offsets spatial reuse – a key advantage of using directional antennas.

In [16], Elbatt *et al.* propose an interesting idea – they use RTS/CTS to inform the neighborhood about the beam indices to be used for the imminent communication. Based on this information, neighbors of the communicating nodes decide which beams may be used for initiating their own RTS packets. In other words, neighbors of a communicating node can now refrain from using specific beams in order to avoid interfering with the ongoing communication. Bandyopadhyay *et al.* [17] present another MAC protocol that informs neighborhood nodes about ongoing communications through additional control messages. They propose to send proactive messages in the neighborhood based on the active flows in the network. To achieve this, the protocol assumes knowledge of network traffic.

Takai *et al.* have described the DVCS mechanism in [18], where a node performs directional virtual carrier sensing. Directional virtual carrier sensing is a useful optimization that increases the spatial reuse of the channel because nodes are required to defer transmission only in the direction in which another communication is in progress. We have independently proposed the notion of virtual carrier sensing as discussed later in this chapter.

Ramanathan [19] presents an interesting discussion on several issues arising from directional communication. [19] discusses the possibilities for taking advantage of higher transmission range of beamforming antennas. For example, the authors discuss the challenges in directional neighbor discovery in which communicating to a DD-neighbor proves to be a difficult problem. The author also discuss the challenges with power control, and how they might impact MAC and network layer protocols. The paper motivates comprehensive research in the field of utilizing directional antennas for wireless multihop networks. Streenstrup [20] visits the problem of neighbor discovery using directional antennas. Three key ideas have been proposed in this context. The basic idea

proposes that nodes choose a mode (scanning or listening) based on their identifiers. Once a scan has been performed, half the nodes (those in listening mode) would be expected to have discovered the other half (those in transmitting mode). In the next round, the similar choice of modes would lead to further discoveries among those that were in the same mode in the first round – leading to an average of log(N) scans for discovering all nodes on average. The second idea proposes a stochastic algorithm, whereby each node chooses its mode according to an uniform distribution. The paper shows that the average case performance of the stochastic algorithm can be better than the deterministic algorithm. The paper alludes to a *rendezvous* scheme in which far away nodes may schedule to discover each other in future, via multiple hop communications (the concept of DD communications using DO hops). This dissertation aims to broaden the scope of earlier research to a generalized set of antennas, that we call beamforming antennas. We aim to study the impact of beamforming antennas on wireless multihop networks, and design solutions that exploit these antennas to their full potentials.

The capacity of wireless networks with directional antennas have been studied analytically. [21] presents results on the lower bounds of capacity improvement (over Gupta and Kumar model) when using directional antennas. They show that improvements of $O(\sqrt{(}2\pi/\theta))$ can be achieved when either the transmission or the reception is directional. When both transmission and reception is directional, the authors show that the improvements can be $O(2\pi/\sqrt{(}\theta\beta))$, where θ and β are the beamwidths of transmit and receive beams. The intuition in the authors argument is that when using directional beams, the interference footprint reduces with beamwidth, and therefore the capacity improvements are of the order of square root of this reduction. The interference model ignores side lobes for the purposes of simplification.

In [22], authors study spatial reuse of antennas by assuming a directional antenna model (including side lobes). The authors argue that in certain directions the interference impinging on the receiver is larger (due to the mainlobe of the receiver pointing in these directions), while from other directions, the impinging interference is lower. Appropriately weighing the interfering signals, the authors determine the effective interference at each node, and show that the net reduction in

interference leads to higher spatial reuse and capacity. The authors also claim that, contrary to intuition, narrower mainlobes lead only to marginal benefits.

IEEE 802.11 DCF

In the IEEE 802.11 MAC protocol for omnidirectional antennas, an exchange of RTS/CTS precedes DATA communication. Both RTS and CTS packets contain the proposed duration of data transmission. Nodes located in the vicinity of communicating nodes, that overhear either (or both) of these control packets, must themselves defer transmission for this proposed duration. This is called *Virtual Carrier Sensing* and is implemented by using a variable called the *Network Allocation Vector* (NAV). A node updates the value of the NAV with the duration field specified in the RTS or CTS. Thus the area covered by the transmission range of the sender and receiver is reserved for data transfer. We term this region as the *silenced region* in view of the fact that nodes located in this region do not initiate any transmission while communication is in progress. By silencing nodes in the *silenced region*, the RTS/CTS mechanism overcomes the hidden terminal problem [23].

The IEEE 802.11 (DCF) standard uses a backoff interval to resolve channel contention. Before initiating transmission, a node S chooses a random *backoff interval* from a range [0, CW], where CW is called the *Contention Window*. Node S then decrements the backoff counter once every clock "tick". When the clock reaches 0, node S transmits the packet. If the transmission from S collides with some other transmission, S chooses a new backoff interval from a range [0, $2 \times$ CW] and attempts retransmission. If a node senses a transmission during its backoff stage, it freezes the backoff counter. When the channel is once again idle, it continues counting down from its previous (frozen) value.

802.11 also specifies a duration for which nodes must perform (physical) carrier-sense before transmitting a packet. This duration of time is called the *Inter Frame Spacing* (IFS). Nodes wait for DIFS (*Distributed IFS*) before transmitting a RTS. Nodes wait SIFS (*Short IFS*) before transmitting CTS, DATA and ACK. The duration of DIFS is greater than the duration of SIFS.

3.1.2 Basic Directional MAC Protocol (DMAC)

We extend IEEE 802.11 to a simple directional version, named DMAC. While presenting the design of DMAC, we will assume that nodes are capable of *directional neighbor discovery*[1]. Therefore, every packet passed down to the DMAC protocol, is associated to a *transceiver profile* that specifies the direction in which DMAC must beamform to transmit the packet. The key operations in DMAC generalize the ideas in [14], [13], [18], and bear resemblance to the IEEE 802.11 protocol. However, DMAC incorporates suitable modifications to utilize the features of directional antennas.

Channel Reservation

Channel reservation in basic DMAC is performed using an RTS/CTS handshake, both being transmitted directionally. The details of these operations are central to several issues that arise with directional medium access. In describing the details of the operations, we refer to a sender node as node S and the receiver node as node R.

RTS Transmission

The MAC layer at node S receives a packet from its upper layers, along with the transceiver profile T^p to be used for transmission. Having received this, DMAC requests the physical layer to beamform according to the transceiver profile T^p. Let us denote this beamform by B^R, since this beam points in the direction of node R. To detect whether it is safe to transmit using beam B^R, node S now performs physical carrier sensing using B^R. If the channel is sensed idle, DMAC checks its *Directional NAV Table* to determine whether it must defer transmitting in the direction of node R. The DNAV table, as explained later, maintains a *virtual carrier sense* status for every direction from which a signal has recently arrived. Once node S finds that it is safe to transmit using B^R, it enters the *backoff* phase (similar to 802.11). When the backoff counter counts down to zero, DMAC at node S sends down the RTS control packet to the physical layer, meant to be transmitted to node R, using beam B^R.

[1] We will present directional neighbor discovery schemes later in this chapter.

RTS reception and CTS transmission

An idle node (like R) remains in the *Omni* mode, listening to the channel omnidirectionally. Once the RTS signal from S arrives at R, node R receives the RTS omnidirectionally and determines the direction of arrival (DoA) of the incoming signal. Let us denote this DoA direction as D^{SR}. If the RTS was received successfully, node R responds with a CTS by forming a beam, B^S, in the reverse direction of D^{SR}, namely D^{RS}. Neighbors of nodes S and R, that overhear the RTS and/or CTS, remember the directions in which the signals were received, and defer transmissions in these directions while the communication between S and R is in progress. In all other directions, the neighbors may initiate transmissions. Thus, the channel gets reserved only along the direction of the line joining nodes S and R[2].

CTS Reception and DATA/ACK Exchange

After transmitting the RTS, the sender node S waits for the CTS using the beam B^R that it originally used to send the RTS. If the CTS does not come back within a *CTS-timeout* duration (calculated as in 802.11), then S backs off for a larger duration of time to cope up with the potential congestion, and schedules a retransmission of the RTS. However, if S receives the CTS, it initiates the transmission of DATA using beam B^R. Node R, on receiving the DATA successfully, transmits an ACK using beam B^S.

Directional NAV (DNAV) Table

The Network Allocation Vector (NAV) is a status variable maintained by the IEEE 802.11 MAC for virtual carrier sensing. The value of the NAV is updated from the "duration" field of overheard packets, and indicates the duration for which the node must defer transmission to avoid interfering with some other communication in the vicinity.

Using directional antennas requires that the NAV be directional as well [14] [19] [24] [18]. Put differently, if a node receives an RTS or a CTS from a certain direction, then it needs to defer only those transmissions that are directed in (and around) that direction. A transmission intended towards some other direction need not be deferred. The *Directional NAV Table (DNAV)* is a table

[2] In multipath environments, the reserved region may not be along the line joining S and R.

that keeps track of the directions (and the corresponding durations) toward which a node must defer transmissions. DMAC consults the DNAV table before initiating a transmission, and thereby appears to handle the hidden terminal problems.

We illustrate DNAV with a simple example with reference to Figure 3.2. Assume that communication between nodes A and B is in progress, with A transmitting to B. During this time if node E has a packet to send to F, it must check its DNAV table to see if it is safe to transmit in the direction of F. However, note that node E would have already received a CTS from node B and updated its DNAV table. Thus on checking the DNAV table, E finds that it is unsafe to transmit in the direction of F. However, if E had a packet to send to D, the DNAV check would indicate that it is safe to transmit. In that case E can proceed to transmit towards D. At this point we would like to mention that even if F was a little displaced and the line joining EF did not coincide exactly with the line joining EB, E might still need to defer transmission to F in order to avoid interfering at B. Specifically, E should defer transmission towards F if the angle between EF and EB is less than some threshold, ϵ. We illustrate ϵ with the help of Figure 3.2. We assume that all nodes transmit with a beamwidth of 2β degrees. Now if node E has received an RTS or CTS from node B, it updates the DNAV table with the direction D^{EB}. Now if E wants to send a packet to node H, then it must beamform in the direction of D^{EH} to transmit the RTS. Node E checks whether, the angle between D^{EB} and D^{EH} is greater than a threshold ϵ, where ϵ is as following:

$$\epsilon = 2\beta + \theta$$

where θ is the angular separation of the edges of the two beamforms. If θ is negative then the two beamforms may overlap. For our simulations we have assumed θ to be zero.

3.1.3 Problems With Basic DMAC

In this section we discuss some key channel access problems with Basic DMAC, that have not been identified in the past. Although characteristic of DMAC, these problems may arise with other CSMA/CA-based directional MAC protocols as well.

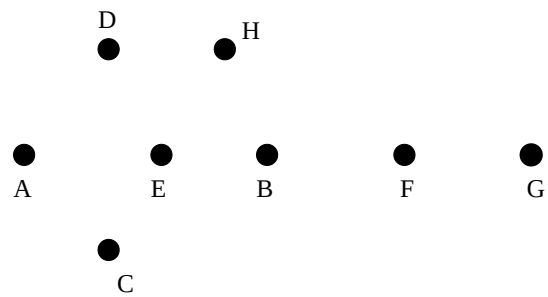

Figure 3.1: An example topology scenario

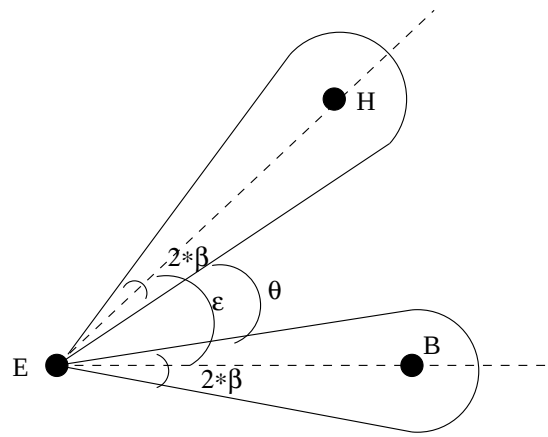

Figure 3.2: An example where node E wishes to send a packet to H while it has a DNAV entry for direction D^{EB}. The Figure shows a condition where node E can safely transmit to node H, since it would not interfere with node B's transmission

Neighbor Discovery

In the design of DMAC, we assumed that a higher layer provides DMAC with the *transceiver profile* necessary to beamform toward a desired destination. Recall that the transceiver profile includes the direction in which the node must beamform for the imminent transmission. In reality, determining the direction of transmission is a non-trivial task. For example, when a node enters the network, it needs to discover its neighbors using a distributed mechanism. Once the neighbors have been discovered (and beamforming directions associated to each of them), the node needs to maintain the mapping between neighbors and corresponding directions. In the face of node mobility, channel fluctuations, and multipath environments, such a mechanism is not trivial.

Although not the case with DMAC, other beamforming protocols may intend to exploit directional-

directional (DD) communication, especially when the transmitter and receivers are located far apart. In CSMA/CA protocols, where transmissions are not scheduled *a priori*, the transmitter may have to request the receiver to beamform towards it in order to accomplish actual data communication. However, observe that delivering such a request to the receiver (who is in the omnidirectional mode) will be difficult. As a result, the receiver remains unaware of when to beamform toward the transmitter. An efficient neighbor discovery protocol is clearly necessary to meet these "bootstrapping" requirements.

Self-Interference

Due to the higher gain of beamforming antennas, the interference footprint extends over longer distances, although in a specific direction. The silenced region is approximately conical, as opposed to the omnidirectional footprint, which is roughly circular. As a result, a specific communication on a particular link can interfere with several downstream links, if the downstream links are aligned in similar directions. Put differently, the impact of the shape of the *silenced regions* on the overall network performance is not straightforward, but it is indirectly dependent on the topology characteristics, traffic patterns and beamwidth of directional transmission.

Hidden terminal problems

The well-known hidden terminal problem in multi hop wireless networks can be resolved by the exchange of RTS/CTS control packets as in MACA [23] and 802.11. However, the RTS/CTS exchange assumes that these packets are transmitted omnidirectionally. Directional transmission of RTS/CTS introduces two new kinds of hidden terminals, discussed below.

(1) Hidden terminal due to asymmetry in gain:

Consider Figure 3.3. Assume that all nodes in this figure are currently idle. Nodes in their idle state (i.e., in *Omni* mode) have a gain of G^o. Now assume that node B transmits a Directional RTS (DRTS) to node F, and F responds with a Directional CTS (DCTS). Assume that node A (still in *omni mode*) is too far from node F to receive this DCTS. DATA transmission begins from node B to F, both nodes pointing their transmission and reception beams at each other, with a gain G^d.

Observe that node A cannot sense this transmission because the transmission from B arrives at A through the low-gain side lobes of B. While this communication is in progress, assume that node A has a packet to send to node B. Node A senses the channel with a directional beam (of gain G^d) pointed towards B and concludes that the channel is idle. Subsequently, it sends a DRTS to node B. However, since node F is receiving DATA with a gain G^d and a beam pointed toward node B (and node A), it is possible that the DRTS from A interferes at node F. In other words, sender and receiver nodes with transmit and receive gain of G^d and G^o, respectively, may be out of each other's range, but may be within range if they both transmit and receive with gain G^d (note that G^d is greater than G^o).

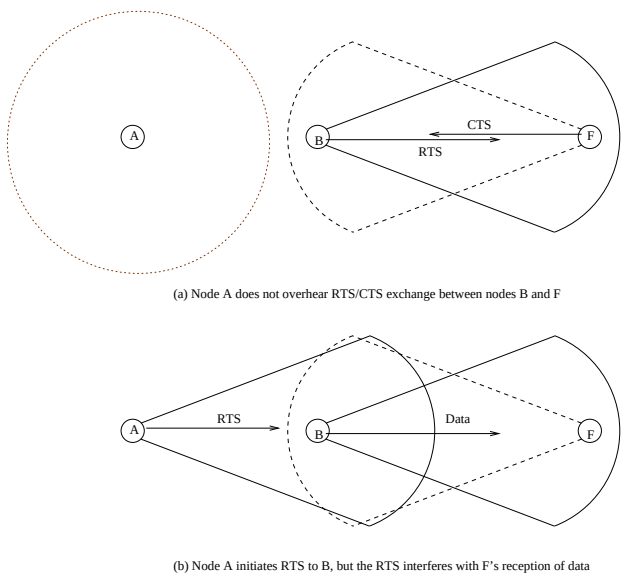

Figure 3.3: Hidden terminal problems arising when nodes change antenna patterns for communication.

While such hidden terminal problems are possible, closer examination revealed that they may not be frequent. Considering the example in Figure 3.1, node A may not always interfere with packet reception at node F (from B) because node A will almost certainly be further away from F than B. If node F locks on to signals from node B, then interfering at F may not be frequent - to interfere, A must be far enough from B to not sense B's transmission through B's backlobes, and yet needs to be close enough to F to cause a collision. With front-to-back ratios of 20dB, and receive threshold being twice sensing threshold, such occurrences will be rare.

(2) Hidden terminal due to unheard RTS/CTS:

Assume that in Figure 3.1, E is transmitting a packet to D. While this transmission is in progress, node B sends an RTS to node F. F responds with a CTS. Although E may be within the transmission range of F, it does not receive the CTS, since it is beamformed in the direction of D. On receiving the CTS from F, B starts transmitting data to F. While communication between B and F is in progress, assume that E finishes transmitting to D and now wants to transmit to F (or any other node in the direction of F). E's DNAV table indicates that it is free to transmit towards F and on performing physical carrier sense, E finds that the channel is idle. E transmits the RTS and a collision occurs at F (since F's receiving beam is pointing in the direction of E). Such kind of hidden terminal problems can be frequent with directional communications, but may not arise in the case of omnidirectional transmissions. With omnidirectional transmissions, node B will typically be aware of the ongoing transmission from E. This would prevent B from transmitting the RTS while E is engaged in communication. A potential tradeoff appears between spatial reuse and collisions (due to hidden terminals) when using directional antennas.

Deafness

Another drawback of directional beamforming is *deafness*. We explain deafness using the scenario in Figure 3.4. Assume that node C and B have packets to send to node A (node A being the common receiver for both the flows). At any given instance, if A is receiving B's packet, C would be unaware of it (when using Basic DMAC) and may transmit an RTS meant for A. Since A is beamformed in the direction of B, A does not receive the RTS and consequently does not respond with a CTS. Node C, on receiving no response from A, doubles its contention window, chooses a new backoff interval from the new window, and begins counting down. When the countdown reaches zero, C retransmits the RTS. Retransmissions can go on multiple times until A has finished the dialog with B, and has switched back to the omni mode. Unproductive retransmissions from C is an outcome of deafness.

Now consider the case where B has multiple packets to send to A. Once B has finished trans-

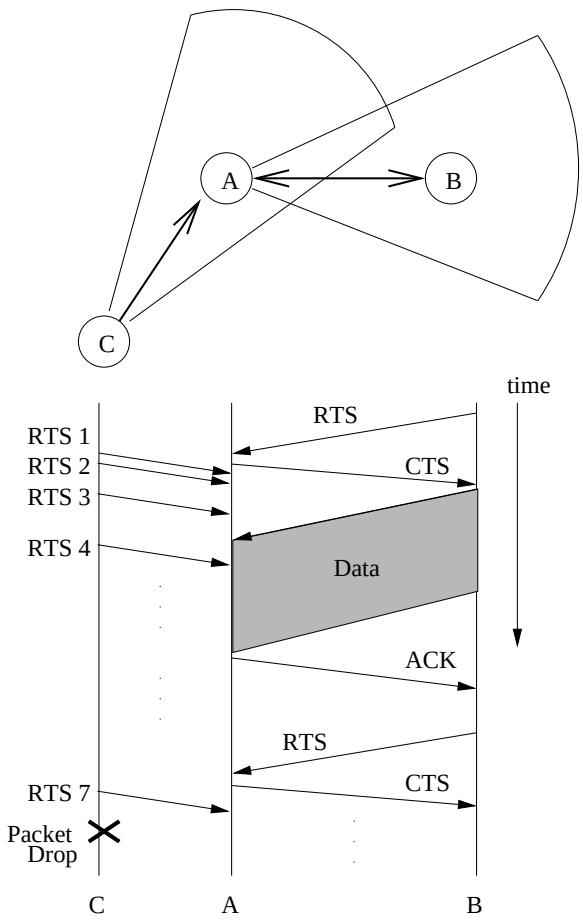

Figure 3.4: An example scenario depicting the possibility of deafness

mitting the first packet, it immediately prepares to transmit the next packet by choosing a backoff interval from the minimum contention window [0, CW_{min}] [3] It is likely that node C is still engaged in the backoff phase when B finishes counting down its small backoff value for the second packet.B acquires channel access. This can continue for a long time, causing C to drop multiple packets before it gets fortunate enough to steal channel access from B. After this point B encounters repeated deafness, and continues to drop packets whenever it exceeds the retransmission threshold. Multiple packet drops at the source node, without actual congestion or link failure, can adversely affect performance. Higher layer protocols, that use packet drops as indicators of the network condition, may be misled. End-to-end throughput and latency can degrade. Deafness also leads to short-term unfairness between flows that share a common receiver.

Now consider another case in which node C intends to send packets to A, and A intends to initiate packets meant for B. In other words, node A is both a receiver and a traffic originator. Also assume that both the originators of the flows (C and A) are always backlogged. Observe that once A has finished transmitting a packet, it beamforms immediately in the direction of B, carrier senses, and begins counting down the backoff interval in the directional mode (as required by Basic DMAC). Clearly, while backlogged, A would almost always remain in the directional mode, forcing C to drop all packets during the interval. Since C continues to remain beamformed towards A while attempting transmission, another node, say D, trying to communicate to C, may also experience prolonged deafness. A "chain" is possible in which none of the nodes communicate successfully – a "deadlock". This is a serious problem, caused when the intended receiver of a node is itself a traffic originator.

MAC-Layer Capture

In the design of basic DMAC, we assumed that nodes remain in the omnidirectional mode when they are idle. This is necessary because, with CSMA protocols, a node does not know *a priori* which neighbor may transmit to it next. In view of this, a node needs to listen in all directions (i.e., remain omnidirectional) so that it can beamform toward any (sufficiently strong) arriving signal. Once the signal has been detected, the node beamforms toward it and receives it.

Observe that such a scheme may not be able to extract the fullest benefits from beamforming. This is because, a node will be prone to receiving every (sufficiently strong) signal that impinges on its transceiver. Many of these signals will not be intended for the node, and the received packet will eventually be dropped at the MAC layer. We term this problem as "capture", to conve node'sty to lity to initiate useful communication while it is engaged in receiving an unproductive packet. The problem arises mainly because a node remains omnidirectional during its idle state, and therefore cannot avoid getting "captured". We argue that controlling the antenna during transmission and reception alone is not sufficient – the antenna patterns need to be regulated when nodes are idle as well. This in turn motivates protocols that are *capture-aware*, and thus can improve the spatial reuse in the network.

	Aggregate Throughput (Mbps)	
Traffic (per flow)	IEEE 802.11	DMAC
CBR 0.5 Mbps	0.512	0.512
CBR 1.0 Mbps	1.018	1.017
CBR 1.5 Mbps	1.294	1.104
CBR 2.0 Mbps	1.425	1.193
CBR 2.5 Mbps	1.427	1.192

Table 3.1: DMAC affected by deafness on a 2-hop flow

3.1.4 Evaluating Basic DMAC

In this section we evaluate the performance of Basic DMAC, and compare it to the IEEE 802.11 standard. We discuss results reflecting the pros and cons of directional communication. For our simulations we have used the Qualnet Simulator, version 3.6 [16]. The directional gain for the mainlobe is 10 dB and the sidelobe is -10 dB. The approximate transmission range when using IEEE 802.11 is 250 meters. The approximate transmission range of a directional link is approximately 400 meters. We use the *two-ray* propagation model. We do not consider node mobility in our simulation scenarios.

Simulation Results

Our main findings from the simulation of DMAC can be summarized as follows.

(1) The performance of DMAC depends on the topology and the traffic patterns in the network – nodes placed in a straight line may suffer due to alignment in directional interference. However, when nodes are more randomly placed, the performance of DMAC shows improvements over 802.11.

(2) Deafness and capture are serious problems, and can offset the benefits of spatial reuse available from beamforming antennas.

(3) The higher communication range of beamforming antennas allow for lower end-to-end delay, because of the possibility of finding fewer-hop routes from the source to the destination. This benefit needs to be exploited through appropriate routing protocols that utilize the higher gain of beamforming antennas.

The details of the experiments are presented next.

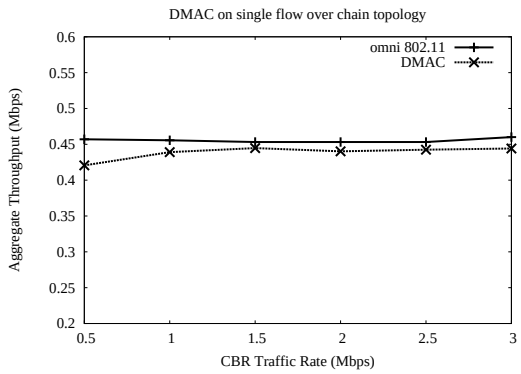

Figure 3.5: End to end throughput on a chain topology

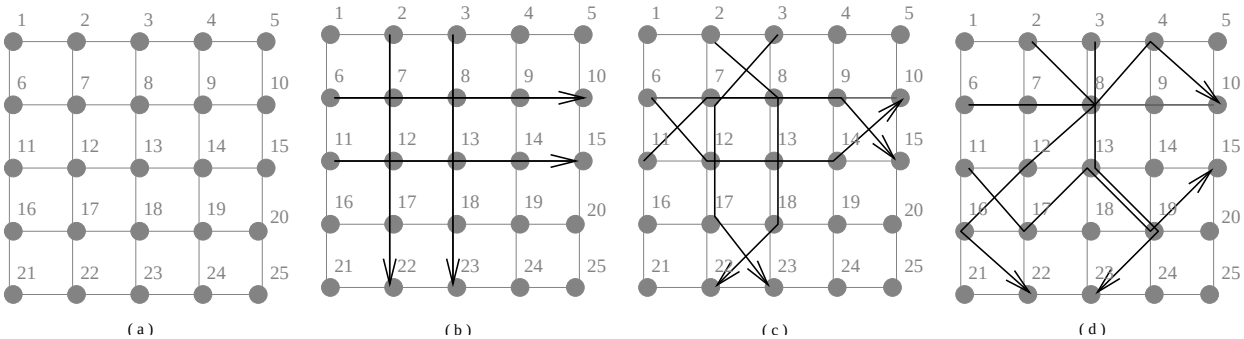

(a) 5x5 grid topology. (b) Aligned routes. (c) Less alignment. (d) Random routes.

Figure 3.6: Example scenarios showing the decreasing degree of alignment in the choice of routes.

Simple Topologies

To understand the impact of several factors that arise with beamforming antennas, we simulate simple toy topologies, and compare DMAC and 802.11 on each of them. Table I compares the end to end throughput for the simple scenario shown in Figure 3.4. As evident from the table, aggregate throughput achieved by DMAC is lower than 802.11. This is an attribute of deafness in which the channel remains underutilized when DMAC backs off exponentially. The problem does not arise with 802.11 since node C remains aware of the dialog between A and B, and does not backoff unnecessarily. In larger networks, the impact of deafness can be severe. We visit *deafness* in detail in Chapter 4, along with quantification of its impact and possible approaches to resolve it.

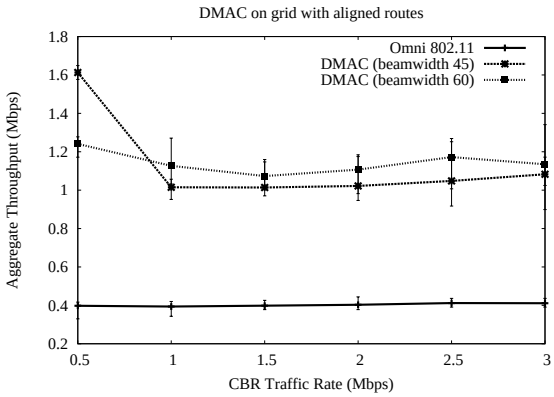

(a) Fully aligned routes in 6X6 grid

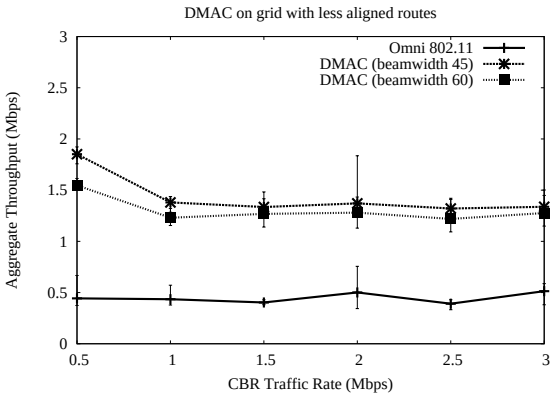

(b) Less aligned routes in 6X6 grid

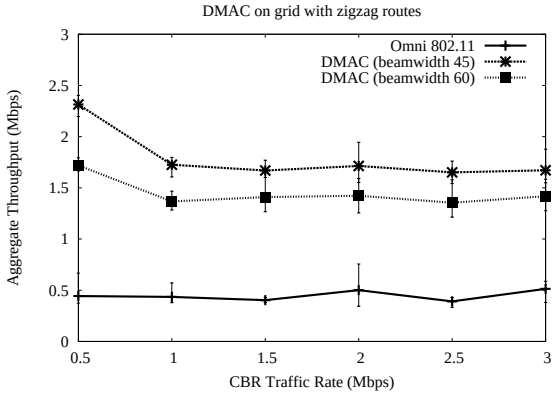

(c) Zigzag routes in 6X6 grid

Figure 3.7: Aggregate throughput for varying degree of route alignment

We simulate chain topologies to understand the performance of DMAC on longer flows. While we evaluate for a number of chain topologies, Figure 3.5 is one example where DMAC performs worse than 802.11. As mentioned earlier, this is a result of the higher interference of DMAC in the direction in which the main lobe is pointed. Observe that the communication range of 802.11 is lower, and thereby downstream links on the chain topology are capable of communicating simultaneously. In contrast, DMAC is unable to achieve such spatial reuse (since downstream links are within the interference footprint of the mainlobe).

Grid Topologies

We now simulate grid topologies to study the benefits of spatial reuse with beamforming antennas. Adjacent nodes on the grid are 200 meters apart. Observe that the previous topologies were disadvantageous to DMAC since they constrained DMAC from exhibiting its ability to achieve simultaneous transmissions in space. With grids of the size of 4X4, 6X6, and 8X8, DMAC can be expected to obtain higher reuse. For example, links on two different routes may be able to communicate in parallel, even though they are in the spatial vicinity of each other; 802.11 may not be able to leverage such possibilities. However, note that DMAC would still retain the issues of deafness and higher interference over the grid (since grids can be viewed as a conjunction of several chain topologies). The extent to which these issues affect performance depends on the type of routes chosen – a route that is linear (like a chain) is likely to affect DMAC more than routes zig-zag. To understand the overall performance of DMAC under grid topologies, we simulate various route patterns – examples are shown in Figure 3.6. Observe that the routes across Figures 3.6 (b), (c), and (d), are decreasingly aligned. In fact, the routes in Figure 3.6 (d) are hand-constructed to deliberately reduce alignment of downstream links.

We present results from 6X6 grid topologies in Figure 3.7. The simulations are run with 10 seeds each and averaged values are plotted with errorbars[3]. Observe that DMAC performs better than 802.11 in all the scenarios. Moreover, the performance gap increases as the alignment in the routes decreases.

In Figure 3.7 (a), DMAC performs better than 802.11 since a greater number of nodes can communicate simultaneously in comparison to a single chain. This makes DMAC better than 802.11. However, the benefit from simultaneous communication is offset by the higher interference of beamforming antennas as discussed earlier. As the routes become less aligned, the effect of directional interference reduces, and consequently spatial reuse improves. Thus, the performance gap increases for Figures 3.7 (b) and (c). However, DMAC is still affected by the serious effects of deafness and MAC-layer capture that together offset the potentially higher gains from beamforming antennas.

Random Topologies

[3]Errorbars imply the maximum and the minimum values obtained from all the 10 simulation runs.

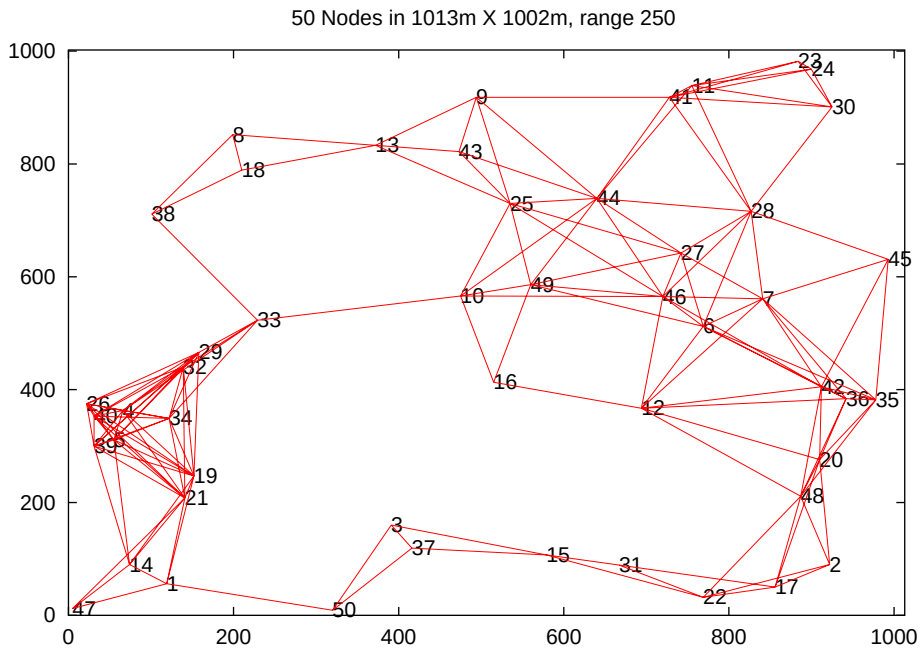

Figure 3.8: An example of a random topology.

We compare the performance of DMAC when simulated over 25 random topologies, each with random routes. Each topology is characterized by 50 nodes in a square region of side 1000m. An example topology has been shown in Figure 3.8. Observe that with randomly located nodes, the degree of route alignment is typically less, indicating the possibility of higher spatial reuse. The simulation results (each averaged over 10 runs) have been shown in Figure 3.9. Since each of the topologies are independently generated, the throughput and delay obtained from the simulations vary across these topologies. However, in all of them DMAC outperforms 802.11 because of higher spatial reuse – the aggregate throughput is higher, and the average end to end delay is lower. However, recall that even though directional interference is not a dominant concern in these random scenarios, the impact of deafness and capture still offsets the potential benefits. The ideal improvements from DMAC can be higher if these issues are addressed efficiently.

Summary

We began with the motivation to understand whether beamforming antennas can offer most of the expected benefits without significant changes to higher layer protocols. We designed a simple

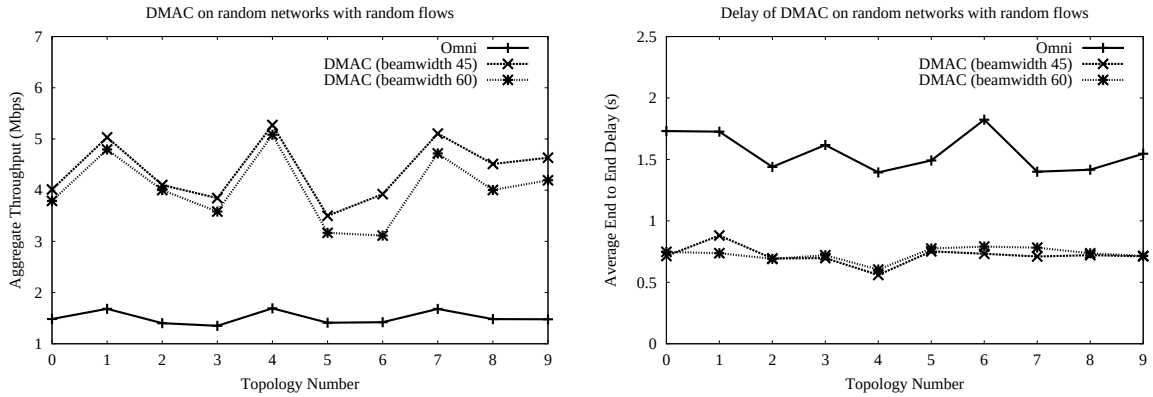

(a) Aggregate throughput in random topologies

(b) Average delay in random networks

Figure 3.9: Average aggregate throughput and average end to end delay for random topologies with random routes.

directional version of 802.11 (called DMAC), and evaluated the protocol for various topologies and traffic patterns. Our study shows that DMAC is impaired by several problems that appear as a result of beamformed communication. These problems, including new kinds of hidden terminals, deafness, and capture, do not appear with omnidirectional MAC protocols. As a result, variants of existing omnidirectional protocols are insufficient to address these problems. Simulation results show that DMAC's problems are pronounced in certain kinds of topologies and flow patterns – aligned routes are particularly unsuitable because higher interference primarily offsets the benefits of spatial reuse. In other unaligned and random topologies, DMAC performs better than 802.11 even though it remains impaired by deafness and capture. The full improvements from beamforming can only be achieved if the problems of deafness and capture are resolved efficiently. Additional performance gains may also be obtained if the higher transmission range of beamforming antennas are better utilized from the network layer. This motivates investigation into the impact of beamforming antennas on the performance of network layer protocols.

3.2 Routing with Beamforming Antennas

Exploiting beamforming capabilities from the network layer may lead to greater performance benefits. For example, the negative impact of higher directional interference at the MAC layer can

be compensated if the network layer exploits this higher communication range to discover fewer hop routes. Moreover, the spatial diversity from beamforming can be leveraged to direct traffic through less congested regions of the network. Understanding the impact of beamforming antennas on wireless multihop routing is important in designing a complete system solution. Motivated from this, we design a basic routing protocol capable of utilizing directional communication. We study the issues underlying directional routing, identify research challenges, and solve them through more refined protocols. We begin by modifying an existing routing protocol, named dynamic source routing (DSR), to suit the needs of beamforming antennas. We begin with a survey of related work in this area.

3.2.1 Related Work

While there has been some work on designing MAC protocols for directional antennas, research on directional routing has been limited [17], [15], [19], [25], [26]. In [15], Nasipuri *et al.* have utilized directional antennas for the purpose of on-demand routing. Their primary aim is to minimize routing overhead by intelligently using directional (sectored) beams for propagating routing information. Their results highlight the large routing overhead incurred by using omnidirectional antennas and show that it is possible to perform better using switched beam antenna systems. However, the authors have equalized the transmission range of directional and omnidirectional antennas, and therefore, have not explored the advantages of higher communication range of beamforming antennas. In our work, we evaluate the effect of various transmission ranges of directional antennas. [27] also assumes a conservative MAC model. We believe that medium access control could be more aggressive with directional communication and accordingly use a suitable MAC protocol below our routing solutions. We show that the interaction between the MAC and routing is important in determining the performance of the network. Optimizing each layer individually may not offer complete benefits. [17] proposes a proactive routing protocol (originally proposed for omnidirectional antennas [28]) over ESPAR antennas. The routing protocol performs well and gives reason to believe that directional antennas may be acceptable for ad hoc networking. However, the MAC protocol is complex, requiring several exchanges of control packets before actual data can be trans-

mitted. In this paper we consider a reactive routing protocol. Also, our MAC protocol is simpler than [17].

Several other papers [26], [25], [15] design directional routing solutions based on existing omnidirectional routing protocols. We present a brief overview of existing routing protocols, and then describe the contributions of these papers. Routing protocols can be broadly classified into "source-routed" or "table-driven" protocols [29], [30]. In source routing [31], the sender of a packet completely specifies the route that the packet must traverse to reach its final destination. Johnson *et al.* proposed dynamic source routing (DSR) [31] in which the sender node floods a route request (RREQ) probe in search of a route to the destination. Intermediate nodes that forward this request probe, append their identifiers to the probe. The probe that arrives first at the destination is assumed to have arrived on the optimal path. DSR uses this path for subsequent communication.

Table-driven routing protocols store routing information locally [32], [33], [26] [17]. Nodes exchange routing messages, either reactively or periodically, to update each other about the status of links in the network. When a node intends to send data packets to another node, it consults its routing tables for a route to the destination. It forwards the data packet to the appropriate neighbor in the route, who in turn consults its own tables to forward the packet further. An intermediate node is often faced with the decision to choose between two of its neighbors, both of which may be equally good for forwarding the packet to the final destination. Ties are broken randomly, without respecting the possibility that one of the nodes may not be suitable for immediate transmission. We believe that anycasting can be useful here – the MAC layer can make educated decisions in such scenarios, leading to potential benefits in performance. In this paper, we would refer to table driven protocols while discussing the details of MAC-layer anycasting. Issues arising from the use of source routing will be discussed separately in Section V.

In [26], the authors present the notion of *maximally disjoint routes.* Routes that do not share the same geographical area with other routes are preferred, thereby minimizing contention at the MAC layer. The authors show that via multi-hop feedback mechanisms, a source might be able

to divert traffic through less congested space. However, as we show later, unless the MAC protocol is equipped to discriminate between productive and unproductive traffic, such a zone-disjoint routing scheme can only offer limited benefits. In [25], authors propose to use directional antennas to prevent networks from getting partitioned. The authors show that efficient power consumption may be achieved by combining the use of omnidirectional and directional modes of antenna systems in a DSR-like routing protocol. Their work highlights the benefit of higher communication range of beamforming antennas. In this dissertation, we identify several issues that require the routing protocol to be directional-antenna-aware, failing which the system performance may degrade.

Some work has also been done on broadcasting in wireless networks using directional antennas. [34] analyzes the convergence properties of broadcast routing by considering several antenna technologies and analyzing the performance of BIP and DBIP protocols. [35] proposes to perform energy efficient broadcast by forming a "relative neighborhood graph (RNG)" at the source node. Each node on the RNG forwards packets to only those nodes that are its neighbors on the RNG, using directional transmissions. Based on the observation that the average degree of relative neighborhood graphs are small (of the order of 2 to 3), the authors claim that energy can be conserved even if multiple transmissions are performed (each on a narrow beamwidth). Assuming that energy varies linearly with the beamwidth of transmission, the savings are significant. In addition, if power control can be achieved for each neighbor on the RNG, further conservation may be possible. Also, using the RNG may reduce the probability of collisions and redundant transmissions.

[36] proposes to mitigate the broadcast storm problem by transmitting packets in the directions in which the packet has not been transmitted earlier. The authors present 3 protocols – (i) On-Off mode, (ii) relay-mode, and (iii) location based broadcast. In on-off mode, a node forwards packet in the directions in which there has been no previous transmission. A random delay is associated to each transmission – if the node does not overhear a packet transmitted in that direction before the expiry of its delay timer, it transmits the packet in that direction. Relay-mode chooses a farthest relay node in each direction, and requires the selected relay to forward the packets. With location based broadcast, a node delays the transmission in a direction based on the additional area that the

transmission would cover. Thus if a relay is close to the source node, then the additional area that it would be able to cover by relaying is small. Such a node selects a higher delay for transmission. Such a scheme can expedite the propagation of the broadcast packets over the entire network.

[37] proposes topology control in ad hoc networks using directional antennas. By identifying the ability of directional antennas to beamform in selected directions, the authors show several properties that can be achieved in the topology. The authors show that there is an implicit equivalence in reducing omnidirectional power and controlling the antenna beamwidth. Using this notion, and existing results from percolation theory, the paper illustrates how a transition phase can be achieved in terms of connectivity, when the beamwidth is greater than a critical threshold. In other words, when the beamwidth and the number of beams increase over pre-calculated thresholds, the network topology becomes connected with high probabilistic guarantee. In addition, since power expenditure is considered proportional to the beamwidth of transmission/reception, the paper claims the possibility of power savings using directional antennas.

3.2.2 Basic Directional Dynamic Source Routing (DDSR)

Several tradeoffs arise when routing protocols (originally designed for omnidirectional antennas) are ported over a directional antenna system. For example, fewer-hop routes may be discovered due to the higher transmission range of directional antennas. However, performing a simple neighborhood broadcast, often required to discover such routes, becomes more complicated. To broadcast a message to all its neighbors, a node will need to sequentially transmit the packet on each of its beam. Consequently, neighbors of a node may receive broadcast packets at different points in time (unlike with omnidirectional antennas). Also, sequential transmission on every beam (called *sweeping*) will incur higher control overhead. More subtle tradeoffs arise. Packets dropped due to deafness may be incorrectly interpreted as a link failure at the network layer – the routing protocol may unnecessarily initiate a new route discovery phase. A routing protocol that remains unaware of the antenna characteristics can degrade system performance significantly. We propose to understand the implications of these tradeoffs on performance of routing protocols, and suggest

modifications and optimizations where applicable. We begin by extending DMAC with a broadcast capability necessary for the network layer above – we call the extended protocol DiMAC.

DiMAC Protocol Description

The proposed DiMAC protocol is very similar to the DMAC described earlier in this chapter. Directional RTS/CTS packets are used for reserving the channel, followed by directional Data and ACK packets. The notion of DNAV, as proposed for DMAC, is also applicable to this MAC protocol. The difference is that DiMAC comprises of a neighbor discovery module and a broadcast module, as described next.

DiMAC maintains a look-up table containing two fields; {*NeighborId, AntennaBeam*}. When communicating to a specific *neighbor*, DiMAC consults the look-up table to obtain the suitable beam of communication. The look-up table is a cache that is updated periodically. As explained later, the look-up table is used to efficiently handle node mobility at the MAC layer. As a result, the network layer remains unaware of the antenna handoffs that happen when nodes move from the coverage area of one beam to another.

DiMAC also implements a *sweeping* mechanism whenever it is required to send out a broadcast packet. In *sweeping*, DiMAC sequentially transmits the same broadcast packet over every beam, emulating an omnidirectional broadcast. While performing a *sweep*, packet transmissions are not preceded by *backing off*. However, with sweeping, the communication range of the broadcast packet can be higher than omnidirectional antennas, and thereby far away neighbors can be reached. We call these neighbors *DO-neighbors* because the link between these neighbors are formed by a directional transmission and an omnidirectional reception. On similar lines, when using omnidirectional antennas, neighbors would be called *OO-neighbors* since the link between the neighbors are formed using omnidirectional transmission and omnidirectional reception.

DSR over DiMAC

We begin our evaluation of ad hoc routing using DSR and later suggest improvements. With DSR, a source node initiates a route discovery process to discover routes to its intended destination node. Route discovery requires the broadcast of RREQ packets, emulated in DiMAC using a *sweep*. As discussed earlier, sweeping discovers far-away neighbors at the expense of greater control overhead.

Once a route has been discovered, the source node initiates unicast packet transmission over the discovered source route. To unicast packets to the downstream neighbor, a node consults its look-up table to obtain the beamforming information. Look-up tables are always updated upon receiving or overhearing packets from the channel. If node A receives or overhears a packet from node B on beam i, then the look-up table is updated with the tuple $< B, i >$. In mobile scenarios, it is possible that information cached in the look-up table gets stale. We incorporate a *scanning* mechanism within DiMAC to handle this. A *scan* is essentially "hello" packets, transmitted sequentially over all antenna beams (i.e., swept), whenever necessary. If node B receives a "hello" from node A, B replies using the same antenna beam with which it received the "hello". The reply from B is delayed to allow node A to first finish transmitting all the hellos (otherwise, A would not receive B's reply due to deafness). *Scanning* is just one of many possible ways to proactively perform directional neighbor discovery.

3.2.3 Evaluating Basic DDSR

We use the Qualnet simulator [38] version 3.6, for simulating our protocols. We simulate our scenarios in a bounded region of 1000 X 1000 square-meters, with CBR traffic for data communication. The tradeoffs we discuss in this section arise from this particular interaction between Directional DSR (DDSR) and DiMAC. However, we believe that our broad observations will apply to other protocols as well.

The main tradeoff analyzed in this evaluation appears as a result of sweeping. Sweeping incurs higher delay because several copies of the same packet have to be sent sequentially over each beam. Moreover, there is a high control overhead associated to sweeping. However, on the positive

side, sweeping allows a source node to discover shorter routes. Also, due to the use of directional antennas, several nodes in the wireless vicinity are capable of simultaneously forwarding the RREQ packet. For example, if a node X receives a RREQ from node Y, X can begin its own sweep even though Y might not have finished its own sweep for that RREQ packet. As a result, the forwarding of RREQ packets are pipelined between nodes, partly compensating for the sweeping delay mentioned earlier. To evaluate these tradeoffs, we simulate several static and mobile scenarios, with different flow patterns. We present our findings below.

Main Simulation Results

We compare the performance of 3 protocols – DMAC, DDSR, DDSR-Manual. DDSR-Manual is introduced to investigate the performance that can be achieved if the best-possible routes are used with DMAC. We determine the best possible routes through a centralized manner based on the distances between nodes, and the communication ranges of the beamforming antennas. Observe that the performance gap between DDSR-Manual and DDSR represents the degradation due to *sweeping* and *scanning*. The gap between DDSR-Manual and DMAC is a measure of improvement available from using shorter hop routes with beamforming antennas.

Chain and Grid Topologies

We begin by simulating static chain and grid topologies (similar to the ones in DMAC). Figure 3.10(a) and (b) shows the aggregate throughput for DMAC, DDSR, and DDSR-Manual, for chains and grid topologies, respectively. For both the figures, DDSR and DDSR-Manual performs better than DMAC because they are both capable of finding fewer hop routes. However, DDSR-Manual performs better than DDSR because the latter suffers from deafness and capture during the phase of route discovery. Observe that broadcast packets, such as route request (RREQ) packets, are not acknowledged. As a result, if a neighbor of a node, X, is deaf during an ongoing broadcast, or is captured by another ongoing communication, the neighbor may not receive the broadcast at all. If this neighbor happens to be on the best route from X, then the best route will not be discovered. Therefore, DDSR often uses sub-optimal routes. Since DDSR-Manual always uses the optimal routes (in terms of hop-count), it achieves higher throughput in comparison to DDSR.

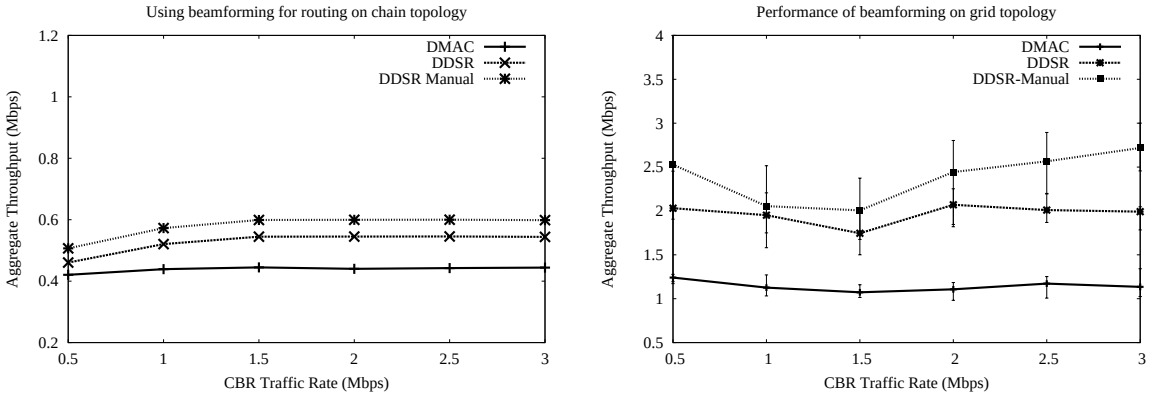

(a) Throughput in chain topology

(b) Throughput in grid topology

Figure 3.10: Performance improvement for chain and grid topologies, when routing using beamforming antennas.

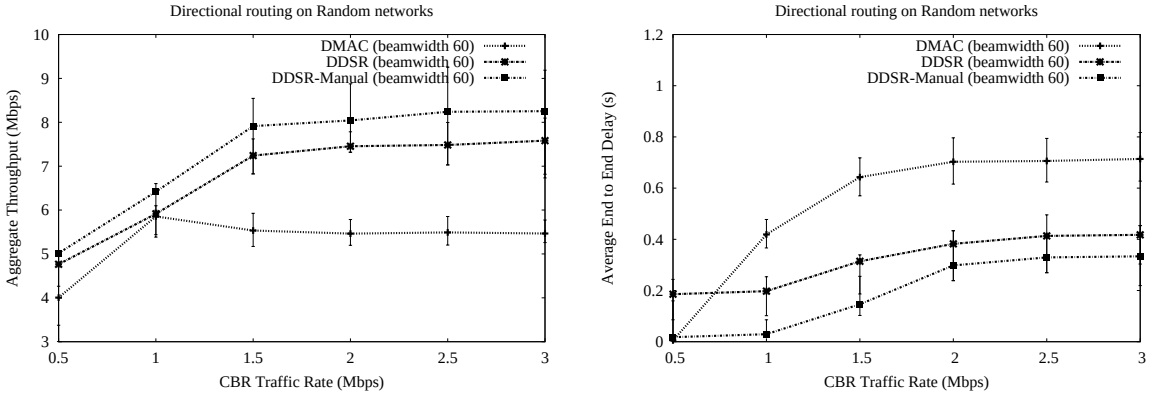

(a) Throughput benefits from beamformed routing

(b) Reduction in latency due to shorter routes

Figure 3.11: Aggregate throughput and average delay in random topologies with directional routing.

Random Topologies

We also simulate DMAC, DDSR, and DDSR-Manual for random topologies and traffic patterns. As with the earlier simulations, we use 50 nodes in a bounded square of side 1000m. We introduce multihop flows between randomly generated source-destination pairs. Figure 3.11 (a) and (b) shows the throughput and delay results from the simulations. While the performance of DDSR improves

over DMAC in random topologies, the performance gap between DDSR-Manual and DDSR continues to be non-marginal. Close examination exposes several other network layer issues that impact DDSR, in addition to the overheads of *sweeping*. We discuss these issues next.

(1) We notice that in denser regions of the network, the route selection suffers severely with DDSR. The hop-count of DDSR is higher than DDSR-Manual, even when simulations are performed without fading. We observe that this happens because RREQ collisions increase in denser regions since a large number of RREQs are unreliably sent within a small time window. The collision of RREQs results in the elimination of good paths, forcing the destination node to choose a suboptimal route. The impact is particularly pronounced when 2 or 3 hop routes are feasible between the source and the destination. In such cases, choosing routes that have slightly higher hop-count can lead to substantial degradation. For example, DDSR performed 33% worse than DDSR-Manual for the network in Figure 3.13(a), and only 13% worse in the network of Figure 3.13(b). DDSR-Manual does not suffer from this problem, and therefore performs better than DDSR – the performance gap is non-marginal, showing that the impact of RREQ collisions (as a result of sweeping) is serious[4]. Clearly, *sweeping* needs to be addressed in order to reduce the indirect effects of RREQ collisions in the network.

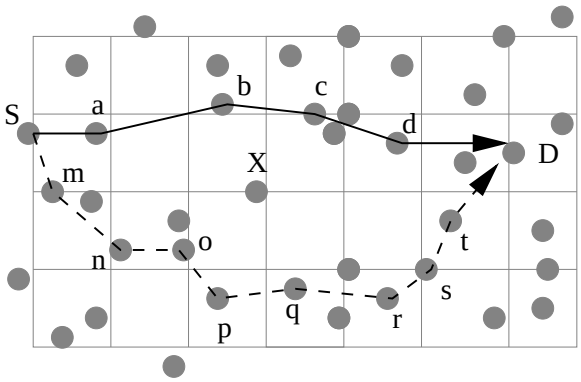

Figure 3.12: DDSR's ability to connect holes in the network.

[4]In our discussions on the design of routing protocol, we will describe how deafness and capture also contributes to choice of sub-optimal routes in DDSR.

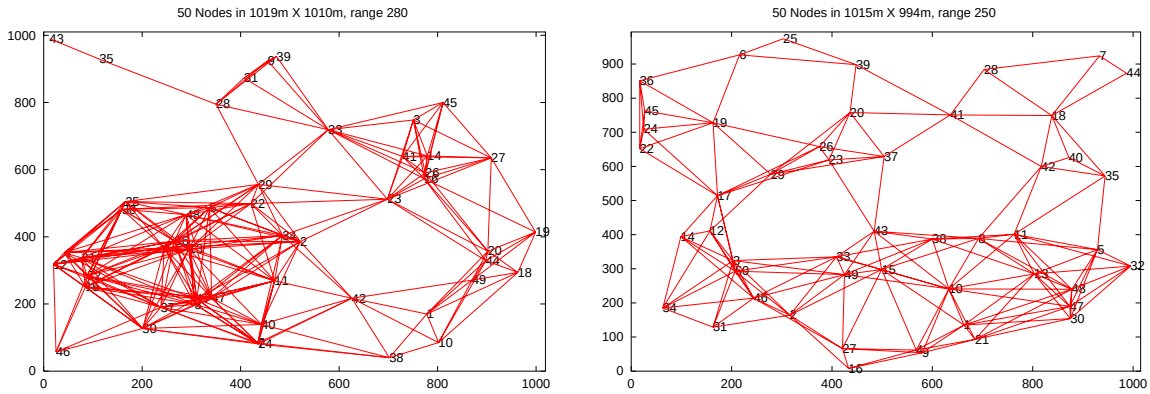

(a) RREQ collisions and sub-optimal routes prevalent

(b) Less RREQ collisions due to sparser node distribution

Figure 3.13: DDSR shows better performance for the topology on the right.

(2) While we studied the negative impacts of sweeping, we also noticed that in an ideal case, the benefits of higher transmission range was significant. This was evident from the high throughput and low latency achieved by DDSR-Manual in Figure 3.11. Simulation traces also show that the benefits are pronounced when the network is sparse, or when the flows are between source-destination pairs far apart. In sparser regions, the higher transmission range of beamforming antennas prove to be effective in bridging holes in the network. Figure 3.12 shows how DDSR is capable of bridging holes, whereas DMAC adopts a longer route around the periphery of the hole. In several cases, where the network partitioned with DMAC, DDSR managed to maintain connectivity. The result of such improved connectivity reflects on the performance, especially in the end to end latency achieved by DDSR-Manual. This justifies the need for using the higher transmission range of beamforming antennas, especially where the node density is not extremely dense. In view of this, we will visit optimizations to sweeping in our design of a new routing protocol for beamforming antennas.

(3) We then study the impact of node mobility on the performance of DDSR. We could not simulate DMAC and DDSR-Manual for these experiments because the routes for these protocols are statically specified, and mobile scenarios would require routes to change constantly. The results

from the mobility experiments are presented in Figure 3.14. Clearly, the performance of DDSR degraded significantly with mobility. The degradation primarily happened because DDSR initiated route discoveries every time a neighbor of a node moved away into the coverage area of a different beam. Clearly, the network protocol needs to be aware than link failures do not indicate disconnection when beamforming antennas are in use. We revisit this issue in Chapter 5, and modify DDSR to accommodate mobility more gracefully.

(4) Beamforming antennas improve robustness to link failures. Earlier studies have shown that when using DSR, links on a discovered route often proved to be unstable [39] for (large-sized) data packet communications. This is because routes are discovered using small-sized RREQ packets transmitted at low data rates; a link might offer sufficient SINR to transport these RREQ packets, but may not be strong enough for large-sized data packets transmitted at higher data rates. With DDSR, such events were rare since routes discovered by RREQ packets comprised of DO-links. However, when larger unicast packets were transmitted over these links, both the transmitter and the receiver beamformed toward each other (DD-link), increasing the SINR at the receiver. Link failures due to channel fluctuations and fading reduced, resulting in more reliable communication. This improved stability in route selection resulted in higher delivery ratio of data packets. The results for this experiment are presented later.

(5) We also study the impact of beamwidth on the performance of DDSR. Counter to intuition (which might suggest that a narrower beamwidth should offer better performance) we encounter a new tradeoff. When beams are extremely narrow, we find that the impact of deafness and the overheads of sweeping together offset the benefits of spatial reuse and range extension. On the other hand, at large beamwidths, spatial reuse degrades significantly, reducing the aggregate throughput in the network. We evaluat the variation of throughput for increasing beamwidth, and observe that 45° appear the best, when complete sweeping is performed at each node. Figure 3.14(b) presents the variation of performance with beamwidth for static random topologies. Interestingly, when we evaluate the same performance with node mobility, results appear different. We observe that slightly narrower beamwidths are better (around 30°). This is because with narrower beamwidths,

the communication range increases. A higher communication range is useful in mobile scenarios because the lifetime of a link increases with higher transmission ranges. As a result, routes break less frequently, leading to fewer route discoveries. These issues are revisited in Chapter 5 for a more detailed study.

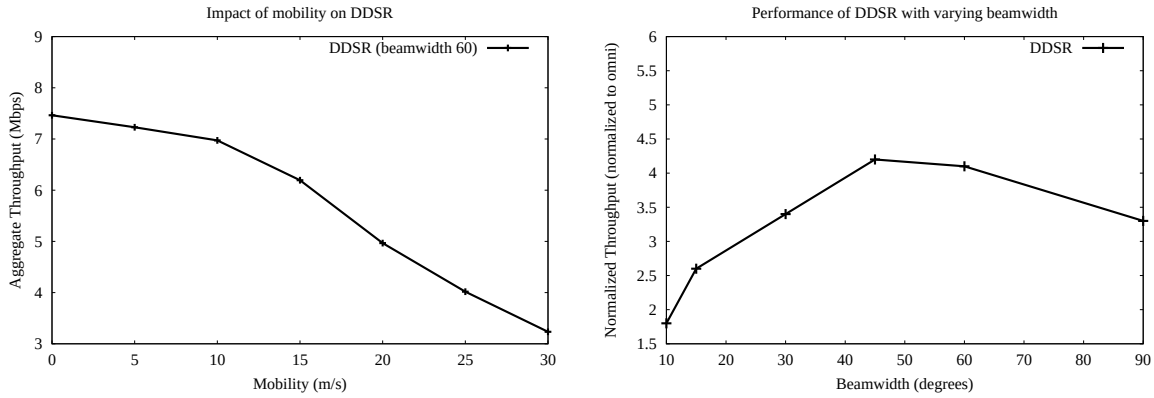

(a) Impact of mobility

(b) Performance variation with beamwidth in static networks

Figure 3.14: (a) Mobility affects performance due to frequent link breakages. (b) Performance varies with variation in beamwidth.

Summary

To motivate the need for antenna-aware routing protocols, we proposed to execute a simple routing protocol on a directional MAC protocol. We incorporate a neighbor discovery module and a broadcast module necessary for the routing functionalities. We called this directional routing protocol DDSR. The evaluation of DDSR revealed the benefits of higher transmission range of beamforming antennas. Results show that because of its ability to discover fewer-hop routes, DDSR can achieve higher throughput and lower latency. Benefits may be possible even in mobile scenarios because higher communication ranges increase the lifetime of a link, leading to fewer route breakages. However, the overhead of broadcasting with DDSR is significant, and unless optimized, can offset the benefits in certain traffic scenarios. The performance of DDSR suffers under high node

density, primarily because of broadcast packet collisions. Unless addressed, these issues can prevent a directional routing protocol from achieving its fullest performance benefits.

Chapter 4

MAC PROTOCOLS USING BEAMFORMING ANTENNAS

Attaching beamforming antennas, without the support of higher layer protocols, is insufficient for wireless multihop networks. Our investigation of a simple MAC protocol in Chapter 3 shows that several issues arise when simple variants of existing (omnidirectional) protocols are used in conjunction with beamforming antennas. Initial performance evaluation shows that these issues can offset the benefits from beamforming. In particular, *deafness* and *capture* emerge to be the key challenges at the MAC layer. Unless protocols are designed to specifically address them, the benefits of beamforming may not be fully achieved. Resolving these deficiencies can lead to significant improvement in network performance. This chapter presents a deeper evaluation of the impact of *deafness* and *capture*, and proposes protocols that carefully address these issues at the MAC layer. The protocols, namely ToneDMAC and CaDMAC, are CSMA/CA based protocols that better exploit the potentials of beamforming antennas. The protocols aim to maximize spatial reuse in the network.

We begin with a discussion on *deafness* and its detailed impact on the performance of the network. Based on the insights, we design a new MAC protocol, named "ToneDMAC", that alleviates deafness, while upholding the spatial reuse available from beamforming antennas. We evaluate the performance of ToneDMAC, and discuss issues with its implementation. Then, we visit the problem of MAC-layer *capture* and follow a similar organization – we study its impact on network performance, develop a protocol named "Capture-Aware DMAC (CaDMAC)", and

evaluate its efficacy to reduce both capture and deafness. Simulation results demonstrate the performance improvements available from ToneDMAC and CaDMAC.

4.1 Deafness

The ability to beamform with beamforming antennas is beneficial in terms of spatial reuse and/or range extension. However, as mentioned briefly in the previous chapter, several tradeoffs arise when designing protocols for such antennas. One such tradeoff is *deafness*. Briefly, deafness is caused when a node, C, attempts to initiate dialog with a node, A, while A is engaged in communication with another node, B, as shown in Figure 4.1. Node A fails to receive signals from C since A remains beamformed towards B over the duration of communication. Node C interprets the absence of a reply from A as indicative of a collision at A[1], and retransmits the packet. This can repeat multiple times, until node A has finished the dialog with B. The impact of deafness can be cascading, and degrade network throughput as well as fairness. This section evaluates the impact of *deafness* in detail, and proposes ToneDMAC (an enhancement to DMAC) to mitigate the problem.

The problem of deafness does not appear with omnidirectional antennas. When using omnidirectional antennas, the neighbors of a communicating node remain aware of the ongoing communication. For example, in Figure 4.1, node C would have received an RTS from node A, and thereby would not initiate transmission to A until the dialog between A and B is over. However, spatial reuse gets affected due to omnidirectional communication, resulting in a performance tradeoff. We propose a tone-based directional MAC protocol (ToneDMAC) to address this tradeoff. Simulation results show that addressing deafness can lead to significant benefits, thereby improving the performance of our basic directional MAC protocol.

4.1.1 Related Work

Recently, several MAC protocols have been propose that suitably adapt 802.11 for beamforming antennas. Ko *et al.* [14] propose DMAC that can benefit from transmitting only in desired di-

[1]IEEE 802.11 and many other CSMA/CA protocols, including those using directional antennas, are designed based on this assumption.

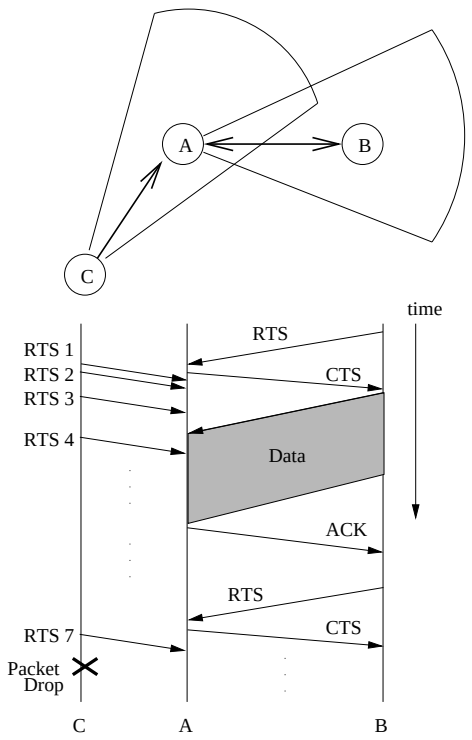

Figure 4.1: A scenario illustrating the problem of deafness

rections. Nasipuri *et al.* [27] also propose a similar protocol in which RTS/CTS exchanges were carried out omnidirectionally, while Data/ACK exchanges were accomplished directionally. Omnidirectional RTS/CTS clearly offsets spatial reuse – a key advantage of using directional antennas. Besides assuming conservative antennas, these protocols also suffer from the problem of deafness, and thereby their improvement is limited. In view of deafness, Elbatt *et al.* [16] propose an interesting idea – they use RTS/CTS to inform the neighborhood about the beam indices to be used for the imminent communication. However, the RTS/CTS are omnidirectional and thereby spatial reuse has to be traded off to resolve deafness. Trading off spatial reuse can be expensive in term of network throughput and latency, as shown in [40]. Authors in [40] propose to evaluate the performance of several MAC strategies for combination of both broadcast and unicast traffic. They claim that several papers have evaluated the traffics individually – realistic scenarios require both traffic models to co-exist. The papers presents three variations of MAC protocol – ORTS-OCTS, DRTS-OCTS, and DRTS-DCTS, where ORTS indicates omnidirectional RTS and DRTS indicates directional RTS. The results show that aggressive channel access (using DRTS-DCTS) yields better performance, primarily due to the ability of higher spatial reuse. Bandyopadhyay *et al.* [17] present

another MAC protocol that informs neighborhood nodes about ongoing communications through additional control messages. They propose to send proactive messages in the neighborhood based on the active flows in the network. This protocol handles deafness partially – however, they implicitly assume knowledge of network traffic.

Some existing proposals have utilized "busy tones" for omnidirectional, and more recently, for directional antennas. In [41], Deng *et al.* have proposed DBTMA, that uses (omnidirectional) transmit and receive busy tones to avoid hidden terminal problems. Huang *et al.* [42] have extended this idea to the case of directional antennas. The protocol assumes multiple transceivers, capable of transmitting data packets as well as busy tones directionally. Moreover, the protocol suffers from the problem of deafness. Similar to DBTMA, and the protocol in [42], our proposal for deafness (ToneDMAC) uses an additional control channel for signaling. However, ToneDMAC requires a single transceiver. Tones are assigned on the control channel, and need not be transmitted in parallel with data packets. In this respect, tones are different from "busy tones" – the latter is typically transmitted (using a second transmitter) to indicate that the node is "busy" with ongoing transmission/reception.

Korakis *et al.* [43] attempt to address the problems arising from directional antennas, including deafness. To inform neighbors of a communicating node, the authors propose to transmit directional RTS packets on every beam. Transmitting multiple RTS packets for each transmitted data packet can drastically increase the control overhead of communication. For example, with 6 beams, 6 RTS packets would need to be transmitted before the transmitter can expect a CTS. If the CTS is also not transmitted in all directions, then the problem of deafness is not fully solved[2], while transmitting 6 CTSs adds further to the overhead. In addition, for the protocol in [43], the duration of carrier sensing is increased to equal the transmission time of 6 RTS packets. Using parameters from IEEE 802.11b, the transmission of each RTS packet is $207\mu s$ (i.e., $192\mu s$ for the physical layer PLCP preamble and header, and $15\mu s$ for the 20-byte RTS at 11 Mbps data rate). Therefore,

[2] [43] proposes to send only one directional CTS. As a result, several neighbors of the CTS-sender are likely to remain unaware of the ongoing communication, leading to the possibility of deafness. We discuss this scenario in detail, in Section IV, and show how ToneDMAC can address such problems.

a transmitter must carrier sense for $6 \times 207 = 1242\mu s$, before it can transmit a packet. When data packets of 1024 bytes require transmission time of only $936\mu s$ (including 192 μs for PLCP preamble and header), carrier sensing for $1242\mu s$ can be detrimental. Throughput degradation and increased latency will be observed as a result of such modifications. [43] shows that the degradation in throughput when using their proposal of multiple RTSs is only 3% to 7%, as compared to a DMAC scheme that uses a single RTS. This result appears to have been obtained with parameters that do not account for RTS overhead accurately. Let us consider a simple example where a node A transmits to node B, and compare the channel utilization using simple calculations. Assume that the protocol in [43] uses N beams, and thus transmits N RTSs. We denote the relevant terms with the following symbols: Duration DIFS is δ, average backoff is β, SIFS is s, and the transmission time of RTS, CTS, Data and ACK packets are represented as R, C, D and A respectively[3]. Then the total time to transmit a packet, using a single RTS, is given by

$$TotalTxTime = \delta + \beta + R + s + C + s + D + s + A$$

While using the proposal in [43], the total time to transmit a packet must account for the time of multiple RTSs and the duration of longer carrier-sensing. This is given by

$$\begin{aligned} TotalTxTime &= (N \times R) + \beta + (N \times R) + s + \\ &\quad C + s + D + s + A \end{aligned}$$

Using respective values of $TotalTxTime$, we calculate the channel utilization according to the formula $Utilization = \frac{D}{TotalTxTime}$. For 11 Mbps data rate, the utilization for 802.11, as well as DMAC, is 0.38, while that for the scheme in [43] with N = 4, is 0.22 – a degradation of almost 42% compared to DMAC. With data rate of 2 Mbps, the degradation for [43] (with N = 4) is approximately 26%. Degradation of 3% to 7%, as shown in [43], is probably a result of omitting the physical layer PLCP preambles and headers in their simulation. ToneDMAC uses omnidirec-

[3]The values of each of the durations are obtained from the IEEE 802.11b specifications [3]: A *slot time* = $20\mu s, DIFS\delta = 50\mu s$, *average backoff* $\beta = 15slots = 300\mu s, SIFS = 20\mu s$ and the duration of transmission of packets are calculated according to the formula $(T_{preamble+header} + \frac{PacketSize}{DataRate})$, where $T_{preamble+header} = 192\mu s$ for 802.11b.

tional tones to mitigate the problems of deafness, while keeping the associated control overhead low.

As detailed later, the problem of deafness does not appear when communicating nodes manage to inform their surrounding neighbors about the imminent communication. For example, in Figure 4.1, C may not attempt communication to A, if C was aware of the ongoing dialog between B and A. However, we later discuss why informing the neighborhood can be expensive, especially under heavy traffic. We argue that not informing neighbors explicitly of an imminent communication can favor performance improvement. This motivates the design of ToneDMAC.

In [44], authors have proposed the use of multiple beam antennas to address the tradeoff between deafness and spatial reuse. Observing the history of traffic patterns, a node attempts to determine its set of neighbors that are likely to experience deafness from it. The node tries to form a multi-mainlobe beam that covers as many as these nodes, and transmits the RTS/CTS packet. The proposal works if the degree-of-freedom of the antenna is in most cases more than the set of identified neighbors. In addition, the time of switching between one multi-mainlobe beam to the other is not negligible, and can make per-packet switching infeasible for practical antennas.

[45] propose a MAC protocol based on adaptive beamforming. They present an idea in which a transmitter node sends a tone for a duration required for beamforming at the receiver. The receiver, based on this tone, beamforms in a way that maximizes the SINR of the signal of interest, while steering nulls to the direction of potential interferers. Using M elements for an adaptive antenna, $(M-1)$ degrees of freedoms are possible, and therefore the sum of the number of beams and nulls can be $(M-1)$, each pointing effectively in $(M-1)$ directions. Once the tone has been transmitted, the transmitter then sends the data packet (RTS/CTS packets, as used in 802.11, is omitted). On receiving the data packet correctly, the receiver sends back an ACK. Clearly, a problem with this scheme is that a node may be involved in other communications, or may have steered a null in the direction of the potential transmitter. This can cause acute deafness – detectable only after a packet has been transmitted. In addition, since nodes beamform (adaptively) to the stronger tones, it is possible that far away node pairs experience unfair channel access.

4.1.2 Impact of Deafness

While exploring the space of directional MAC protocols, we gained insight into the impact of deafness. Careful analysis and simulation results indicated that deafness can be a serious problem, offsetting some of the key benefits of beamforming. While some approaches appeared to resolve deafness, the penalties they incurred deemed them impractical. This section presents our findings, and motivates our proposal of ToneDMAC.

We consider directional MAC protocols that are motivated from the basic structure of 802.11. While deafness may not be a problem for TDMA-based protocols, we confine ourselves to the regime of 802.11-like CSMA/CA protocols only. We consider the design choices for the following protocol steps:

(1) *Carrier Sensing*

(2) *Backing Off*

(3) *RTS/CTS Handshake*

(4) *Data/ACK Exchange*

For each of the above steps, we discuss the ramifications of using the directional versus omnidirectional modes of operation. We identify how existing protocols adopt some of these design choices, and argue why they may not be sufficient to address deafness. We will refer to the scenario in Figure 4.2 for examples and explanations. In the figure, we will assume that node B intends to transmit data to node A.

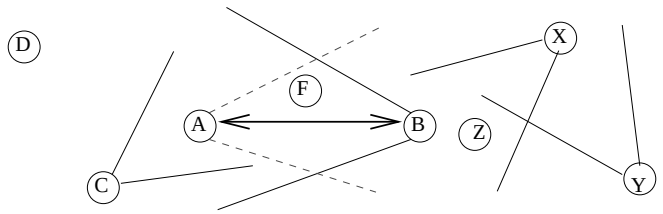

Figure 4.2: An example scenario in multi-hop ad hoc networks.

Carrier Sensing

Omnidirectional carrier sensing does not seem appropriate, when beamforming antennas are available. This is because node B (in Figure 4.2) would sense a transmission from X to Z, and defer its own transmission to A, even though the two transmissions are in different directions, and are not in conflict with each other. Clearly, this is unnecessary. However, if B sensed the carrier directionally, it would not sense the signal from X to Z, and could simultaneously begin communication to A. Simultaneous transmissions between B-A and X-Z improves network throughput.

Backing Off

Based on the arguments on carrier sensing, it appears that a node must remain in the directional mode, while backing off as well. Remaining in the directional mode can prevent a node from getting unnecessarily captured by surrounding communications. However, an acute problem can arise. A node B, that backs off in the directional mode (and communicates subsequent data packets directionally as well) would perpetually remain in the directional mode while it has packets queued for transmission. Put differently, if each packet requires node B to carrier-sense, backoff, and communicate directionally, then multiple back-to-back packets would require B to be in the directional mode for a continued period of time. This implies that a neighbor of node B, say X, intending to communicate to B, would experience complete deafness over this entire duration. Notice that in attempting multiple transmissions, X would also remain in the directional mode as long as it has more packets to transmit. A neighbor of X, say Y, intending to communicate to X, would similarly experience acute deafness from X. A chain is possible in which none of the nodes, except the first one, can communicate successfully – a "deadlock". DMAC [46] and the protocol proposed in [17] propose directional backoff and fail to identify the possibility of "deadlocks".

RTS/CTS Handshake

Many variations have been proposed for exchanging RTS/CTS with directional antennas. We discuss the ramifications of these variations in view of the deafness problem.

(i) Omnidirectional RTS/CTS was proposed by early protocols, and directional capabilities

were utilized only for Data/ACK packets. Transmitting omnidirectional RTS/CTS limits the communication range of beamforming antennas. Node B may no longer communicate directly to node A, otherwise possible when using directional RTS/CTS transmissions. In addition, omnidirectional RTS/CTS packets may reduce the benefits of spatial reuse. Referring to Figure 4.3, assume that nodes B and F perform omni RTS/CTS and begin communicating (the circles depict omni communication range). To protect this ongoing communication, Node Z must not initiate transmissions that may interfere at the receiver F. Notice that an omnidirectional communication from Z (meant for X) may interfere with F's reception, and is therefore not appropriate.

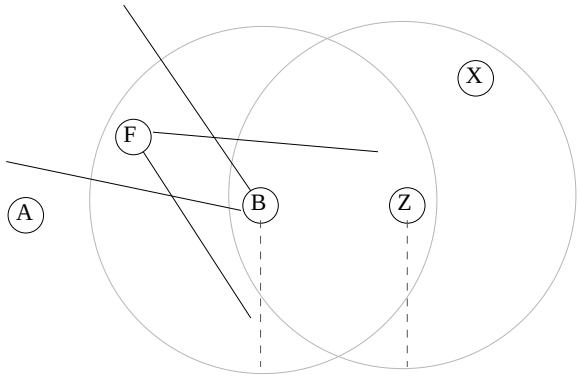

Figure 4.3: Omnidirectional transmissions affect spatial reuse.

(ii) Directional RTS/CTS [16] [42], [46], [17] allows for long-range, simultaneous transmissions, but suffer the adversities of deafness. Referring to Figure 4.4, if node X initiates communication to B while B is communicating to A, then X would experience deafness, as explained earlier. The cumulative impact of deafness would be clear when we discuss the issues related to directional Data/ACK packets.

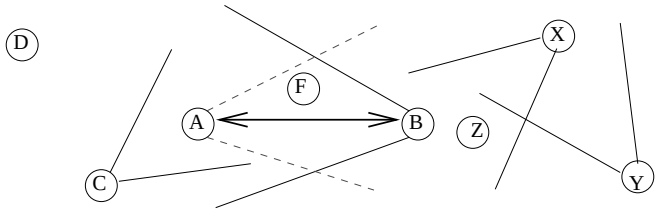

Figure 4.4: An example scenario in multi-hop ad hoc networks.

(iii) Multiple Directional RTS/CTS: Another option is to transmit multiple DRTS/DCTS packets (one on each beam) so that all neighbors of the transmitter and receiver are informed of the imminent dialog [43]. This strategy can eliminate deafness, while offering spatial reuse and range extension. However, the overhead can be extremely high if each communication needs to be preceded by N control packets[4]. Circular-DMAC [43] is a proposal that adopts this design choice as a remedy to deafness, but incurs large overhead, as evident from our simulations later in this section.

The above discussion leads us to believe that single directional RTS/CTS is probably the appropriate design choice, even though deafness appears as a tradeoff to spatial reuse and range extension. ToneDMAC aims to explicitly resolve deafness, thereby retaining the benefits of reuse and extension.

Data/ACK Transmission

The tradeoffs associated with RTS/CTS, are associated to Data/ACK packets as well. Omnidirectional Data/ACK transmissions reduce communication range and spatial reuse. Sacrificing these benefits is probably inappropriate, and therefore proposals in [14], [16], [43], [27], [46], almost unanimously adopt directional Data/ACK transmissions. However, choosing to transmit the Data/ACK packets directionally leads to deafness as well.

While B and A (in Figure 4.4) are engaged in directional Data/ACK, node X may attempt multiple transmissions to B, each transmission preceded by an exponentially longer backoff duration. Now, once B finishes transmitting a packet to A, it must choose a new backoff value in preparation of transmitting the next queued packet. Since X has exponentially increased its backoff duration multiple times, it is highly probable that the backoff value B chooses for the next transmission, is smaller than X's remaining backoff duration. As a result, B can finish counting its backoff earlier than X, and proceed with the next directional Data/ACK communication. Observe that when X finally finishes counting down, B has initiated a new dialog (DRTS/DCTS/DData/DACK), caus-

[4]Observe that with 802.11b, the time to transmit RTS or CTS packets is around one-third of the total time for sending 512-Byte data packets. When number of beams, N, are of the order of six or more, overheads can be intolerable.

ing X to fail again. After a *threshold* number of unsuccessful attempts, X drops the packet, and informs the upper layer of a link failure. Higher layer protocols will initiate elaborate mechanisms to reroute packets, and trigger congestion control algorithms. Performance can thus deteriorate significantly.

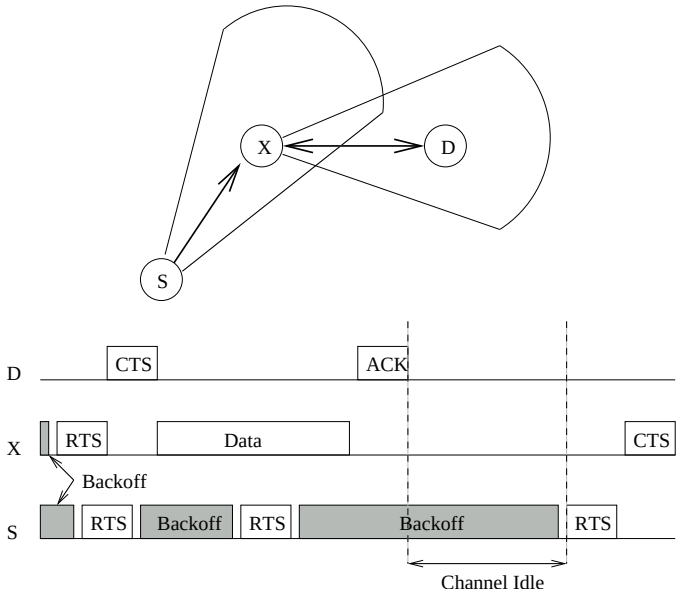

Figure 4.5: A scenario illustrating channel under-utilization due to deafness

Consider another consequence of deafness. Assume that node S intends to send packets to D using X as an intermediate node (Figure 4.5). When packets are large, X would remain engaged in transmission to D for a longer duration, during which S may attempt multiple retransmissions (each transmission preceded by an exponentially increasing backoff duration). When X finally completes packet delivery to D, and is ready to receive the next packet from S, it is likely that S would be counting down a large backoff counter. The channel remains idle over this entire duration (as shown earlier in Figure 4.5) – communication resumes only when the backoff period expires and S initiates the next transmission to X.

4.1.3 ToneDMAC – Protocol Description

To be able to handle deafness, a directional MAC protocol seems to require two main features.

1. The ability to communicate directionally (in the interest of higher spatial reuse)

2. The ability to inform all surrounding neighbors about a communication (in the interest of addressing deafness).

In addition, the above goals need to be achieved without a significant performance penalty.

Our proposal of ToneDMAC aims to achieve these design goals. In the beginning, a node performs directional communication (including DRTS/DCTS), and does not try to inform its neighbors of its intended communication. However, once the dialog is over, the node transmits an out-of-band tone to inform its neighbors that it had been engaged in communication over the recent past. Tones are transmitted omnidirectionally. On detecting the tone, a neighbor that had been failing to establish communication can deduce that the cause of previous failures was deafness (and not congestion). The node cancels its remaining backoff, chooses a new, relatively smaller, random backoff value, and reattempts transmission. The channel utilization increases, packet drops and transmission failure decreases, while spatial reuse and range extension is still maintained.

However, observe that tones do not contain information bits, and thus are not demodulated – they are only detected (through energy estimation) on the corresponding narrow frequency band. This implies that a neighbor may not be able to determine the sender of a tone, unless nodes are associated to specific tones. ToneDMAC proposes to hash node identifiers to tone-frequencies using a simple static hash function. The details of ToneDMAC are presented next.

ToneDMAC Carrier Sensing and Backoff

When using ToneDMAC, a transmitter (B) beamforms in the direction of its intended receiver (A) and carrier senses in the directional mode. If the channel is idle, the transmitter selects a backoff duration from an interval of $[0, CW_{min}]$ (similar to 802.11). Now, ToneDMAC requires B to switch the antenna to the omnidirectional mode while performing the backoff countdown. However, a node does not cancel its backoff on sensing any arbitrary signal. Instead, when the node senses a signal, it performs the azimuthal beam scan to determine the direction of the arriving signal. If this direction is in a different direction from that of A, then B continues with the countdown; otherwise B cancels the timer until the channel becomes idle. Observe that if the arriving signal is strong enough for B to receive it, then B cancels the backoff timer, and decodes the packet. If the packet is meant for B, then B can reply to the packet, alleviating the possibility of deadlock and deafness.

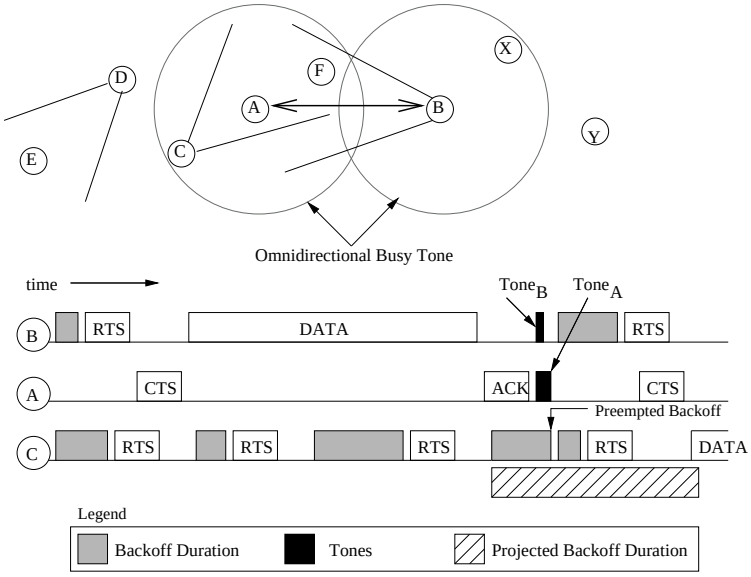

Figure 4.6: An example scenario and timeline of the ToneDMAC protocol.

ToneDMAC DRTS/DCTS/DData/DACK

The RTS/CTS handshake is carried out directionally when using ToneDMAC. Once the handshake is established successfully, Data and ACK packets are transmitted directionally as well. Nodes that overhear the DRTS, DCTS, or DData packets, update their records suitably to avoid interfering with communications in those directions. Both B and A switch back to the omnidirectional mode once the DData/DACK exchange has been accomplished.

ToneDMAC Tone Transmission

In ToneDMAC, the channel is split into two sub-channels: a data channel and a narrow control channel. RTS, CTS, Data, and ACK packets are transmitted on the data channel. The tones (essentially sinusoids with sufficient spectral separation) are assigned to the control channel. Once DData/DACK is over, ToneDMAC requires B and A to transmit tones omnidirectionally. The frequency of the tone, and the duration for which the tone is transmitted depends on the transmitter of the tone. A node i chooses a tone of frequency f_i from a set of K frequencies, and an integer time duration t_i from an interval $[t_{min}, T]$. The tuple (f_i,t_i) serves as the signature of node i. This allows a neighbor of node i, that detects the tone, to deduce the transmitter of the tone with some

degree of accuracy. The values of f_i and t_i are static hash functions of the node's identifier. We assume that a higher layer is capable of assigning consecutive identifier's to nodes, so that the tuple (f_i,t_i) can be uniformly distributed over the node identifier space. A simple hash function is used to assign (f_i,t_i) to a node i.

$$f_i = (i \bmod K) + 1$$

$$t_i = (i \bmod (T - 1)) + 2$$

Nodes B and A transmit their corresponding tones omnidirectionally. Since transmission in the omnidirectional mode may not reach all the directional neighbors, the transmit power of the tones are suitably increased to equal the range of directional transmissions. This ensures that nodes that may potentially be encountering deafness from B and A, are within the transmission range of the tones.

ToneDMAC Tone Reception

A neighbor of A, say C in Figure 4.6, that has been not been able to communicate to A because of A's communication to B, would receive the tone transmitted by A. To be able to conclude deafness as the cause of communication failure, C must first check whether the originator of the tone matches its intended receiver. Therefore C applies the static hash function on its intended receiver, calculates the *(f,t)* signature, and compares it with the frequency and duration of the received tone. If the comparison is positive, C concludes that the reason for communication failure was indeed deafness. In order to avoid backing off unnecessarily, it preempts its current backoff timer, reselects a new smaller backoff from $[0, CW_{min}]$ (as in 802.11), and reattempts transmission. Figure 4.6 shows the sequence of events where C successfully initiates transmission before B can initiate the second DRTS to A. Observe in the figure that, without the tones, C's projected backoff could be longer than B's second backoff, and therefore C would experienced deafness repeatedly.

Tone Aliasing

Observe that the hashing mechanism proposed for ToneDMAC does not guarantee that the trans-

mitter of the tone will always be deduced accurately. Aliasing occurs when two neighbors of a node are hashed to the same *(f,t)* signature (e.g., $f_D = f_A$, and $t_D = t_A$). A node like C, that has both A and D as its neighbors, will not be able to deduce the tone-sender correctly. We alleviate this problem by utilizing the direction in which the tones are received [47].

Another case that can cause incorrect deduction is when two nodes, hashed to the same tone-frequency, transmit tones that overlap in time. For example, assume that $f_D = f_A = f$, and $t_D \neq t_A$. It is possible that C would receive a tone at frequency f for a duration, τ, such that $\tau > t_D + t_A$. In such cases, an incorrect deduction occurs.

Although incorrect deduction is possible, we argue that the consequences may not be detrimental [47].

Performance Evaluation

We use the Qualnet simulator [38], version 3.6, for simulating ToneDMAC and comparing its performance with DMAC [46] and circular-DMAC [43]. We consider several scenarios and systematically analyze the impact of deafness on directional medium access control. The assumed data rate is 11 Mbps, with packet sizes of 512 bytes. We use both UDP and TCP traffic in our experiments. Sources are always backlogged unless mentioned otherwise. We do not consider mobility.

We begin by using three simple scenarios in order to understand the individual benefits of each design choice made in ToneDMAC. The scenarios are shown in Figure 4.7 – dotted lines indicate flows and solid lines indicate links. We briefly explain our findings for these three simple scenarios.

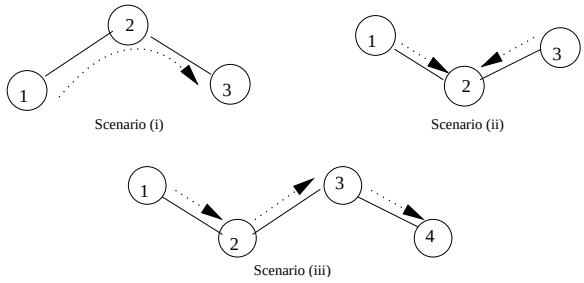

Figure 4.7: Scenarios affected by deafness.

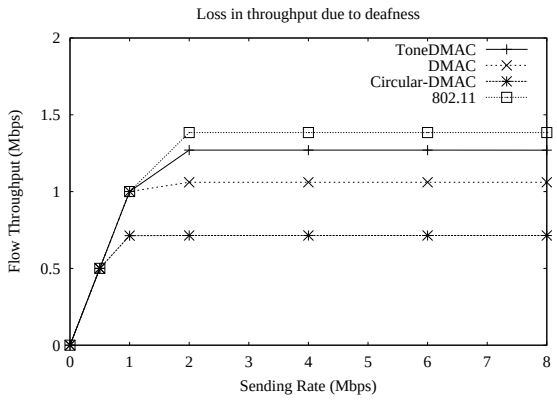

Figure 4.8: Throughput comparison for scenario (i).

Scenario (i)

Scenario (i) shows a two-hop flow between node 1 to 3 via node 2. As discussed earlier, node 1 experiences deafness when node 2 forwards the packet to node 3, which in turn results in channel under-utilization. Figure 4.8 quantifies the degradation in throughput due to deafness in scenario (i)[5].

Scenario (ii)

Scenario (ii) is designed to measure the cascading effects of deafness in terms of packet losses, delay variances, etc. Both nodes 1 and 3 intend to communicate to 2. Now, when one of the nodes (say node 3) initiates successful communication, the other would experience deafness. Due to the impact of deafness, node 3 continues to win channel contention for a long duration, while node 1 remains unfortunate with large backoff values. However, once node 1 gets lucky enough to steal channel access from 3 (possibly after dropping many packets), node 3 begins to suffer from sustained deafness. The alternation continues – the winning node transmits multiple packets once it grabs the channel, while the loosing node drops multiple packets. Figures 4.9 and 4.10 show the variation of backoff windows (and successful transmissions) with DMAC and ToneDMAC. Clearly, backoff values remain low, and channel access fairness is reasonably maintained with ToneDMAC.

Due to the alternation of channel access, delays experienced by nodes 1 and 3 can vary signifi-

[5]Note that scenario (i) measures only the losses suffered from deafness – we do not expect potential benefits of directional communication because there is no scope for spatial reuse.

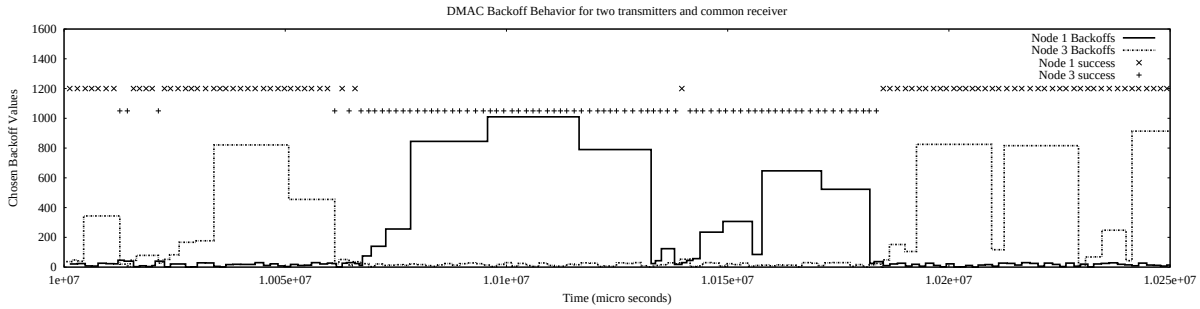

Figure 4.9: Variation of backoff values for scenario (ii) using DMAC.

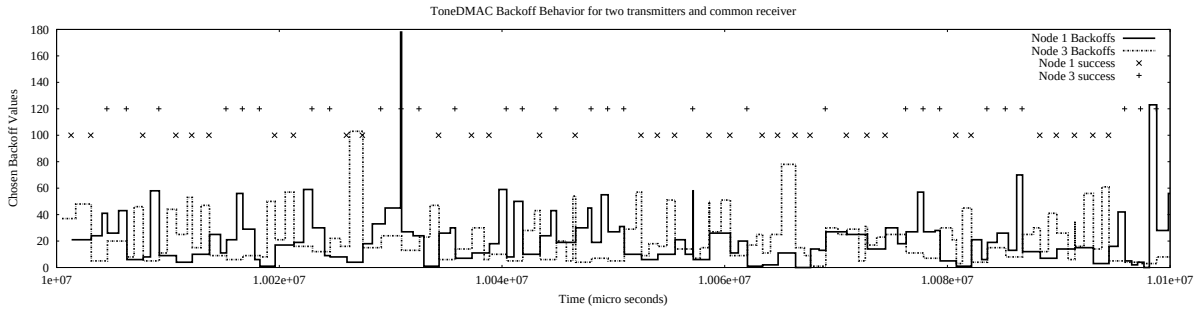

Figure 4.10: Variation of backoff values for scenario (ii) using ToneDMAC.

cantly in scenario (ii). A node that loses channel access to another node, might have to wait for a long duration before it can re-establish successful communication – a case of short-termed unfairness. Although the "average" delay over all the packets is comparable, "variance" in packet delays is often a metric of interest (to transport protocols like TCP, for example). Figure 4.11 compares the per-packet delay experienced by node 1 when using DMAC and ToneDMAC. While DMAC's delay is characterized by large fluctuations, ToneDMAC maintains a smaller range of per-packet delay. This reduces the "variance", which in turn improves TCP throughput (presented in [47]).

Scenario (iii)

Scenario (iii) aims to evaluate the "deadlock" problem discussed earlier in this section. Recall that "deadlock" occurs if the directional MAC protocol chooses to backoff in the directional mode, as proposed in DMAC [46]. As explained earlier, all nodes on a chain may fail to communicate except for the first node (node 3 in this case). However, when using ToneDMAC, nodes would apply the combination of *omnidirectional backoff* and *out-of-band tones* to handle deafness. To understand the individual impact of each of these factors, we design a variation of ToneDMAC,

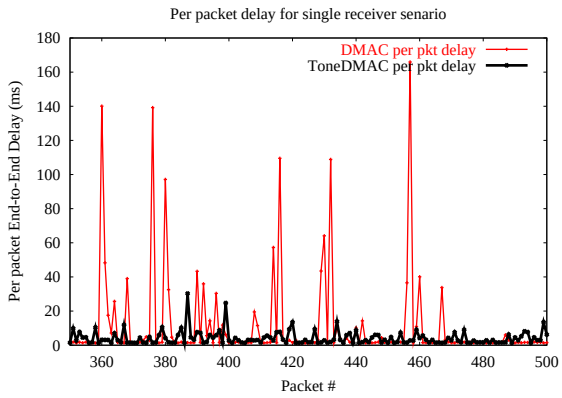

Figure 4.11: Comparison of delay for scenario (ii).

called *ZeroToneDMAC*, which does not use tones at all. In other words, ZeroToneDMAC is a design to reflect the individual benefits of omnidirectional backing off.

We quantify the impact of deadlock by measuring the aggregate packet drops in the network. Due to "deadlock", DMAC experiences severe rate of packet drops at nodes 1 and 2. However, as evident from Figure 4.12, ToneDMAC exhibits far less packet drops, proving the efficacy of tones to avoid deafness.

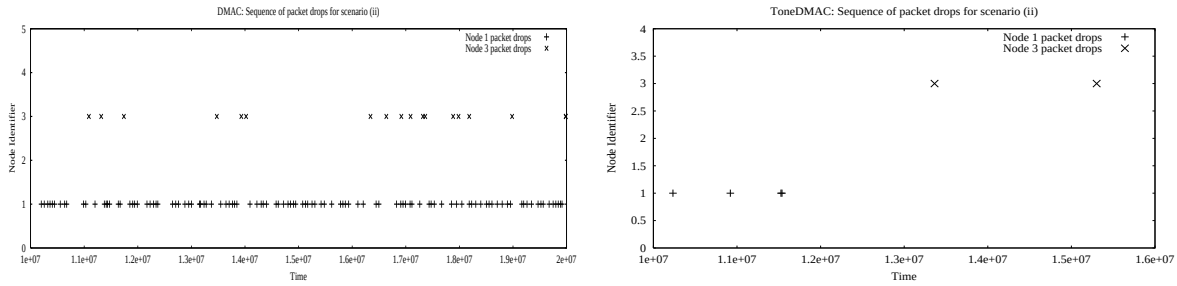

(a) Packet drops with DMAC

(b) Packet drops with ToneDMAC

Figure 4.12: Packet drop due to deafness in DMAC and ToneDMAC

Chain and Grid Topologies

With reasonable understanding of how deafness impacts the simple scenarios, we simulate chain and grid topologies. Figure 4.13(a) shows the throughput achieved by ToneDMAC in comparison to DMAC. Observe that ToneDMAC does not exhibit benefits – in a chain of nodes X, Y, and Z, while Y and Z are communicating, node X is aware of the communication because it can receive the

CTS from Z to Y. However, when we simulate grid topologies, the improvement with ToneDMAC becomes evident. Since grid topologies are characterized by scenarios such as in Figure 4.13(b), ToneDMAC achieves higher throughput in contrast to DMAC. Figure 4.13 quantifies the improvements – clearly, ToneDMAC outperforms 802.11 as well as DMAC.

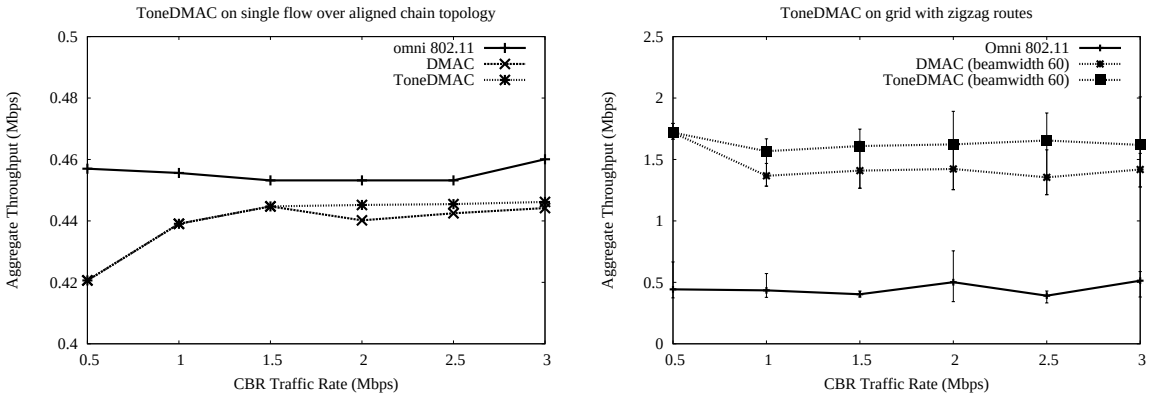

(a) ToneDMAC does not benefit aligned chain topologies

(b) ToneDMAC improves the throughput in grid topologies

Figure 4.13: Performance of ToneDMAC in chain and grid topologies

Random Topologies

We now experiment with random topologies. 50 nodes are placed randomly in a region of 1000 X 1000 square-meters. Source-destination pairs are selected randomly, and routes are statically assigned to each of the flows. We use $K = 4$ and $T = 4$ (i.e., 4 tone-frequencies and maximum 4 slots for transmission) for ToneDMAC. The simulation results are averaged over 15 runs. As described earlier, we include results for ZeroToneDMAC in order to discern the benefits of *tones* from *omnidirectional backing off*.

Figure 4.14 shows the performance of ToneDMAC when UDP traffic is used on a specific random topology. Evident from the graphs, ToneDMAC outperforms both DMAC, ZeroToneDMAC, and Circular-DMAC mainly because of its ability to reduce channel-idle time and resolve deadlocks. DMAC performs worse than ZeroToneDMAC and ToneDMAC because of its inability to address ei-

ther of the issues above. The benefit of omnidirectional backing off is evident from the performance improvement of ZeroToneDMAC over DMAC. The benefit of tones, to reduce idle-channel time, is evident from the improvement of ToneDMAC over ZeroToneDMAC. Circular-DMAC achieves the least throughput not only because of the multiple RTS transmissions, but also because deafness appears due to the transmission of a single directional CTS packet (as proposed in [43], to reduce overheads). The performance over multiple random topologies are shown in Figure 4.15. As evident from the figure, ToneDMAC consistently achieves higher throughput in comparison to DMAC.

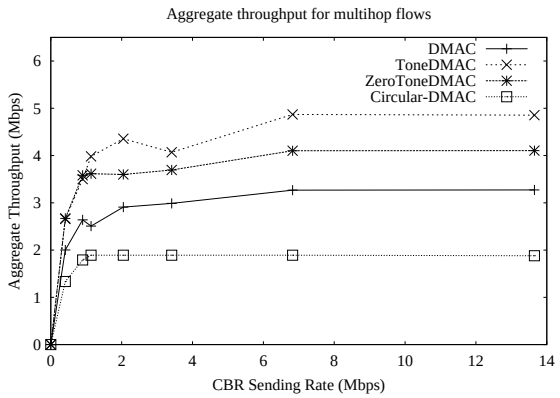

Figure 4.14: Throughput in a random network for increasing traffic.

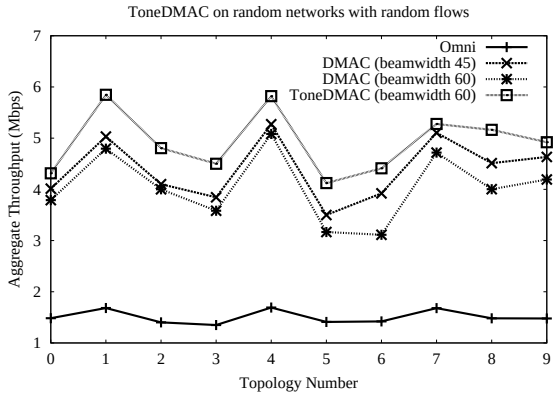

Figure 4.15: Throughput comparison for different random networks.

The outcome of mitigating deafness may provide additional advantages when network layer issues are considered. For example, fewer packet-drops may trigger fewer route errors. Suboptimal route selection, caused due to undelivered route-discovery packets, affects end-to-end throughput

ToneDMAC	ZeroTone	DMAC	C-DMAC
2.64	1.82	1.41	1.23

Table 4.1: Aggregate TCP multihop throughput

and delay [48]. ToneDMAC may alleviate such problems. The results presented in this section do not reflect such potential performance gains. We believe that when a suitable routing protocol is executed, the performance gains of ToneDMAC would be even more pronounced.

While evaluating scenario (ii) in Figure 4.7, we observe large fluctuations in delay for DMAC transmissions. To evaluate the impact of high fluctuations (causing higher variance), we simulate flows with TCP traffic. As shown in Table I, the aggregate throughput achieved by ToneDMAC is higher than ZeroToneDMAC, which in turn is higher than DMAC and Circular-DMAC. This happens because the round trip time (RTT) estimation of TCP depends on the variance in end-to-end transmission latency. When using DMAC, the higher variance and recurring packet-drops degrade system performance. Tone-based feedbacks reasonably limit delay-fluctuations, in addition to increasing channel utilization. The net benefits are encouraging.

Discussion

Most directional MAC protocols, including ToneDMAC, have adopted the CSMA/CA principles of IEEE 802.11. However, it is unclear that modifying 802.11 is the optimal approach to designing directonal protocols. While binary exponential backoff (BEB) may be suitable for omnidirectional antennas, such conservative backing off may not be the best choice for beamforming antennas. We intend to design new backing off schemes as a part of our future work.

An issue with ToneDMAC pertains to successful tone detection, especially in the presence of fading and multipath. Since tones are transmitted on a narrow bandwidth channel, and the durations of transmissions are short, tones may be detected partially, or may not be detected at all. Multipath factors can also cause a tone to arrive from a different direction than the known direction of a neighbor. Clearly, both these effects can cause a node to misclassify the cause of transmission failure. However, notice that the consequences of misclassification may not be severe – a node may initiate its RTS too early or transmit its RTS too late. Performance penalties associated to these

consequences do not appear to be severe.

Omnidirectional backing off, as opposed to directional, has a subtle tradeoff. While in the omnidirectional mode, a node, S, may not be able to overhear an RTS/CTS packet from a distant communication. If the node later beamforms in the direction of this communication, and initiates an RTS, then in some cases collisions may occur. However, such cases are rare since the receiver of the distant communication is not only outside the omnidirectional range of S, but also points its beam in the direction of its transmitter. As a result, the signal to interference plus noise ratio (SINR) at the receiver is likely to remain sufficiently high for successful reception.

Summary

We investigated the problem of deafness, that arise when MAC protocols are designed for beamforming antennas. We evaluated several cases in which deafness can lead to channel under-utilization. We also identified scenarios in which the impact can be cascading, leading to packet drops, delay fluctuations, unfairness, etc. Our proposal of ToneDMAC uses a tone-based notification mechanism that allows neighbors of a node to classify congestion from deafness, and react appropriately. Simulation results show that ToneDMAC can be effective in mitigating the adverse effects of deafness, while retaining the benefits of directional beamforming.

4.2 MAC-Layer Capture

Beamforming antennas have been shown to improve spatial reuse in wireless networks. Protocols that aim to exploit beamforming antennas have leveraged benefits from directional transmission as well as directional reception. Yet, the improvements from such protocols have not met the analytical expectations. We argue that utilizing antenna characteristics for transmission and reception alone, is not sufficient. Controlling the antenna during the idle state of a node also plays an important role in attaining the expected benefits. For example, a node that does not control its antenna during the idle state, will be prone to receiving every (sufficiently strong) signal that impinges on its transceiver. Observe that many of these signals will not be intended for the node, and the

received packet will eventually be dropped at the MAC layer. We term this problem as *"capture"*, to convey a node's inability to initiate useful communication while it is engaged (or *captured*) in receiving an unproductive packet.

This section presents a MAC protocol that utilizes beamforming antennas to minimize the impact of capture. The MAC protocol characterizes traffic patterns based on their directions, and prescribes suitable antenna patterns that can spatially filter out unproductive signals. We show later that the network layer might be able to assist the MAC layer by using a new *capture-aware routing metric* that prefers routes less affected by capture. Simulation results show that reducing the impact of capture can improve the achievable performance of wireless multihop networks, such as mesh networks.

4.2.1 Problem Formulation

Theoretical results have shown that directional antennas have the potential to improve spatial reuse of a wireless network by a factor of $O(\frac{2\pi}{\sqrt{(\alpha\beta)}})$ [21], where α and β are the beamwidths of the transmit and receive beams. A study of the existing directional protocols shows that almost all of them leverage substantial benefits by transmitting and receiving directionally. However, the improvements still remain below the achievable limits, the reasons for which are intuitively argued below.

Observe that transmission and reception can be associated to specific directions toward which a node must point its beam. In contrast, there is no particular direction associated to the idle state of a node. While in the idle state, a node remains unaware of directions from which a subsequent packet might arrive. In the absence of such information, most directional antenna protocols require idle nodes to be in the omnidirectional mode. Existence of such omnidirectional beam-patterns, in an otherwise directional protocol, can limit the achievable spatial reuse of the network. This is because the omnidirectional mode is severely vulnerable to a problem we refer to as "capture".

Definition of Capture

We define "capture" as a problem in which idle nodes get engaged in receiving an unproductive packet, and thereby cannot utilize that time duration for its own (useful) communication. Capture is a prominent problem in CSMA protocols because nodes are expected to remain idle for long durations (e.g, while backing off, waiting for packet arrival, etc.). The problem is most pronounced when these idle nodes are required to be omnidirectional. In such a case, the benefits of directional beamforming cannot be fully leveraged. The example in Figure 4.16 illustrates the point.

In Figure 4.16, assume that node A intends to transmit to node B, and node C to node D. Capacity analysis assumes that node D has the ability to directionally receive from C, while node A is transmitting to node B. This form of directional transmission and reception intuitively leads to the benefits of $O(\sqrt{\frac{2\pi}{\alpha}})$ and $O(\sqrt{\frac{2\pi}{\beta}})$, respectively, over omnidirectional antennas. While designing CSMA protocols for beamforming antennas, such parallelism may not be easily achieved. Since node D is required to remain omnidirectional (while idle), it gets captured by the transmission from node A. Observe that if node C initiates transmission first, node B will get captured; in neither case will the two dialogs proceed simultaneously. Clearly, capture limits the spatial reuse ideally achievable with directional antenna protocols.

In addition to affecting the spatial reuse, capture also leads to deafness. Observe that when node C initiates communication to node D, the latter remains captured by the ongoing dialog between A and B. As a result, node D is unable to receive the message from node C, and thereby does not reply to it. Clearly, node C experiences deafness. Since several nodes lie within the interference footprint of an ongoing communication, they get captured – this increases the occurrence of deafness in the environment. Alleviating capture not only improves the spatial reuse of the network, but also reduces deafness by a non-marginal amount.

Addressing Capture – Design Choices

Let us consider some design choices to address the problem of capture.

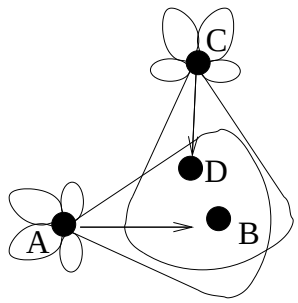

Figure 4.16: Capture prevents parallel dialog between A-B and C-D

(1) *Adaptive beamforming* at the physical layer has been shown to be effective in minimizing interference at the receiver. However, such beamforming techniques may often lock on to an unproductive signal (assuming that it is the signal-of-interest, SoI) and receive it completely. Observe that the signal-of-interest at the physical layer is not necessarily the packet-of-interest (PoI) at the MAC layer. Since the physical layer does not recognize the notion of a PoI, it may not be equipped to mitigate our notion of MAC-layer capture.
More complex techniques, such as *multi-user detection, etc.*, are capable of distinguishing between a productive and an unproductive packet. However, such techniques incur significant processing delay and cost, deeming them impractical for many 802.11-based mesh networking applications.

(2) *Pure TDMA based protocols* might appear to be a candidate to mitigate capture. In Figure 4.16, if node D is aware of a scheduled transmission from node C, then node D might beamform *a priori* toward node C, thereby avoiding capture from A. However, in multihop networks, distributed scheduling has been shown to be difficult, especially when traffic patterns are changing with time. The complexities only increase when employing directional scheduling. While several other approaches appear to alleviate capture, the penalties they incur deem them impractical.

(3) We propose to address capture by combining partial ideas from TDMA schemes with a directional CSMA protocol. The key idea is to perform *MAC-layer beam control*, assisted by capture-aware route selection from the network layer. More specifically, when the MAC layer observes persistent capture traffic from a particular direction, it prescribes a beam-pattern that filters out signals from the corresponding direction. However, such a beam-pattern cannot be used

indefinitely, in view of the possibility that productive traffic might arrive from around the same direction sometime in future. This leads to the design of a MAC protocol that must *multiplex between omnidirectional and directional modes even while a node is idle.* The network layer may further assists the MAC layer by choosing routes on which links are less susceptible to capture. The contributions of our work is summarized as follows.

- We identify the problem of *capture* as a limiting factor to higher spatial reuse in existing directional antenna protocols[6].

- We propose a MAC protocol that implements a beam selection policy to reduce the impact of capture and deafness. In addition, we suggest an optimization called *Early Destination Detection*, wherein the *destination field* in a packet is detected early to avoid capture on small time scales.

- We propose a routing protocol that uses a *capture-aware* metric to select routes on which links are less prone to capture. This results in high-throughput routes since each route is less "distracted" by traffic from other routes. We show that capture-awareness at the MAC and network layer can complement each other, enhancing network performance.

Related Work

The benefits of using beamforming antennas have been explored at several layers of the protocol stack. At the physical layer, signal processing algorithms have been proposed for adaptive beamforming [2], [49]. Smart array antennas, that exploit these algorithms, are rapidly appearing in the commercial market [50], [51], and offer considerable gain in terms of link quality. At the MAC layer, several protocols have been designed that aim to exploit directional capabilities [14], [27], [18], [46], [47], [52]. Neighbor discovery, topology control, and multicast have been studied in [53], [37], [54]. At the network layer, beamforming antennas have been shown to find fewer

[6]Observe that when using omnidirectional antennas, the motivation to mitigate capture is limited because captured nodes are anyway not expected to initiate simultaneous communication.

hop routes [48], energy-efficient routes [15] and interference-aware routes [55]. Several researchers have also investigated the use of multi-beam antennas [56], [54]. Some projects have also initiated testbed implementations toward utilizing directional antennas for multihop networks [57], [58]. We discuss some of the existing work from this rich body of research, and explain how our contributions improve the state-of-the-art.

The early work on directional MAC protocols improve spatial reuse by enabling nearby communications to progress in parallel [14], [27]. Further investigation introduces the notion of directional virtual carrier sensing through the directional NAV (DNAV) mechanism [18]. However, several tradeoffs are also identified, including new kinds of hidden terminal problems, and deafness [46], that limit the benefits of directional antennas in multihop scenarios. Subsequent protocols address some of these problems and raise the limits on achievable performance [43], [47], [52].

The impact of beamforming antennas on multihop routing has also emerged to be of interest [37], [55]. The longer transmission range of beamforming antennas holds promise of discovering shorter routes, in addition to the prospect of localizing queries only toward the direction of the destination [48]. Moreover, the ability to control the topology using directional antennas [37] makes it a candidate for selecting less congested routes [15], [55]. From the perspective of energy awareness, directional antennas are also shown to be beneficial for routing, broadcast, and multicast traffic [54], [55].

In a more relevant work [26], the authors consider the possibility of *zone disjoint routing*, wherein, routes are chosen such that they are spatially disjoint. More precisely, a source prefers routes in which intermediate nodes are least "loaded" by traffic from other sources. However, we notice that the intermediate nodes can still be susceptible to capture, because they may be located near another "loaded" node. Our approach in this dissertation is motivated from the observation that nearby nodes in a network should be capable of communicating simultaneously, as long as their directions of communications are non-overlapping. As will be clear later in this section, our protocol encompasses the benefits available from *zone disjoint routing*, and enhances it with

capture-awareness.

The opportunity to enhance spatial reuse, by eliminating capture, has not been identified with existing directional protocols. Of course, the notion of capture has been extensively studied in physical layer research, where capture refers to the successful demodulation of the stronger signal, when two signals collide at a common receiver [59]. However, as will be clear later, *our notion of "capture" applies more to a MAC-layer phenomenon (and does not necessarily involve the incidence of multiple signals).* By mitigating this form of capture, from both the MAC and network layer, we demonstrate the possibility of performance improvement.

Motivation and Intuition

Most CSMA-based directional MAC protocols require an idle node to be in the omnidirectional mode [14]– [46], [17], [47]. This is because when using CSMA/CA protocols like IEEE 802.11, an idle node is not expected to know the direction from which signals may arrive. Once a signal impinges on the receiver, different protocols assume different techniques to beamform toward the signal. Once beamformed, the rest of the dialog is typically accomplished in the directional mode. When the dialog is over, the node returns to the omnidirectional mode, waiting for the next packet. Observe that this model of operation lends itself to capture. *While an idle node is in the omnidirectional mode, it is susceptible to get captured by every (sufficiently strong) unproductive signal that impinges upon its transceiver.*

This observation motivates protocols that are "capture-aware". A capture-aware MAC protocol can employ an appropriate beam-pattern (a spatial filter) that filters out unproductive signals, periodically. Nodes may now be able to participate in productive communication while the spatial filter is active. To respect the possibility that productive signals may arrive from the filtered-out directions, a node must refrain from using the filter for indefinitely long periods of time. Although alleviated, capture is certainly not eliminated by using such a scheme. Observe that a node will continue to get captured if both productive and capturing signals arrive from around the same

direction, in which case, the MAC layer will have to live with the problem.

The network layer might be capable of resolving this problem. A node that observes oncoming capture traffic from a certain direction, may avoid selecting routes through that direction. Introducing angular separation between directions of useful-communication and capture allows the MAC layer to filter out the capture-causing signals with higher efficacy. As a result, routes admitted to the network are less "distracted" by cross traffic from other routes. Motivated by these observations, we propose the design of capture-aware MAC and routing protocols, in this and the next chapter, respectively.

4.2.2 Capture-Aware Directional MAC (CaDMAC)

We divide time into recurring cycles of duration T_{cycle} as shown in Figure 4.17. Each cycle is further subdivided into an *ON-Duration* and an *OFF-Duration*, of lengths τ_{on} and τ_{off}, respectively. For our proposed protocol, we assume that nodes are clock-synchronized (we show later that synchronization need not be tight, and can be within the limits of current accuracy levels). Nodes enter the ON and OFF durations synchronously, and the cycle repeats. The *ON-Duration* denotes the time period during which all nodes in the network have all their beams turned on, i.e., all nodes are in the omnidirectional mode. The *OFF-Duration* denotes the time period during which a subset of the beams may be turned off. Note that during the *OFF-Duration*, the transceiver is not switched off and the node continues communicating – only some of the antenna beams may be turned off (as detailed below).

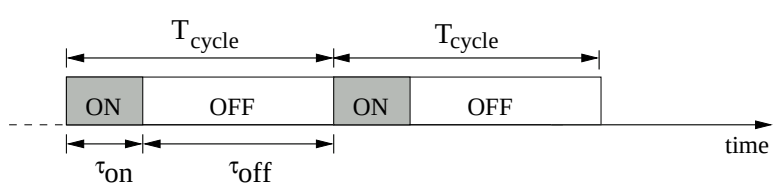

Figure 4.17: Time divided into cycles of ON and OFF-Durations

While in the ON-Duration, the MAC layer at each node passively monitors the ongoing traffic on a per-beam basis. Monitoring ongoing traffic involves segregating incoming packets into 2 cate-

gories – (i) productive traffic and (ii) capture traffic. Productive traffic at a node includes unicast packets meant for that node, as well as broadcast packets. Capture traffic at a node includes overheard packets that are intended for other nodes. If a beam proves to be the receiver of only capture traffic[7], then the beam is "black-listed". At the end of the ON-Duration, the MAC layer decides to turn off all the black-listed beams for the next OFF-Duration. The resulting beam-pattern used by node i, during the ensuing OFF-Duration, is denoted by β^i_{off}. The beam β^i_{off} can be viewed as a spatial filter that makes node i less susceptible to capture during the OFF-Duration. As a result, node i is left with more time to participate in its own productive communication. The example in Figure 4.18 illustrates this simple scheme.

In Figure 4.18, assume that node A communicates with node C via node B, while node E communicates with node G via node F. While in the ON-Duration, nodes identify their respective beams that must be turned off. As shown in Figure 4.18(a), node B turns off its north and south pointing beams, while node F turns off the east and west pointing beams. As a result, when B communicates with nodes A and C, node F also communicates with E and G, using their respective β_{off} beam-patterns. Once the nodes become idle, they do not switch back to their omnidirectional modes. Instead, they continue to listen to the medium using their respective β_{off} beams, waiting for the next packet to arrive. Spatial reuse increases in this simple scenario.

Observe that the beam patterns shown in Figure 4.18(a) will have sidelobes and backlobes. However, parallel communications will still be achievable, provided F and B are not too close to interfere each other significantly, along the direction of low-gain side lobes.

The above benefits of spatial reuse in CaDMAC arises from the ability of a node to use beam-pattern β_{off} while idle. However, using the same beam-pattern for communication may not be efficient. Consider Figure 4.19. If node C initiates a flow destined to node D using β^C_{off}, then the communication between C-D and A-B cannot occur simultaneously. This is because node C's transmission to D captures node B, and prevents the latter from simultaneously communicating to A. The vice versa happens when node B communicates to node A.

[7]Recall that each packet arrives at the MAC layer along with the beam-of-arrival information.

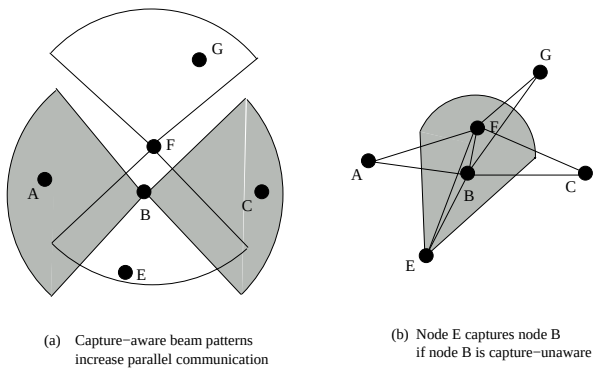

Figure 4.18: CaDMAC enables parallel communication by avoiding mutual capture between nearby nodes.

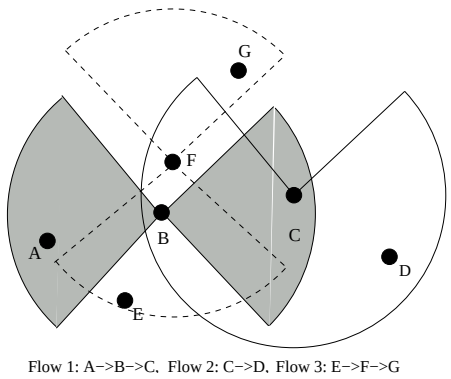

Figure 4.19: Communicating with beam β_{off} is inefficient (nodes B and C capture each other because of higher interference).

Observe that using a single beam for transmission can avoid such a problem (i.e., C points its conical mainlobe toward D while B points its conical mainlobe toward A). We adopt this feature in CaDMAC – a node always uses (both in ON and OFF durations) a single beam to communicate with its upstream and downstream nodes. Once a dialog is over, an idle node i returns to the beam-pattern β^i_{off} (pre-selected for that ongoing OFF-Duration), and waits for the next dialog to begin. Referring to Figure 4.20, notice that by controlling beam-patterns between busy and idle states, CaDMAC can support all 3 communications simultaneously, namely communications between B-A, C-D, and F-G (or E-F). Existing directional protocols such as DMAC [46], Circular-DMAC [43], etc. support at most two out of these three dialogs in parallel, while omnidirectional protocols support only one.

Directional virtual carrier sensing (DVCS) is a useful mechanism that requires a node to refrain

from initiating transmissions toward a direction from which it has overheard an RTS, CTS, or a Data packet. However, observe that the DVCS mechanism is not designed to address the problem of capture, because, even while implementing DVCS, a node will get captured by the RTS packet in the first place. More importantly, a node that overhears an RTS from a specific direction does not attempt to filter out the unwanted data packet that is likely to arrive from the same direction in the immediate future. When the unwanted data packet arrives, the node gets captured again, and thereby cannot schedule its own useful communication.

We enhance the DVCS mechanism to make it capture-aware. When a node, X, overhears an RTS or a CTS packet on a particular beam, CaDMAC recommends the physical layer to turn off that beam for the proposed duration of the imminent dialog. As a result, X can remain idle over the duration of the dialog, and can potentially initiate (or respond to) a new communication. Since the duration of dialogs can be long, the benefits from avoiding capture at such time scales can be non-marginal.

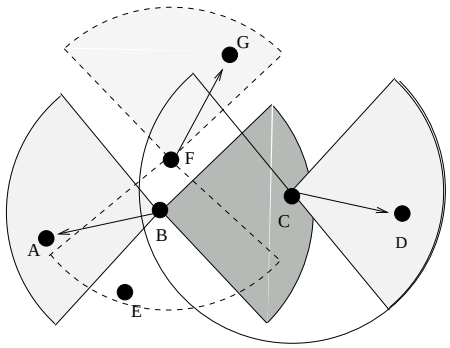

Figure 4.20: Maximal spatial reuse available from CaDMAC

CaDMAC does not assume that directions of signal arrival obey line-of-sight propagation. Observe that environments can be rich in multipath effects, and signals from a node can arrive after multiple reflections. CaDMAC only records the beams on which the strongest component of the signal arrives, and turns them off if they satisfy the criteria. We assume that network nodes are stationary and thus the multipath environment, although rich, does not change heavily over short

time scales of few seconds. This assumption has been validated in [58].

Early Destination Detection (EDD): An Optimization

CaDMAC avoids capture in two time-scales; (i) at the time-scale of OFF-Durations, by choosing suitable beam patterns, and (ii) at the time-scale of "RTS/CTS/DATA/ACK" dialogs, by extending the DVCS mechanism. However, capture is still possible at the time-scale of single packets, e.g., a node may get captured by a single RTS or CTS packet. More importantly, when RTS/CTS packets are not used, a node may be captured by the entire DATA packet. To minimize capture on the time-scale of single packets, we present an optimization that we call *Early Destination Detection (EDD).*

Observe that while a packet reception is in progress, if the MAC layer can quickly detect that the packet is not meant for it, then it might be possible to preempt reception of the rest of the packet. The beam-of-arrival for this packet can be turned off for the remaining duration of the packet (calculated using information from the packet header), thereby freeing the node from capture. The quicker the *destination detection*, the higher are the benefits (especially when packets are large). We incorporate this EDD optimization into CaDMAC. For this optimization, we need to ensure that a packet's destination address is included within the initial bits of the packet header, so that the decoder can decode this information and pass it to the MAC layer, before the entire packet is received. Support for such capabilities is already available in a transceiver prototype from Lucent [60].

CaDMAC's Dependence on Routing Protocol

While CaDMAC is able to mitigate capture, an important concern appears from the routing perspective. With reactive protocols such as DSR [31] or AODV [61], if a route discovery is initiated during the OFF-Duration, discovered routes may prove to be sub-optimal. In some cases, a route may not be discovered at all, even though a route exists (during the ON-Duration). This can happen because, during the OFF-Duration, the topology may be less strongly connected compared to the original topology (since nodes may turn off a subset of their beams, and thereby remove

links from the original topology).

We propose simple modifications to the routing protocol to address this issue. The routing protocol, even if antenna-unaware, could be made aware of the ON and OFF durations. Route discoveries initiated during the ON-Duration require no modification. However, if the route-discovery is initiated during the OFF-Duration, then the routing protocol sets a flag to remember this action. If a route is discovered during the OFF-Duration, the routing protocol begins sending packets on this route (even though the route might be sub-optimal). If a route is not discovered, it buffers the arriving packets and waits for the next ON-Duration. If the flag is set, at the beginning of the next ON-Duration, it re-initiates route discovery to find an optimal route, and redirects traffic over it. Of course, the latency of finding the optimal route increases on an average. Yet, network throughput and average end-to-end delay improves. This is because the gains from avoiding capture over subsequent OFF-Durations, compensate adequately for the losses from initial sub-optimality. Moreover, sub-optimal routes discovered during OFF-Durations are short-lived, and often not too bad in comparison to the optimal. To validate this, we evaluate the impact of ON and OFF durations on network performance later in this section.

A given routing protocol can be executed over CaDMAC with the above modifications. However, the routing protocol will remain capture-unaware. In the next chapter we illustrate that incorporating capture-awareness in the routing protocol can further improve performance. Our aim is to select routes based on a capture-aware metric, such that links on that route are less susceptible to capture. In the remainder of this chapter, we evaluate CaDMAC and quantify the benefits of capture-awareness.

Performance Evaluation

We use the Qualnet simulator, version 3.6, for simulating our proposed protocols, and comparing them with 802.11 and other beamforming antenna protocols. We consider several scenarios, and systematically analyze the factors that impact protocol performance. We use a transmission data

rate of 11 Mbps. We use a two-ray ground reflection model with Rayleigh fading. We use constant bit rate (CBR) traffic at each source node.

The antenna beams used for our simulations are characterized with sidelobes and backlobes, obtained from realistic antenna patterns [49]. We evaluate the performance of our protocol for varying beamwidth. We begin with an evaluation of our proposed capture-aware MAC protocol (CaDMAC), and evaluate the optimizations later. For our *Early Destination Detection (EDD)* optimization, we assume that the destination field of the packet is available to the MAC layer at the time when the PLCP and the physical layer header has been completely received.

Simple and Grid Topologies

We begin by using a simple scenario shown in Figure 4.21 – node A communicates with node C via node B, while node E communicates with node G via node F. Routes are specified statically for this simple scenario. The ON and OFF durations are assigned as 1 and 3 seconds respectively, and the beamwidth is specified to be 90°. We compare the end to end throughput of CaDMAC against DMAC [46] and Omnidirectional 802.11.

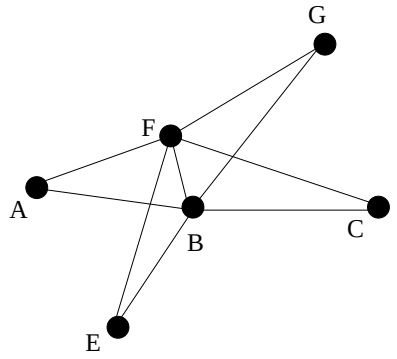

Figure 4.21: A simple scenario with 2 flows; A-B-C and E-F-G.

At low sending rates, the benefits of capture-awareness are not evident. As the sending rates increase, the network begins to saturate earlier with DMAC. This is because node B often gets captured by communications between nodes F-G or nodes E-F. Observe that node F also gets frequently captured by transmissions from nodes A and C. Capture limits the available spatial reuse in DMAC and Circular-DMAC, which translates into lower aggregate throughput. The difference

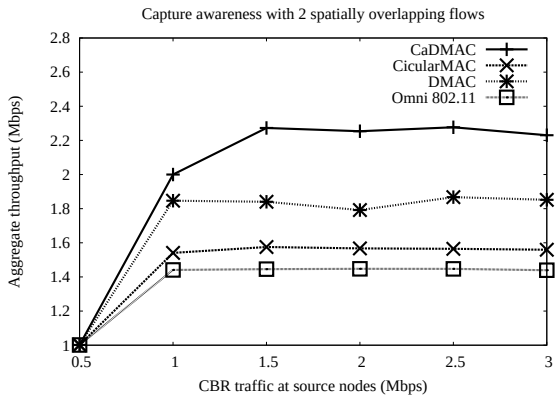

Figure 4.22: CaDMAC improves throughput for the simple scenario

in DMAC and CaDMAC graphs is an indicator of the benefits available from capture-awareness. We also compare the performance of CaDMAC for grid topologies. In grid topologies with aligned routes (i.e. horizontal or vertical routes), the performance of CaDMAC is comparable to DMAC. This is because there is not too much directional diversity with aligned routes, and as a result, beams that are utilized for communication, are also the ones that get captured by unproductive traffic. Consequently, CaDMAC cannot turn this beams off, leading to capture in several cases. We visit this issue in Chapter 5, and show how choosing routes using a capture-aware metric can mitigate this problem. However, when we assigned zigzag routes on the grid, we observed that CaDMAC improved over DMAC. The improvements are evident from Figure 4.23. The improvements result from CaDMAC's ability to turn off beams that are not used for communication, but are captured by unproductive communication.

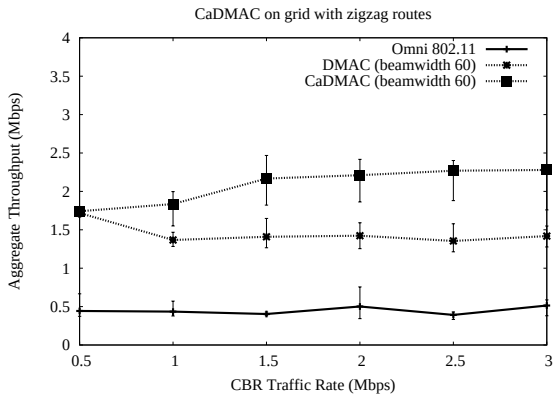

Figure 4.23: CaDMAC improves throughput for grid topologies with zigzag flows.

Large Random Scenarios

We simulate large networks, with 50 nodes randomly placed in a square region of side 1000m. We select random pairs of source and destination nodes, and routes are statically assigned to each pair. The routes are not selected to favor capture-aware routes, and thus, the graphs in Figure 4.24 reflect the MAC-layer benefits alone. The ON and OFF durations are specified to be 1 and 3 seconds respectively, while the beamwidth is 60°. The results are an average of 15 topologies.

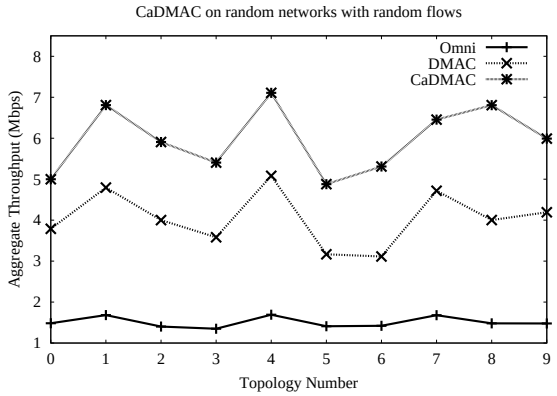

Figure 4.24: Throughput benefits from capture-aware MAC protocols in large random networks, with beamwidth of 60°

As evident from Figure 4.24, the relative benefits from CaDMAC are significantly high in several topologies. Recall that the improvement in CaDMAC is higher when the route patterns in the network are less aligned, and therefore a node can steer nulls toward unproductive directions while listening in productive directions. Since random topologies with congested traffic is likely to have this diversity, CaDMAC's improvements are expected. The EDD optimization with CaDMAC improves performance slightly. The slight improvement is because of using RTS/CTS packets, as well as small data packets (512 Bytes).

Delay in Random Networks

Figure 4.25 shows the impact of capture-awareness on end to end packet delivery latency. The latency improves because nodes are less captured, and thereby do not waste time in receiving unproductive packets from its neighbors. This "free" time is invested in productive communication, and therefore, the communication latency reduces.

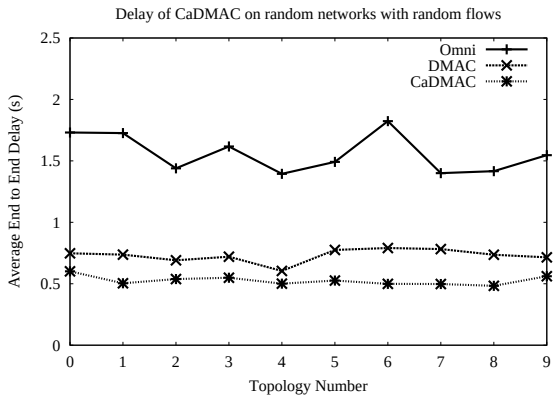

Figure 4.25: Reduction in end to end latency with capture-aware MAC protocols in random networks.

Discussion

Backhaul and Mesh Networks

Directional antenna protocols are proving to be an attractive design choice for wireless backhaul and mesh networks. Such networks are typically characterized by stationary nodes and persistent traffic. The stationarity of nodes justifies our assumption that the captured beams do not vary over short time scales. However, multipath effects may still cause variations, although such variations are typically over longer time scales. The variation in traffic patterns may also require the CaDMAC multi-beam patterns to change. However, for mesh networks, where the nodes are essentially traffic aggregation points, the traffic is not likely to observe fluctuations in the time scale of few seconds. Changes in traffic patterns over the time scale of tens of seconds can be handled by CaDMAC. Our proposed protocols are therefore particularly amenable to wireless backhaul and mesh type of networks.

Summary

We identified the problem of capture in idle nodes, and proposed a simple MAC protocol (CaDMAC) that can filter out capture-causing signals. Alleviating capture enables nodes to participate in their respective productive communication, increasing the reuse in the network. The impact of deafness is also reduced. However, CaDMAC cannot eliminate capture completely – several cases are possible in which a beam receives both productive and capture traffic. In such scenarios the beam cannot be

turned off, and thereby, the corresponding node will continues to get captured by communications in the vicinity. To further reduce the impact of capture in wireless multihop networks, a routing solution is necessary that can choose routes least affected by capture. However, even without assistance from a routing solution, our simulations show that CaDMAC can improve the network performance for stationary networks, like wireless backhaul or mesh networks.

Chapter 5

ROUTING PROTOCOLS USING BEAMFORMING ANTENNAS

In Chapter 3 we argued the need for higher layer protocols to exploit the characteristics of beamforming antennas. In Chapter 4, we evaluated the impact of beamforming on the MAC layer, and proposed new MAC protocols to address them. While antenna-aware MAC protocols offer performance improvements, they might not achieve the complete gains available from beamforming. As mentioned earlier, the higher transmission range can be leveraged if the network layer utilizes this capability to find shorter hop routes. Nodes may be free of capture if the routing protocol assists the MAC protocol by choosing capture-aware routes. The overall improvements from an antenna-aware network layer can be worthwhile, especially when used in conjunction with an appropriate MAC protocol. Motivated from this, we discuss the design and evaluation of a new routing protocol that is capable of exploiting beamforming antennas in wireless multihop networks.

The rest of this chapter is organized as follows. Section 5.1 revisits the directional DSR protocol, and addresses the issues on directional broadcast identified in Chapter 3. Several optimizations, including *partial sweep* and *delayed route reply* are discussed in subsections, that together improve broadcasting functionalities with beamforming antennas. In Section 5.2, a new routing protocol is designed with the capability to perform capture-aware routing. The optimizations of directional broadcast is combined with capture-awareness to form a consolidated protocol named *Capture-aware Routing Protocol (CaRP)*. Optimizations are proposed, and the protocol is evaluated through simulations.

5.1 Directional DSR (DDSR)

Recall directional DSR (DDSR) – a simple extension of DSR with the capability to broadcast with beamforming antennas. In addition, DDSR includes a proactive neighbor discovery scheme, where nodes transmit *hello* packets to establish connectivity with its *DO-neighbors*. Earlier discussions on DDSR showed that routing with beamforming antennas resulted in one key tradeoff. On one hand, broadcasting with beamforming antennas (implemented through *sweeping*) leads to multiple transmissions of the same control packet, leading to high control overhead. Also, *sweeping* is affected by deafness and capture (arising from DiMAC) that further impairs DDSR from achieving the full benefits of beamforming. On the other hand, the discovery of fewer-hop routes via *sweeping* improves the throughput and delay performance of the network. The tradeoff needs to be addressed carefully to extract benefits from directional routing. Motivated from this, we propose optimizations to Directional DSR (DDSR).

5.1.1 Improving DDSR

Delayed Route Reply Optimization

Simulation results revealed that DDSR often chose suboptimal routes in terms of hop-count. This phenomenon occurred even in the absence of other traffic, implying that collisions and capture were not responsible in this cases. Close examination of DDSR explained the occurrences. Observe that with DDSR, neighbors of a node, X, receive the RREQ from X at different points in time. Consequently, neighbors that receive the RREQs earlier, have a higher probability of delivering the RREQs earlier to the destination. However, earlier delivery in this case may not indicate a shorter-hop route (several omnidirectional routing protocols assume that the earliest arriving RREQ traversed the shortest-hop route). If the routing protocol requires the destination to reply only to the earliest arriving RREQ, a suboptimal route can get selected in DDSR. Prolonged use of such a suboptimal route may result in performance degradation. Therefore, replying to the

earliest arriving RREQ may not be suitable when *sweeping* is performed. Interestingly, replying to all RREQs (as optionally performed by DSR) is not suitable either. We discuss the reasons for this next, and present an optimization called *"Delayed Route Reply"* to handle the problem.

Consider the point of time when a destination node receives the first RREQ. Assuming DSR's option to reply to all RREQs, the destination must immediately prepare to reply back with a route reply (RREP) for this received RREQ. With DDSR, the destination must beamform in the direction of the node from which it received the RREQ, and transmit the RREP back to it. The RREP is a unicast packet, and needs to be preceded by an RTS/CTS handshake at the MAC layer, and followed by an acknowledgment. Observe that over the entire duration of this dialog, the destination is deaf to all its other neighbors. If the other neighbors broadcast an RREQ, the destination will fail to hear it. Clearly, if one of the unheard RREQs had traversed the optimal path, the destination would not be able to choose it. Therefore, replying to all RREQs does not ensure choice of the optimal path[1].

Our *"Delayed Route Reply"* optimization requires the destination node to delay sending the route reply (RREP) by a time duration T, calculated from the time it received the first RREQ. This allows the destination node to choose the best among all the routes that arrive within this time T. We specify T as $T = \rho \times T_{sweep}$, where T_{sweep} is the time taken to complete one full sweep, and ρ is a configuration parameter. Waiting for a sufficiently long time duration T minimizes this possibility. We have observed that with this simple optimization, the quality of selected routes improve in several scenarios.

Partial Sweeping Optimization

The overhead associated with sweeping degrades the performance of DDSR, especially when the network is dense or route failures are frequent. Assuming N beams, sweeping requires N sequential transmission of the same packet in N different directions. We intend to compare the control

[1]Note that even if ToneDMAC is used at the MAC layer, this problem persists. This is because ToneDMAC mitigates deafness in unicast communication. Eliminating broadcast for broadcast communication, such as in route discovery, is a hard problem. Therefore, robust route selection mechanism, in the presence of underlying deafness, needs to be incorporated at the network layer.

overhead of DSR and DDSR. A traditional metric based on the "number of transmitted control packets" does not seem appropriate. A more suitable metric is one that characterizes the amount of network resources consumed due to this control traffic. An example might explain this point better. Two protocols, X and Y, may transmit equal number of control packets to achieve identical objectives, although X may need to transmit packets at half the power level of Y. Clearly, X is superior to Y in view of its ability to achieve higher spatial reuse in the network. A metric that is restricted to the number of control packets alone, will not be able to capture the distinction between X and Y. Based on this observation, we design the following overhead metric α.

$$\alpha = \frac{\sum(NumberofControlPackets \times AreaBlockedbyEachPacket)}{\sum(NumberofDataPackets)}$$

where *AreaBlockedbyEachPacket* is the approximate conical area covered by the radiation pattern of a single beam. Intuitively, network capacity consumed due to the transmission of each control packet is proportional to the region silenced by that transmission. The total control overhead is thus equivalent to *Number of Control Packets* $\times$ *Area Blocked by each packet.* This overhead is in exchange of the total *Number of Data Packets* received by sink nodes in the network. We use the ratio of these two quantities as the metric α. Clearly, the overhead of route discovery in DDSR can be significantly higher than DSR. To reduce this overhead in DDSR, we propose the following optimization.

Partial Sweeping:

In Figure 5.1(a), when node Y receives a RREQ from node X, it may not be necessary to transmit that packet back in the direction of node X. This is because, nodes lying in the coverage of node X's beam 3, can receive the packet from X and need not receive a copy of the same packet from Y again. In omnidirectional scenarios, the problem of redundant reception of RREQs has been referred to as the *broadcast storm* problem [36], since control packets are unnecessarily broadcast in the channel. Using directional antennas, we consider an optimization whereby a node forwards a control packet with only n $(n \leq N)$ beams. A node initiating the control packet, however, forwards it using all beams.

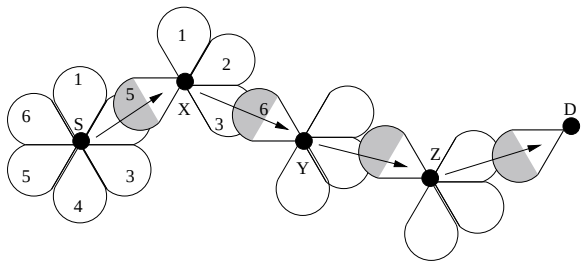

(a) An example scenario showing the propagation of RREQ using the *partial sweeping* optimization. The beams with which nodes receive RREQs are shaded.

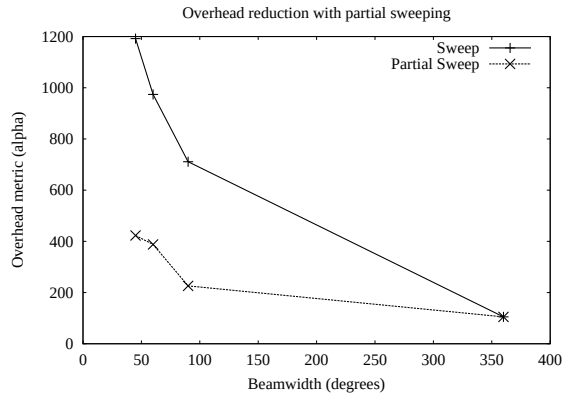

(b) Routing overhead α for optimized and unoptimized DDSR.

Figure 5.1: Partial sweeping optimization with directional DSR

We define $\gamma = \lceil \frac{n}{N} \rceil$. Parameter γ may be treated as a configuration parameter. We have used $\gamma = \frac{1}{2}$ for our simulations. Also, in our simulations, the n beams used to forward control packets are the ones that are diagonally opposite to the beam with which the control packet was received. More specifically, if $\gamma = \frac{1}{2}$, N = 6 (i.e., 6 beams numbered from 1 to 6), and if a node receives a RREQ using beam 5, then antenna beams 1, 2 and 3 are used to forward the RREQ. The choice of the value of γ and the use of diagonally opposite beams are motivated by fact that the RREQs can be forwarded outward from the source node. Figure 5.1(a) shows how a RREQ may propagate using $\gamma = \frac{1}{2}$ and N = 6. Observe that node X receives the RREQ using beam 5 and forwards it using beams 1, 2 and 3.

Figure 5.1(b) plots routing overhead α against beamwidth of directional antennas. α was calculated from 25 simulations of random, static topologies. Evident from the figure, DDSR's routing overhead reduces due to the partial sweeping optimization, although it remains greater than DSR (which corresponds to the beamwidth of 360°). The connectivity of the graph does not get affected; in fact in some cases, we observe that better routes are discovered because of the fewer RREQ collisions in the network.

Addressing Node Mobility

In the presence of mobility, a node's neighbor may move out of the radiation pattern of one beam to that of another. Look-up tables, that maintain the mapping between neighbors and beams, will get stale as a result of such mobility. To track a mobile neighbor, a node needs to switch between different beams over time, and accordingly refresh the look-up table. The act of switching to the appropriate beam is called *handoff*. The efficiency of performing handoff can influence the performance of the network, since poor handoff can lead to frequent link breaks, which will unnecessarily result in a new route discovery phase. An efficient handoff mechanism is clearly associated with overheads, since nodes must detect the need for a handoff, and then establish renewed communication with its neighbor after a successful handoff. Based on these non-trivial overheads, one might observe that beamforming antennas are unsuitable for mobile networks.

We argue that there are incentives to use beamforming antennas even in mobile scenarios. As a tradeoff to the overheads of handoff, beamforming antennas offer the possibility of greater link lifetime. The higher transmission range of a directional beam allows a node to communicate to its neighbor for a longer duration, when the neighbor is moving away from the node. While omnidirectional antennas (with a shorter communication range) would have to initiate new route discovery, beamforming antennas can maintain an active link for a longer timespan. This reduces the frequency of route errors, compensating for the losses from handoff.

We have handled the problem of handoff at the MAC layer and DDSR remains unaware of it. DiMAC handles handoffs as follows. If DiMAC does not receive a reply from a neighboring node after a small number of attempts, it assumes handoff has occurred, and attempts to find the appropriate beam that could re-establish communication. To find such a beam, the node performs a partial *scan*, whereby the transmitter node sends "hello" packets using K beams that are adjacent to the beam that was in use previously. As an example, if K = 2 and beam i was used previously, then "hello" packets are transmitted over beams $i-1$ and $i+1$, sequentially. Nodes that have moved to the coverage area of an adjacent beam ($i-1$ or $i+1$ in this example), can now re-establish communication, by replying to the "hello" packet. The look-up table at the transmitter is correspondingly updated with this new beam. If the link cannot be established using any of the

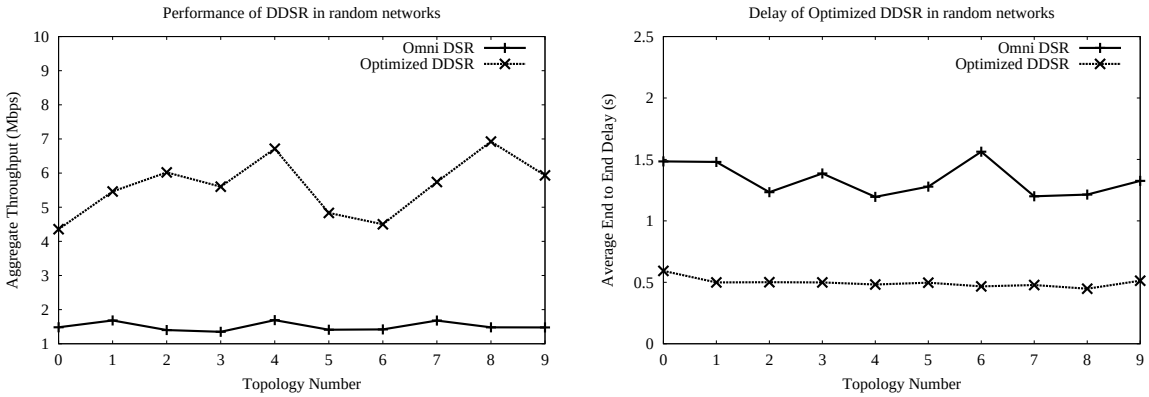

(a) DDSR throughput with incorporated optimizations

(b) DDSR end to end delay with incorporated optimizations

Figure 5.2: Performance of DDSR on static random networks.

K beams, DiMAC informs DDSR of a link failure and a route error packet (RERR) is generated. Thus, handoffs do not always trigger route errors. We have assigned K = 2 in our simulations – a higher value of K may not be very useful since it is unlikely that nodes would move across the coverage area of multiple beams within the time frame of a few packet transmissions.

Performance Evaluation

We evaluate the optimizations of DDSR using the Qualnet 3.6 simulator. We compare optimized DDSR against omnidirectional DSR, and thereby characterize the improvements of higher transmission range of beamforming antennas, when exploited from the network layer. Figure 5.2 shows the results of our simulations. Optimized DDSR performs considerably better than omnidirectional DSR, both in terms of throughput and end to end latency. The optimizations improve DDSR over the unoptimized version of DDSR that we discussed in Chapter 3.

Figure 5.2(b) shows the measurement of average end to end latency for data packets in random networks. The average packet delivery latency is smaller with optimized DDSR in comparison to omnidirectional DSR. The improvements appear as a result of fewer hop routes (although part of

the benefits are offset by the sweeping delay). The end to end delay is also lower than the basic DDSR (proposed in Chapter 5). This happens because the proposed optimization reduces the sweeping delay, and also selects better routes in comparison to basic DDSR.

We also compare the impact of mobility on the performance of the proposed optimizations, as shown in Figure 5.3. Observe that the throughput of both (basic) DDSR and optimized DDSR, degrade with mobility. This happens because the time for route discovery is almost a constant. To illustrate this consider an example where a simulation is performed for 100 seconds using $mobility = 10m/s$. Also assume that over the entire simulation, exactly one route error occurs at time = 20 seconds. Assume that a fresh route is discovered and data transmission resumes at time = 25 seconds. Now consider the scenario using $mobility = 20m/s$ and corresponding simulation time = 50 seconds. Observe that the route error in this scenario occurs at 10 seconds through the simulation and ends at 15 seconds. Consequently, the throughput at higher mobility is expected to decrease since a smaller fraction of the time (45/50) is now utilized for useful data communication (observe that at $mobility = 10m/s$, the fraction of time used for transmitting useful data is 95/100). While performance degrades for both basic and optimized version of DDSR, observe that the degradation is more graceful for optimized DDSR. This is because communication failures due to handoff are handled at the MAC layer by our proposed *scanning* mechanism. This leads to fewer route errors as the mobility in the network increases, which in turn results in higher performance. Also, when route errors occur, and route discovery is initiated, the partial sweeping optimization alleviates some of the negative impacts of directional broadcast. Consequently, optimized DDSR achieves higher performance.

Summary

We identified several issues that arise when beamforming antennas are utilized in multihop routing. Tradeoffs arise when the higher communication range of beamforming antennas are exploited to obtain fewer-hop routes. To address this tradeoff, we proposed several optimizations to DDSR. Our resulting approaches, although simple, alleviates the overheads associated to directional broadcast

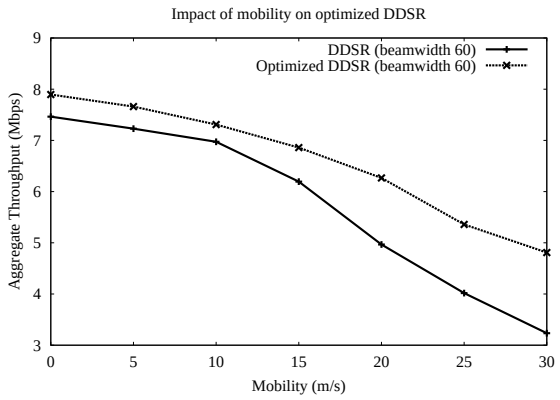

Figure 5.3: Optimized DDSR degrades more gracefully in comparison to DDSR.

and neighbor discovery under mobility. Moreover, we also propose the *delayed route reply* optimization to prevent the selection of sub-optimal routes. Simulation results show that the proposed optimizations can offer additional benefits. These benefits emerge from addressing the tradeoffs associated to the higher communication range of beamforming antennas. However, we believe that there is room for further improvement at the network layer. For example, the ability to receive signals from selective directions (called *capture-awareness* in Chapter 3 and 4) has not been exploited at the network layer. Incorporating capture-awareness in the routing protocol may further increase the achievable performance. Motivated from this, we present capture-aware routing in the next section.

5.2 Capture-Awareness in Routing

Our investigation of MAC protocols using beamforming antennas exposed the problem of *capture* in wireless multihop networks. Recall from Chapter 3 and 4 that capture arises because a node remains engaged in receiving a packet that is not meant for itself. During the entire duration of reception, captured nodes are unable to accomplish their own communication – a wastage of channel capacity. Clearly, capture is unproductive because nodes in the vicinity of an ongoing communication will not be able to accomplish parallel communication. As a result, the full efficacy of beamforming may not be visible unless capture is addressed efficiently.

In Chapter 4, we evaluated the impact of capture and designed a capture-aware MAC protocol

(CaDMAC) to alleviate its effects. Briefly, CaDMAC prescribes the radio layer to turn off a subset of beams that are likely to get captured by unproductive traffic. While CaDMAC reduces capture, we noticed that several occasions remain in which capture persists. We recognized that unless approached from the network layer, the existence of capture cannot be further minimized.

This section presents a "capture-aware" routing protocol, using which selected routes might be less prone to capture. Put differently, routes chosen with an awareness of capture can complement the benefits of CaDMAC. We illustrate this with a simple example in Figure 5.4.

In Figure 5.4, assume that node B routes packets to node A. While such a flow exists, assume that node S intends to discover a route to node G. Routing protocols that use minimum hop-count metrics for selecting a route, may very well choose the route {S-M-F-G}. When using route {S-M-F-G}, observe that CaDMAC at node M must ensure that its east-pointing beam is turned on, in order to communicate with F. However, this very (east-pointing) beam at M is susceptible to capture by transmissions from node B to A. As a consequence, routes {B-A} and {S-M-F-G} will be forced to share the available channel capacity. However, choosing route {S-M-E-G} proves to be beneficial because transmissions on link M-E and B-A are free of mutual capture. The throughput achieved by each of the individual routes increases. *Intuitively, in the presence of more flows, more routes can traverse through this region without having to share the channel between them, as long as there is an available direction that is capture-free from others.* Our capture-aware routing metric is designed to pick out all such capture-free directions (whenever possible), and route traffic along these directions.

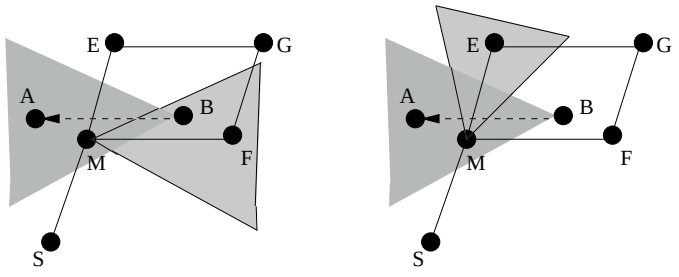

Figure 5.4: Communication on link A-B captures link M-F. Routes that avoid using link M-F (and chooses link M-E instead) extracts greater spatial reuse from the channel

5.2.1 Capture-Aware Routing Protocol (CaRP)

Measuring Route Cost

Our goal is to associate each route with a cost that indicates the "amount of capture" on that route. To achieve this, we determine the capture-cost of each link on the route. Since each link can be represented by a pair of directional beams, quantifying the *amount of capture* on each beam can determine the capture-cost of the end-to-end route.

To be precise, consider an end to end route, R, between nodes S and D. Let nodes i and j be two adjacent nodes on this route. Let B_{ij} be the beam that i uses to communicate to j, and let B_{ji} be the beam that j uses to communicate to i. Let $C_{B_{ij}}$ represent the "amount of capture" on beam B_{ij}. Then, the capture cost of link ij, denoted as κ_{ij}, can be calculated as

$$\kappa_{ij} = C_{B_{ij}} + C_{B_{ji}} \tag{5.1}$$

Consequently, the capture-cost of route R, denoted by κ_R, can be calculated as

$$\kappa_R = \sum_{ij \in R} \kappa_{ij} \tag{5.2}$$

We now describe a simple mechanism to calculate capture-cost of a beam. We use the specific example in Figure 5.4 for our description, and then generalize the mechanism. In Figure 5.4, when there is only a single flow between nodes B and A, beam B_{MB} is affected by capture. Now observe that beam B_{MB} is the very same beam as beam B_{MF}, i.e., the east-pointing beam. Thus, while considering a new route that includes beam B_{MF} in it, the capture-cost of B_{MB} must be taken into account. In other words, the value of $C_{B_{MF}}$, for this example scenario, is equal to $C_{B_{MB}}$. In general,

$$C_{B_{ij}} = \{\sum_k C_{B_{ik}} : \ B_{ik} \ is \ same \ as \ B_{ij}\} \tag{5.3}$$

Using this mechanism, observe that the capture-cost of beam B_{ME} is zero because it is not captured by already existing traffic in the network. Thus, when comparing the cost of alternate

routes from S to G, the cost of route {S-M-F-G} is higher than {S-M-E-G}, making {S-M-E-G} the preferred route with our capture-aware routing metric.

- Incorporating Hop-Count

  The design of our capture-aware routing metric, to this point, is not loop-free. Since the capture-cost of beams can be zero (i.e., when they are not affected by capture), the magnitude of κ_R does not increase monotonically with increasing number of hops on the route, R. Moreover, when using κ_R, a routing protocol can choose a significantly long route if the longer route is less affected by capture. Since hop-count of a route also impacts the achievable throughput, we incorporate hop count into our metric. Each link ij on a route is associated with a cost H_{ij}, Existing metrics such as ETT, that account for link quality, can be used to update the value of H_{ij}.

- Sharing Active Nodes

  Routes that are equivalent in terms of capture and hop-count, may not offer equal throughput. This is because intermediate nodes in one route might participate in other communications, and thereby offer less bandwidth to a newly admitted flow. To favor routes on which nodes are less loaded, we require the routing module at node i to maintain a *participation variable*, P_i, that reflects the node's participation in ongoing communication. Nodes that act as source, intermediate nodes, or sink are assigned a positive value for P_i, and are called "active nodes". Nodes that do not participate in any communication are assigned $P_i = 0$, and are referred to as "inactive nodes". Participation variables at active nodes are associated with timeouts. If an active node does not participate in any communication for a threshold period of time, the node resets P_i to 0.

A Unified Routing Metric

The unified cost of a route, U_{route}, can now be formulated as a weighted sum of all the factors that contribute to the capacity of a route. Thus U_{route} is expressed as

$$U_{route} = \sum_{ij \in R} \omega_\kappa \kappa_{ij} + \omega_p P_i + H_{ij} \tag{5.4}$$

where ω_κ and ω_p are weights used to control the influence of participation and capture, on the choice of routes. In our simulation of this routing metric, we assign $H_{ij} = 1$. For the capture-cost, we assign $C_{B_{ij}} = n$, where n is the number of nodes that capture node i on its beam $C_{B_{ij}}$. For active nodes, we assign $P_i = 1$ when a node is a source or destination of a flow. However, if an active node is an intermediate node of at least one flow, we assign $P_i = 2$, since it bears the onus of both receiving and forwarding.

We are aware that $C_{B_{ij}}$ and P_i are very course-grain representations of the amount of directional-capture and load on a node. However, they have the property of being insensitive to changing traffic patterns in the network. When traffic patterns fluctuate rapidly in time, such binary forms of indicators can be useful. We intend to refine our cost models as a part of our future work, and incorporate them in our proposed framework. In this paper, we only vary the relative influence of *capture* and *participation* on routing, in order to determine suitable values of ω_κ and ω_p.

Protocol Design

We present the design of a source-initiated capture-aware routing protocol (CaRP), similar to DSR [31]. The source initiates a route request (RREQ) flood in the network – each RREQ accumulates the cumulative cost over all the links it traverses. Observe that RREQ packets are always transmitted using omnidirectional beams. The mechanism of cost accumulation in RREQs is performed as follows. Consider the case in which a node X receives a RREQ initiated by a source node, S. The RREQ forwarded by node X contains the following information.

(1) The cumulative *capture-cost* associated with the partial route from S to X, added to the *capture-cost* for each of X's beams. This implies that a tuple of N separate costs are included in the RREQ, where N is the number of beams. The ith element of the tuple reflects the sum of the capture cost of all beams on the partial route {S-X}, added to the capture-cost of X's i^{th} beam.

Figure 5.5 illustrates this with a diagram.

(2) The cumulative *number of active nodes* on the route along which this RREQ has arrived, denoted as A_{sx}.

(3) The intermediate nodes' identifiers (i.e., the partial source route) along which the RREQ has traversed to reach X from S. This indicates the hop-count.

When node Y receives the RREQ, it deduces the capture-cost on the partial route from from S to Y, using the information available in the RREQ. For example, if node X uses beam 2 to communicate to Y, then Y selects the 2nd value from the tuple included in the RREQ, and adds to it the capture-cost of the beam that Y used for receiving the RREQ from X[2]. The result represents the capture-cost for the partial route from S to Y.
If node Y happens to be already participating on another active route, then it increments the active-node-sharing cost, i.e., $A_{sy} = A_{sx} + P_Y$. Then, node Y calculates the hop count using the list of intermediate nodes included in the RREQ. Once the three factors are calculated individually, node Y calculates the unified cost of the partial route from S to Y, U_{SY}, using equation (4). Node Y remembers the cost of this partial route, for reasons explained later. Node Y then updates the RREQ with the required information, and forwards it. Figure 5.5 illustrates the cost accumulation procedure in a propagating RREQ.

When the destination node receives the RREQ, it calculates the cost of the route using equation (4). If the RREQ happens to be the first one for this route discovery phase, the destination node always responds with a route reply (RREP). For RREQs that arrive later, the destination responds with an RREP only if the cost of the received RREQ is lower than all of the previously seen RREQs.

In the basic DSR protocol, an intermediate node does not forward a duplicate RREQ under the assumption that the earlier arriving RREQ traverses the better path. While this may be reason-

[2] We assume a neighbor discovery module that enables node Y to be aware of the beam used by X, to communicate to Y. In stationary networks like mesh networks, this can be achieved with low overheads.

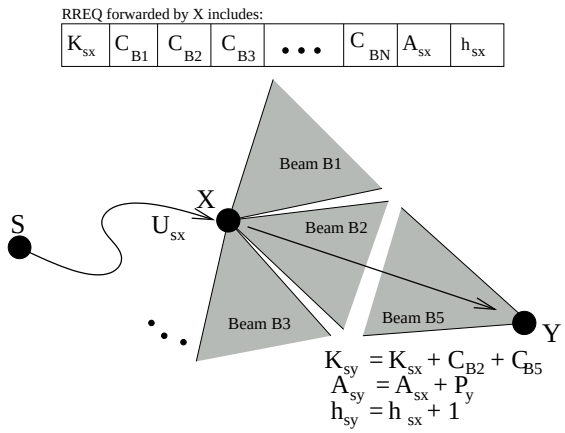

Figure 5.5: Accumulation of route cost in the propagating RREQ

able for discovering minimum-hop routes, it may not be applicable for our purposes. With our proposed capture-aware routing, a longer path may be associated with a lower cost. To account for this possibility, we require intermediate nodes to remember the least cost RREQ that it has forwarded for a particular route discovery. If a lower-cost RREQ arrives later, then we require the intermediate node to forward this RREQ as well.

CaRP Details

- Choosing ON and OFF durations

  The length of ON and OFF durations lead to a performance tradeoff. Longer OFF-Durations enable routes to be less captured by cross traffic, leading to higher performance. However, a sub-optimal route discovered during the OFF-Duration will have to persist proportionally longer, affecting performance. While the optimal ratio between ON and OFF durations is studied later, we observe that for a given optimal ratio, choosing the smallest possible ON duration is desirable, especially in terms of route set-up latency. The smallest possible ON-Duration is a function of the time taken for a source node to discover a route. Knowing the diameter of a network, the length of the ON-Duration can be calculated (with necessary factors of safety). Based on a fixed ON-Duration, we evaluate the impact of the length of OFF-Durations on network performance. The results are presented in Section VI.

- Route synchronization

Recall that if a source node initiates a route discovery in an OFF-Duration, then it reinitiates the discovery process in the subsequent ON-Duration as well. Consequently, several nodes may initiate (pending) route discoveries right at the beginning of an ON-Duration. These route discoveries can get synchronized, and thereby affect the quality of selected routes. We reduce synchronization using jitters, i.e., the pending route discovery is initiated at a time, t, chosen randomly from the ON-Duration. Of course, t is chosen conservatively, so that the initiated route discovery can finish before the node switches to the OFF-Duration.

Also, t is chosen to be at least greater than a minimum threshold for reasons of clock synchronization. In other words, since nodes may not be accurately clock-synchronized, we offer a guard band at the beginning of the ON-Duration. Since the accuracy of clock synchronization can be fairly high (less than 10 μs [**?**]), the guard band remains reasonably small compared to the length of the ON-Duration. The guard band aids CaRP in coping with clock inaccuracies.

Performance Evaluation

We use the Qualnet simulator [38], version 3.6, for simulating our proposed protocols. We present performance comparisons between (CaRP + CaDMAC), (DSR + CaDMAC), and (DSR + 802.11). We show that the CaRP is used in conjunction with CaDMAC, performance can be superior. In other words, the gains from alleviating capture offers significant improvements in throughput and latency.

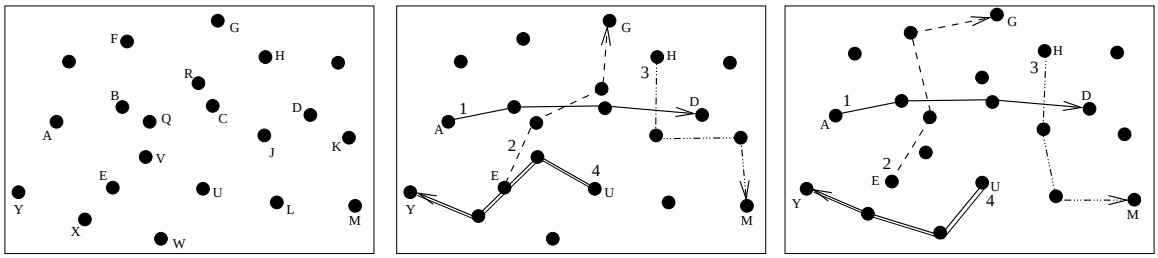

(a) A small network (b) Routes chosen by DSR+CaDMAC (c) Routes chosen by CaRP+CaDMAC

Figure 5.6: Enumeration of routes when using DSR and our capture-aware routing protocol, on a small topology with 4 ordered flows.

Simple Scenarios

We begin with a small network, laying out the detailed impact of choosing capture-aware routes.

Consider Figure 5.6, where a few nodes are placed in a rectangle. In this small network, we introduce four flows in the following order: (1) From A to D, (2) from E to G, (3) from H to M, and (4) from U to Y. More precisely, flow (i) is initiated in the ith ON-Duration, allowing CaRP to choose its optimal route. Figure 5.6 (b) and (c) show the routes chosen by DSR and CaRP respectively (both use CaDMAC as the MAC protocol). The aggregate throughput, as a result of using these routes, is shown in Figure 5.7.

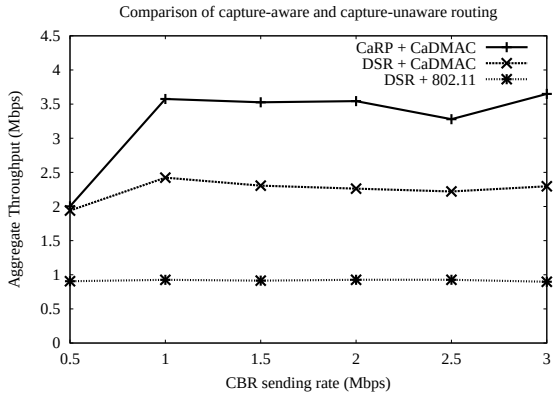

Figure 5.7: Capture-aware routes, chosen by CaRP, offers higher aggregate throughput.

The properties of the chosen routes (in Figure 5.6) explain the differences in achieved throughput. Observe that for flow 1, the routes chosen by DSR and CaRP are identical. This is natural, because, in the absence of any other spatially overlapping traffic, the metrics of DSR and CaRP are influenced only by the hop count of the discovered route. However, when flow 2 is introduced, unlike DSR, CaRP avoids using link Q-R because it is prone to capture from link B-C. Instead it uses link Q-F, which is capture-free from the traffic on route {A-B-C-D}. Similar trends are visible in other cases as well – for flow 3, DSR chooses link J-K that is susceptible to capture from C-D, whereas CaRP chooses J-L. Finally, when flow 4 is introduced, CaRP utilizes the uncongested region of the network by choosing the route through node W. DSR chooses a route through link V-X, which is prone to capture from traffic over link E-Q. Clearly, routes chosen by CaRP can extract higher spatial reuse from the channel, when used in conjunction with CaDMAC. *Observe that there is still more capacity available (in this scenario), since all the available directions have not been saturated.* As a result, in a more congested network, CaRP will still be able to "squeeze

in" more flows that do not share the bandwidth between themselves.

To this point, all data flows were forcibly initiated during ON-Durations. This allowed us to measure the optimistic benefits from capture-aware routing. However, in reality, route discoveries may very well be initiated during the OFF durations. The next experiment takes into account such possibilities, and initiates traffic at random points of time. We compare the performance of DSR with omnidirectional antennas, DSR with CaDMAC, and CaRP with CaDMAC.

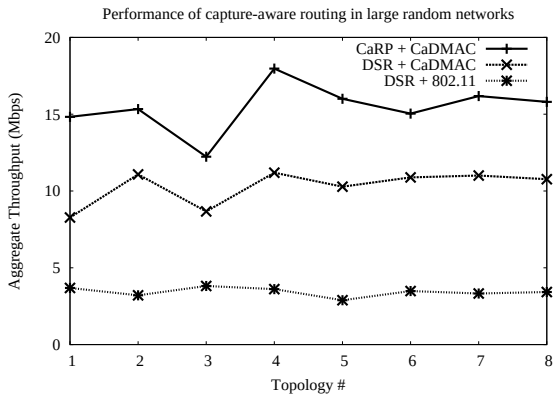

Figure 5.8: Capture-aware routing in large random networks.

We randomly placed 50 nodes within a square of size 2000m, and initiated 20 CBR flows between randomly chosen source and destination nodes. The starting times were chosen randomly as well (without attention to the ON and OFF durations). We simulated our experiments for various topologies. Observe that when using omnidirectional antennas, the route discovery process is always performed on the unaltered topology – DSR chooses the minimum-hop routes in this case. However, when DSR and CaRP are executed over CaDMAC, the route discovery can often be initiated during the OFF-duration, and thereby suboptimal routes are used for the corresponding OFF-Duration. Route discovery is re-initiated in the next ON-Duration, and traffic is redirected if a different route proves to be better. The ON and OFF durations for these experiments are assigned to be 1 and 3 seconds respectively. The beamwidth for this experiment was 60°.

Figure 5.8 shows the results for 8 different topologies. Although many route discoveries were

initiated in OFF-Durations, the benefits of CaRP remain reasonably higher than DSR. This is because CaRP compensates for the sub-optimal routes by choosing a high throughput route in the immediate ON-Duration. Moreover, in moderately dense networks, the route chosen first may not be severely suboptimal because of the availability of multiple nodes that may serve as alternates. With narrower beamwidth, we believe that the improvements can be higher.

- *Impact of Parameters*

  We study the sensitivity of network throughput, and delay, to the length of OFF-Durations. Figure 5.9 presents normalized performance (with respect to DSR+802.11) against the value of $\frac{\tau_{off}}{\tau_{on}}$. Observe that the normalized throughput increases with increase in the OFF-Duration, but after a point (around $\frac{\tau_{off}}{\tau_{on}} = 3$), the curve flattens. Since the duration of CBR flows are specified to be short, very long OFF-Durations force the initially discovered (suboptimal) route to take up a large fraction of the route lifetime, offsetting the gains from capture-awareness. The latency to discover the optimal route has been shown in Figure 5.9. Observe that for more persistent traffic, the benefits from longer OFF-Durations can be higher.

  The normalized end to end delay also improves for small OFF durations – recall that even though sub-optimal routes are used during OFF-Durations, the capture-aware routes selected in the next ON-Duration compensates for the losses. However, the compensation is not sufficient when the OFF-Durations are too long.

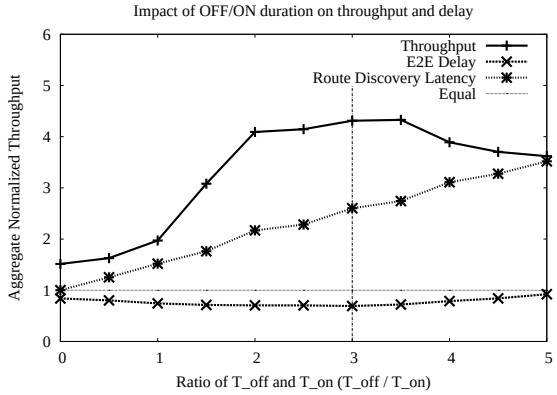

Figure 5.9: Impact of ON and OFF durations on performance.

Observe that Figure 5.9 can be used by a network designer for selecting the length of ON and OFF durations. For this paper, we used $\frac{\tau_{off}}{\tau_{on}} = 3$, since that region appeared to be reasonable both in terms of throughput and delay.

Figure 5.10 shows the sensitivity of normalized network performance to the ratio of $\frac{\omega_p}{\omega_\kappa}$, where $\omega_\kappa = 1$. Observe that when the network is congested, the performance is almost insensitive to ω_p because nodes in two different routes have their *participation variables* almost similarly set. However, when the network has some free nodes (moderately loaded), there is benefit in choosing a longer route if the longer route includes nodes that do not participate in any communication, i.e., increasing ω_p is useful. However, beyond a point, increasing ω_p degrades performance because instead of leveraging the available capacity from a particular region, the routing metric attempts to bypass the congested region, making the hop-count significantly long. In the absence of information about traffic patterns, we choose $\frac{\omega_p}{\omega_\kappa} = 1$ for this paper. As an aside, observe that when $\frac{\omega_p}{\omega_\kappa} = 0$, the benefits in performance are solely from capture.

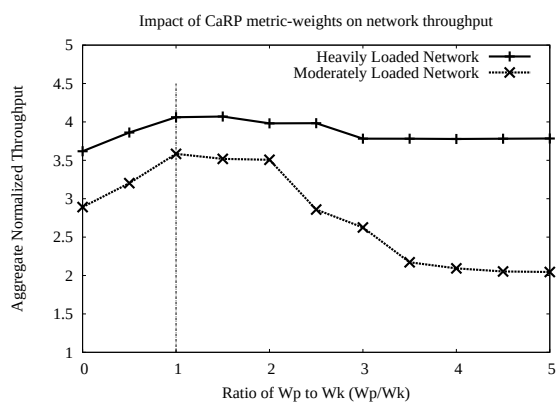

Figure 5.10: Impact of varying weights on the performance of CaRP.

Discussion

With our capture-aware routing protocol, the cost of a route is a function of the traffic pattern in the network. If the traffic pattern changes between time instants t_1 and t_2, a route selected before t_1 may not be the best-possible route after time t_2. As a result of this dependence, spatial reuse of the channel may not be maximal, even with CaRP. One way to address this could be to have

nodes periodically rediscover routes during the ON-Duration.

Caching

In mobile ad hoc networks, DSR utilizes the option of overhearing route reply packets, thereby learning routes to other nodes without having to discover them. Capture-aware protocols reduce the opportunity to overhear packets, reducing the benefits of promiscuous listening. However, in applications such as wireless mesh networks, nodes remain almost stationary – the need for promiscuous listening may not be severe in such cases.

Backhaul and Mesh Networks

Directional antenna protocols are proving to be an attractive design choice for wireless backhaul and mesh networks. Such networks are typically characterized by stationary nodes and persistent traffic. The protocols proposed in this paper can be particularly amenable to these type of networks.

Summary

We identified the problem of capture in idle nodes, and proposed a combination of a MAC and routing protocol, named CaDMAC and CaRP respectively. By bypassing capture over short and long time scales, we demonstrated significant improvements in performance when using beamforming antennas.

Chapter 6

EXPERIMENTAL EVALUATION

6.1 System Description

Implementations of wireless adhoc networks have typically used omnidirectional antennas that come along with the wireless interface cards. With advances in signal processing and antenna technology, it may now be possible to incorporate sophisticated antennas into portable devices, and thereby exploit the benefits of beamforming at the MAC, network, and transport layer. Recent networking research has produced exciting results in this direction. However, most of the research has either been theoretical or based on simulations. In this chapter, we present the implementation of an ad hoc network prototype that uses beamforming antennas – we name it Beamnet. Our implementation quantifies the tradeoffs emerging from the use of beamforming antennas in multihop networks. We report the challenges and discuss lessons learned from our testbed. Through careful measurements, we show how beamforming antennas can be exploited to improve network performance, especially in the context of static applications like mesh networks and wireless backbones. We believe our observations can be a valuable feedback to both antenna designers as well as to wireless protocol designers.

We developed a testbed with 6 IBM T41 laptops, each equipped with *CISCO Aironet 350 series LM Wireless interface Cards.* To the external antenna-ports of these Cisco wireless cards, we connected our beamforming antenna system, called *electronically steerable passive array* (ESPAR) (shown in Figure 6.1). ESPAR [62] is capable of supporting 13 distinct beam patterns, i.e., 12 overlapping beams with 60° half power beamwidths (HPBW) and one omnidirectional beam. To

obtain a desired beam with ESPAR, corresponding voltages must be used to drive the antenna system. We controlled the voltage assignment through software by sending appropriate voltages over the USB port, to a 8 channel Contec D/A converter. The output of the D/A converter was amplified by a simple amplifier and then fed to the elements of ESPAR. The block diagram of the system architecture is shown in Figure 6.2. Using this system prototype, we evaluate the performance of basic MAC and network layer protocols, under various environments.

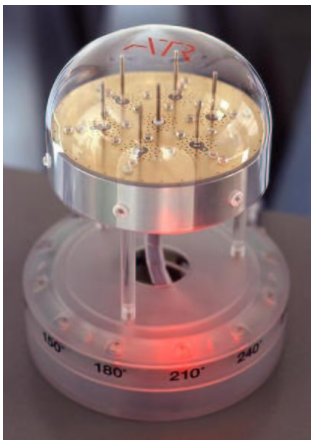

Figure 6.1: ESPAR antenna prototype (for 2.4 GHz band).

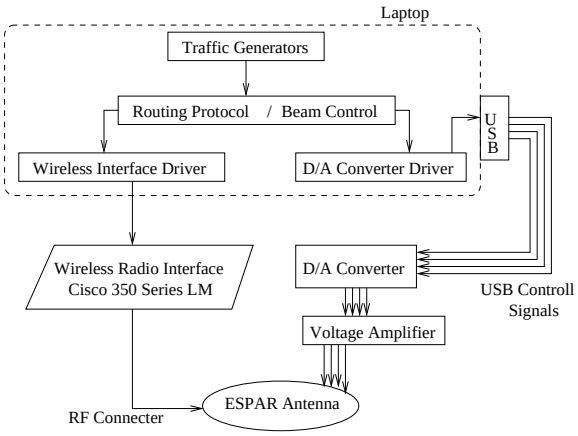

Figure 6.2: Block diagram showing interaction between laptop and ESPAR.

6.2 Related Work

Recent years has seen several testbed implementations that aim to evaluate/validate the performance of ad hoc networks [63], [64], [65], [66], [67], [68]. While most of these testbeds use omni-

directional antennas, only one reported so far have utilized directional antennas. We discuss some of the key observations from these papers, and refer to them later when relevant to our context of directional communications.

In [63], the authors observed that links are inherently transient and, even while present in the routing tables, can cause higher layer traffic-communication to fail. The paper suggests that signal strength based routing can be effective – links chosen based on a quality metric can prove to be reliable for higher layer communications. The authors in [39] also confirm the probabilistic nature of links and evaluated how weak routes can deeply affect packet delivery between source-destination pairs. They propose "transmission count" (a measure of the number of packets received in the near past) as an alternate metric to "fewest-hop" routing. In addition, the authors emphasize the frequent occurrences of link asymmetry, even when nodes forming the link operate at equal power levels.

A testbed comparison of AODV and DSDV was made in [64]. The possibility of intermittent link failures caused the authors to propose a signal-strength based routing metric – *powerwave*. *Powerwave* is a neighbor filter policy that adds only the reliable neighbors to the routing table, thereby improving upon the handoff between routes. Motivated by the non-repeatability of measurements in wireless networks, authors in [66] aimed to produce repeatability in traffic and mobility by choreographing the movements of nodes through scripts. The APE architecture they propose can be viewed as a tool to compare protocols. In [68], the authors have evaluated the associativity based routing (ABR) protocol, and suggested localization techniques to limit routing overheads.

While most testbeds have used omnidirectional antennas, authors in [57] were reportedly the first to utilize directional antennas for ad hoc networking. Including directional neighbor discovery and a power-controlled MAC protocol, the authors were able to realize directional end-to-end communications. Their evaluations indicate the network wide benefits of directional communication, but does not offer much insight into what leads to the overall improvements. In this dissertation, we validate/evaluate the basic primitives of directional communication from a networking perspec-

tive. We build a prototype testbed as a proof-of-concept. We experiment with simple scenarios and measure the achievable gains from individual antenna properties.

6.3 Experiments

6.3.1 Neighbor Discovery

To be able to transmit directionally to its intended receiver, a transmitter must first beamform in the correct direction. In multipath environments, the correct direction may not be in the line of sight (LoS) and may even change with time. To handle this, we implement a simple neighbor discovery protocol that aims to select the appropriate beam, for communicating to each neighbor. A transmitter, T, periodically initiates the neighbor discovery phase (NDP) – the objective is to determine the most suitable beam, with which T must communicate with its neighbors. In the neighbor discovery phase, node T transmits a beacon on each of the beams, sequentially – called a *sweep*. The beam index, β_i, with which the beacon is transmitted, is included in the beacon. Neighbors of T that receive the beacons compare the received signal strengths of each beacon to determine the beam index with which the received signal strength is maximum ($\beta_{max} = max(\beta_i)$). Once this has been determined, each neighbor transmits β_{max} back to T. Node T updates its neighbor table with these beam indexes, and uses these beams to communicate to the corresponding neighbors. Observe that when T initiates the neighbor discovery phase (NDP), only T is able to learn the beam indexes suitable for communicating to its neighbors. Our implementation therefore requires each node to initiate the NDP.

We record the beam selected for each neighbor after one round of beacon transmission. Each recorded value corresponds to a dot on the graph in Figure 6.3. We continue recording the beam for an hour. Evident from the figure, selected beams are often not in the line of sight (LoS) direction, and are instead scattered in several directions. The line of sight direction is shown in black dots, while the other directions are shown in light gray dots. To prove that multipath was indeed the cause for such scatter, we performed the same experiment inside an anechoic chamber. The plot

shown in Figure 6.4 confirms our hypothesis. In the anechoic chamber, the degree of scatter is far less as a result of mostly line of sight communication[1]. Our experiments show that multipath effects underlie wireless communication, and beamforming protocols that rely on location for estimating a neighbor's direction [53], [69] can be erroneous. This implies that our neighbor discovery mechanism must refresh its neighbor directions periodically.

To understand how often the selected beam changes with time, we plot the results of NDP for a 60-second duration – Figure 6.5. We noticed that the scatter is much longer in contrast to Figure 6.3. This indicates that the beam used for communication does not change too frequently with time. In fact some of the dots that are adjacent to the LoS direction, were also received by the LoS beam, although the received strength was lower in comparison to the adjacent beam. Based on this result, we concluded that initiating NDP at a 60 second interval offers reasonable performance. If the overheads are not a major concern, NDP can be initiated at 30 second intervals for even better performance.

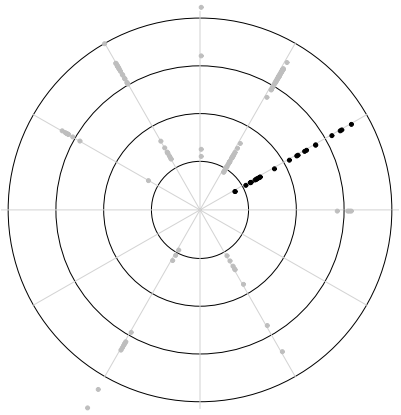

Figure 6.3: Each dot results from one sweep – the angular position of the dot represents the directional beam that was selected in that sweep, and the radial distance is proportional to max received signal strength for that particular sweep. The direction of black dots represents the LoS direction. The graph plots results for 1 hour.

[1]Since the anechoic chamber is not perfect, and because beams are overlapping, sometimes adjacent beams were chosen.

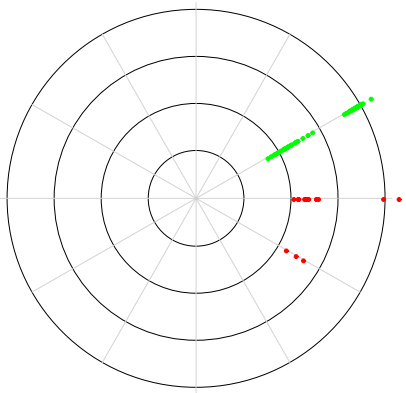

Figure 6.4: Beam selection improves in the anechoic chamber.

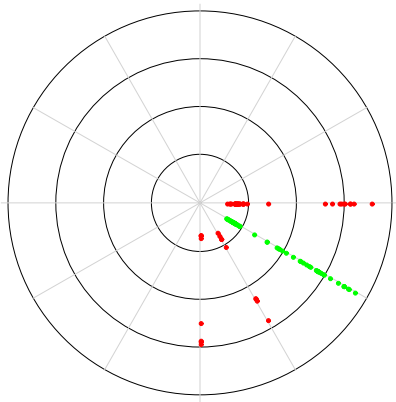

Figure 6.5: Neighbor directions for a plot for 30 seconds.

6.3.2 MAC Layer: IEEE 802.11 and ESPAR

In this section, we discuss our experiences when beamforming antennas are used in conjunction with IEEE 802.11. We were unable to implement a specific directional MAC protocol as most of the radio functionalities in current wireless cards are handled in firmware. However, by controlling the beam through software, we were able to implement a protocol similar to DMAC as proposed in Chapter 3. We explain the experiment configurations next, and discuss the obtained results thereafter.

When a transmitter is ready to send packets to a receiver, we control the ESPAR antenna to beamform in the appropriate direction of the receiver, as prescribed by our neighbor discovery protocol. Observe that the wireless interface cards are off-the-shelf, and incorporates the standard

802.11b MAC protocol. The wireless interface remains unaware of the antenna system that is being controlled from a higher layer. When packets are ready for transmission, the antenna beam is suitably set to accomplish beamformed communication. Using such a set-up, we experiment with different topologies, and validate the feasibility of spatial reuse, link reliability, and range extension. Our findings are presented next.

Spatial Reuse:

To demonstrate spatial reuse, we set up a simple scenario as shown in Figure 6.6. We measure the individual per-link throughput using omnidirectional and directional antennas respectively. Then, we measured the aggregate throughput over the two flows, while the flows are made to communicate simultaneously. The results from this experiment is plotted in Figure 6.7. The graphs represented by OO indicate the throughput when the transmitter and receiver are both omnidirectional, while DD represents directional. While OO and DD achieve comparable individual per-link throughput, in terms of aggregate throughput DD achieves almost double of OO. In other words, due to a smaller interference footprint with beamforming antennas, multiple communications can progress simultaneously resulting in higher channel utilization.

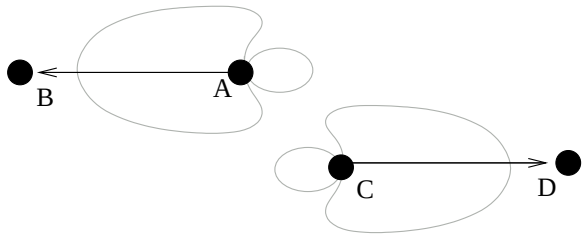

Figure 6.6: Node configurations to exhibit spatial reuse.

Capture:

To validate the prevalence of capture, we set-up a different topology – nodes C and D are placed perpendicular to nodes A and B (similar to Figure 4.16 in Chapter 4). When we initiate both flows simultaneously (i.e., between A-B, and C-D), the aggregate throughput is significantly low indicating that spatial reuse between the flows is not available (Figure 6.8). However, when we force node D to beamform toward C, the achieved aggregate throughput is close to double the previous

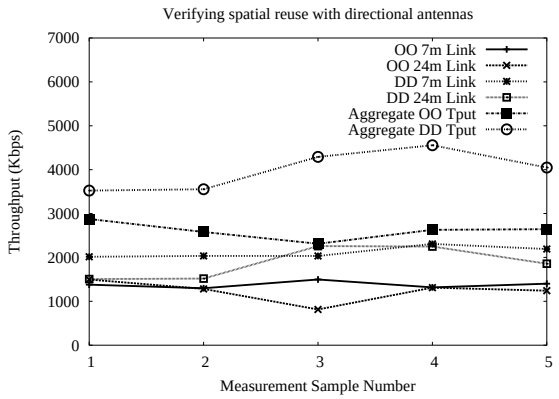

Figure 6.7: Comparing spatial reuse for OO and DD flows.

throughput (Figure 6.9). Recall that this happens because node D gets frequently locked to node A's transmissions to B, and therefore cannot communicate parallely to C. When D beamforms *a priori* toward C, capture can be avoided, leading to parallel communication.

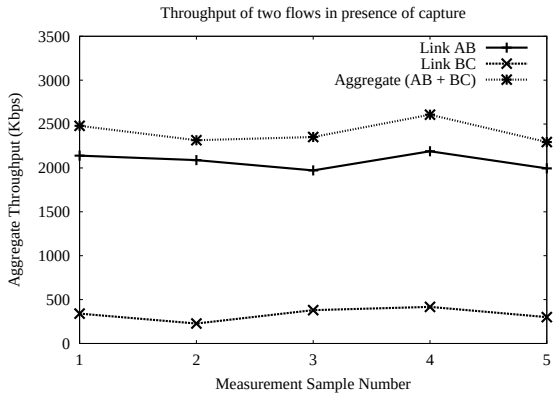

Figure 6.8: Validating capture in a toy topology.

Deafness:

Deafness has been shown to be significant problem with beamforming antennas. Recall that *deafness* is caused when a transmitter fails to communicate to its intended receiver, because the receiver is beamformed towards a direction away from the transmitter. The transmitter is unable to classify the cause of failure, which can either be congestion or deafness. Whenever the cause of failure is incorrectly classified, the MAC protocol incurs performance penalties. We validate the negative impact of deafness through setting up toy topologies in our testbed.

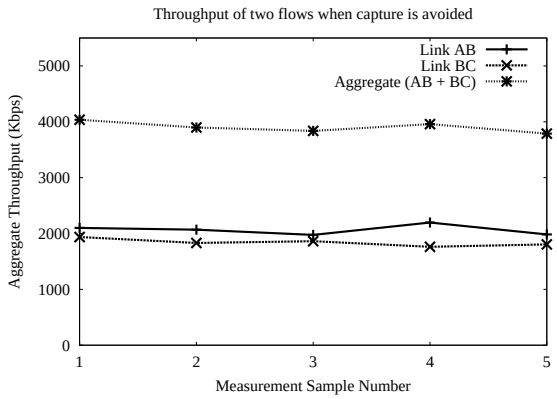

Figure 6.9: Performance when capture is prevented.

In the interest of space, we only present Figure 6.10 to highlight the acute packet drop rate as a consequence of deafness. In the scenario for Figure 6.10, two transmitters (T_1 and T_2) were made to transmit packets to a common receiver R. Node R beamforms towards the node that initiates the dialog first. Once the dialog is over, R returns to the omnidirectional mode and waits for the next packet transmission. While node R is communicating with say T_1, node T_2 will experience deafness if it initiates communication in the meanwhile. When T_2 experiences deafness, it backs off exponentially (as per the 802.11 standard) and attempts retransmission later. However, node T_1 is very likely to initiate communication before T_2 because T_1 backs off for a smaller duration. This can continue several times; finally T_2 may exceed the retransmission limit, and drop the packet. Notice that if T_2 luckily gets channel access, node T_1 will continue to experience deafness. The transmitters frequently drop packets as shown in Figure 6.10. When we turned off the neighbor discovery and beam selection module, forcing nodes to use the default omnidirectional antennas, no packets were dropped. Clearly, deafness is a problem that arises with beamforming, and unless addressed, can offset the net benefits from smarter antennas.

Link Quality:

Experiences from testbeds have shown that links selected by the routing protocol often offer inadequate stability for data communication [39]. This happens because the channel quality during route discovery may happen to be better; at later points in time, the channel quality may fluctuate resulting in link failures. Moreover, the route discovery packets are small in size, and transmitted

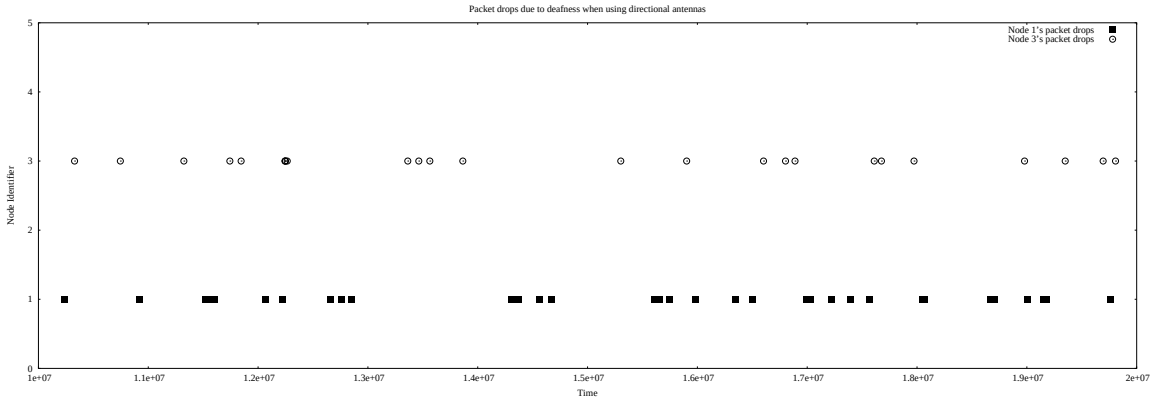

Figure 6.10: Packet drops rates at transmitters due to the impact of deafness

at lower rates. While a link might be good enough to support low data-rate packet transmissions, it may not be so for larger, high data-rate packets. Beamforming antennas are beneficial in such scenarios because links are discovered with directional transmission and omnidirectional reception.However, when data communication actually takes place, the transmitter and the receiver mutually beamform toward each other. This increases the achieved SINR on the link (a factor of safety) making the link more robust to channel fluctuations. Even if weak links are found during route discovery, the links might be able to support future communication because of the added transmission and reception gains of beamforming antennas. Figure 6.11 confirms the performance gains for increasing packet sizes over routes chosen by a link state routing protocol. Clearly, when using omnidirectional antennas, the throughput over that link degrades quickly with increase in packet size, while with beamforming antennas, the degradation is less pronounced.

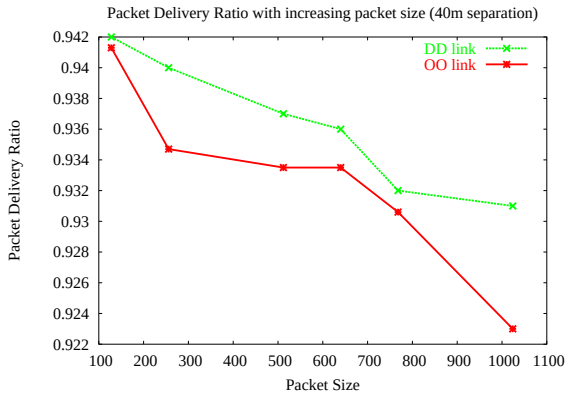

Figure 6.11: Higher packet delivery ratio with increasing packet sizes

6.3.3 Link State Routing

We implemented a simple link state routing protocol. Nodes propagate their neighborhood information to their single-hop neighbors, who in turn propagate it further. On account of such flooding, nodes in the network learn about the entire topology within a warm-up time. In addition, every node builds up its own look-up table that contains the suitable beam indices to communicate to its single-hop neighbors. While such blind flooding of link state packets may not be optimal, we adopt this in view of our small testbed of 6 laptops. We performed multihop packet delivery using this routing protocol. Because of the high communication ranges of the antennas, we performed the experiments on a soccer field. We observed that in several cases, directional antennas formed links between nodes that were connected over two hops when using omnidirectional antennas. Table III shows the end to end throughput of the multihop flow. We expect that over larger networks, the benefits would be much more pronounced.

6.3.4 Discussion and Summary

Our overall experience with a small prototype shows that in reality, channel behavior is heavily fluctuating. Routes that seem to be stable at one instant of time, may not be suitable for continued use in the future. While performance-optimization is important, we strongly felt the need to embed a high level of resilience in ad hoc networking protocols. Existing ideas like proactive route searching, opportunistic/diversity forwarding, anycasting, etc., should probably be revisited. We also observed significant link asymmetry and multipath effects in our measurements of the testbed prototype. Protocols that rely on these assumptions may suffer when executed in real environments. Further, from our experiences with directional MAC protocols, we believe that deeper nulls are as important as narrow beamwidths. We intend communicate this to the antenna design community, so that these considerations could be taken when optimizing for the antenna beamforms.

In summary, we attempted to evaluate the benefits of beamforming antennas in a prototype implementation of ad hoc networks. We showed the potentials of spatial reuse, better link quality, and routing benefits on one hand, and the negative impacts of deafness and multipath on the other.

Although an early evaluation, we believe that our experiences with our testbed will contribute to a better understanding of wireless multihop networking using beamforming antennas.

Chapter 7

SUMMARY AND FUTURE WORK

The tradeoff between *resource* and *performance* appears in almost all areas of engineering. With continuous research and development, resources become less expensive (and more feasible), and thereby the tradeoff continues to be interesting. One such resource is beamforming antennas that have received intense research attention from signal processing, communications, and the antenna design communities. In this dissertation we have visited the problem of utilizing beamforming antennas for wireless multihop networks. We have investigated the impact of beamforming antennas on the performance of wireless multihop networks, and identified several new challenges that need to be addressed. We have proposed MAC and routing protocols to address these challenges. Through simulations and some practical experimentation, we have demonstrated higher levels of performance achievable by beamforming antenna systems.

We began this dissertation by evaluating the effect of replacing an omnidirectional antenna with a beamforming antenna. We designed simple extensions of omnidirectional protocols, namely IEEE 802.11 and DSR (at the MAC and network layer, respectively) to drive the beamforming antennas. We aimed to study the performance of such a system, where protocols are not specifically designed for beamforming antennas. We hypothesized that if the performance of the system met the achievable limits, then research on antenna-aware protocols could be deemed unnecessary. Otherwise, the results of our experiments would serve as the motivation for protocols designed for beamforming antennas.

The results of our experiments with DMAC and DDSR were below expectations. Simulation

traces offered useful insights into the inadequacies of these protocols. At the MAC layer, we identified a wide range of problems that impair the performance of beamforming antennas. *Hidden terminal problems*, *deafness*, *capture* were some of the problems that had never appeared with omnidirectional antennas. Together, these problems significantly offset the benefits of beamforming antennas. At the network layer, broadcast proved to be an important problem – a tradeoff appeared when discovering fewer-hop routes with beamforming antennas resulted in high *sweeping* overheads. The interaction of MAC and routing protocols also needed to be designed carefully. On studying the problem space we observed that the issues were non-trivial, and may not be resolved with variants of existing protocols. New protocols, tailored for beamforming antennas, were necessary in order to achieve further performance benefits.

We designed new MAC protocols that address the most important problems – deafness and capture. To alleviate deafness, we designed ToneDMAC that uses tones as a feedback to notify the neighbors of a node about the node's activity. Unlike earlier protocols that use busy tones, ToneDMAC does not require an additional transceiver. With ToneDMAC, we showed how several scenarios were less affected by deafness. Simulation results showed that the throughput of the network improved for both UDP and TCP traffic. The fairness among flows were also considerably higher when using ToneDMAC.
To address the problem of both deafness and capture, we designed a protocol named CaDMAC. We showed how CaDMAC reduces deafness, and significantly mitigates the impact of MAC-layer capture. This is achieved by treating the beamforming antenna as a spatial filter, and filtering out unproductive signals in space. CaDMAC results in significantly higher spatial reuse in the network, which translates to higher throughput and lower end to end latency. Simulation results proved to be encouraging when improvement in aggregate throughput was more than 300%. However, we observed that CaDMAC did not eliminate capture completely, and thereby, there was scope for performance improvement in the system.

The scope for further improvement led us to approach the system from the network layer as well. We noticed that assistance from the routing protocol could improve overall performance. As

a first step, we utilized the higher communication range of beamforming antennas by incorporating a directional broadcast scheme. With directional broadcast – called *sweeping* – routes were discovered using beamforming antennas, and therefore, links on the route proved to be longer than omnidirectional links. As a result, the hop-count of routes decreased with beamforming antennas, which improved the network performance. However, this protocol, named directional DSR (DDSR), was impaired by deafness and unwanted collisions. We presented optimizations to DDSR to handle these issues, and showed that the benefits from higher transmission ranges can be useful. Even under mobility conditions, DDSR performed reasonably well, because of its ability to track a node longer with its higher communication range.

In addition to utilizing higher communication range for routing, we also incorporated capture-awareness in our routing protocol. Named capture-aware routing protocol (CaRP), the proposed protocol is capable of choosing routes on which links are less captured by other traffic in the network. Put differently, CaRP selects those links that needs to share the channel least with other links in the networks, and therefore can increase the amount of channel reuse in the network. Simulation results proved the efficacy of CaRP and demonstrate performance improvements of more than 400% in random networks.

We implemented part of our proposals in a small prototype testbed in which laptops were equipped with electronically steerable antennas called ESPAR antennas. In our testbed, we implemented a neighbor discovery solution where nodes were required to *sweep* hello packets. The sweeping mechanism enables a neighbor to identify the best beam with which subsequent communication should be initiated. We also implemented a link state routing protocol – nodes collect link states from all other nodes in the network, and populate its own routing table appropriately (based on a shortest hop metric). Although small in size, our prototype testbed validates the possibility of improved performance with beamforming antennas in wireless multihop networks.

References

[1] T. Rappaport, "Wireless communications principles and practice," *Prentice Hall*, 2002.

[2] J. Litva and T. Lo, "Digital beamforming in wireless communications," in *Artech House Publishers*, 1996.

[3] IEEE 802.11, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, 1999.

[4] F. Rashid-Farrokhi, K. Liu, and L. Tassiulas, "Transmit beamforming and power control for cellular wireless systems," vol. 16, pp. 1437–1450, Oct 1998.

[5] S. Swales, M. Beach, D. Edwards, and J. McGeehan, "The performance enhancement of multibeam adaptive base-stationantennas for cellular land mobile radio systems," *IEEE Transactions on Vehicular Technology*, vol. 39, pp. 56–67, Feb 1990.

[6] M. Horneffer and D. Plassmann, "Directed antennas in mobile broadband systems," in *Proceedings of IEEE INFOCOM*, vol. 2, pp. 704–712, April 1996.

[7] T. S. Yum and K. W. Hung, "Design algorithms for multihop packet radio networks with multiple directional antennas," *IEEE Transactions on Communications*, vol. 40, no. 11, pp. 1716–1724, 1992.

[8] A. Chandra, V. Gummalla, and J. O. Limb, "Wireless medium access control protocols," *IEEE Communications. Surveys and Tutorials*, vol. 3, no. 2, pp. 55–64, 2000.

[9] F. Alessandri, M. Mongiardo, and R. Sorrentino, "Computer-aided design of beam forming networks for modern satellite antennas," *IEEE Transactions on Microwave Theory and Applications*, vol. 40, pp. 1117–1127, Jun 1992.

[10] J. Zander, "Slotted ALOHA multihop packet radio networks with directional antennas," *Electronic Letters*, vol. 26, no. 25, pp. 2098–2100, 1990.

[11] J. Ward and R. Crompton, "Improving the performance of slotted ALOHA packet radio networks with adaptive arrays," *IEEE Transactions on Communications*, vol. 40, no. 11, pp. 1716–1724, 1992.

[12] M. Sanchez, T. Giles, and J. Zander, "CSMA/CA with beam forming antennas in multi-hop packet radio," in *Proceedings of Swedish Workshop on Wireless Ad Hoc Networks*, pp. 63–69, March 2001.

[13] C. A. Balanis, "Smart antennas for future reconfigurable wireless communciation networks," tech. rep., Arizona State University, Annual Report, april 2000.

[14] Y. Ko, V. Shankarkumar, and N. H. Vaidya, "Medium access control protocols using directional antennas in ad hoc networks," in *Proceedings of IEEE INFOCOM*, vol. 1, March 2000.

[15] A. Nasipuri, J. Mandava, H. Manchala, and R. Hiromoto, "On demand routing using directional antennas in mobile ad hoc networks," in *Proceedings of IEEE ICCCN*, September 2000.

[16] T. Elbatt, T. Anderson, and B. Ryu, "Performance evaluation of multiple access protocols for ad hoc networks using directional antennas," in *Proceedings of WCNC*, 2003.

[17] S. Bandyopadhyay, K. Hausike, S. Horisawa, and S. Tawara, "An adaptive MAC and directional routing protocol for ad hoc wireless networks using ESPAR antenna," in *Proceedings of ACM MobiHoc*, pp. 243–246, October 2001.

[18] M. Takai, J. Martin, R. Bagrodia, and A. Ren, "Directional virtual carrier sensing for directional antennas in mobile ad hoc networks," in *Proceedings of ACM MobiHoc*, pp. 39–46, June 2002.

[19] R. Ramanathan, "On the performance of ad hoc networks with beamforming antennas," in *Proceedings of ACM MobiHoc*, pp. 95–105, October 2001.

[20] M. Streenstrup, "Neighbor discovery among mobile nodes equipped with smart antennas," in *Porceedings of ADHOC Networks*, May 2003.

[21] S. Yi, Y. Pei, and S. Kalyanaraman, "On the capacity improvement of ad hoc wireless networks using directional antennas," in *Proceedings of Mobihoc*, 2003.

[22] A. Spyropoulos and C. Raghavendra, "Capacity bounds for ad-hoc networks using directional antennas," in *Proceedings of ICC*, May 2003.

[23] P. Karn, "MACA - a new channel access method for packet radio," in *Proceedings of 9th Computer Networking Conference*, 1990.

[24] R. R. Choudhury, X. Yang, and N. H. Vaidya, "Using directional antennas in ad hoc networks," tech. rep., Texas A&M University Final Report submitted to BBN Technologies, July 2001.

[25] A. Saha and D. Johnson, "Routing improvement using directional antennas in mobile ad hoc networks," in *Proceedings of Globecom*, October 2004.

[26] S. Roy, D. Saha, S. Bandyopadhyay, T. Ueda, and S. Tanaka, "A network-aware MAC and routing protocol for effective load balancing in ad hoc wireless networks with directional antenna," in *Proceedings of ACM MobiHoc*, pp. 88–89, June 2003.

[27] A. Nasipuri, S. Ye, J. You, and R. Hiromoto, "A MAC protocol for mobile ad hoc networks using directional antennas," in *Proceedings of IEEE Wireless Communication and Networking Conference*, vol. 3, pp. 1214–1219, September 2000.

[28] R. R. Choudhury, S. Bandyopadhyay, and K. Paul, "MARP: A multi-agent routing protocol for mobile wireless ad hoc architectures," in *Accepted for publication in the Journal of Autonomous Agents and Multi-Agent Systems*, Kluwer Academic Publishers.

[29] J. Broch, D. Maltz, D. Johnson, Y. Hu, and J. Jetcheva, "A performance comparison of multi-hop wireless ad hoc network routing protocols," in *Mobile Computing and Networking*, pp. 85–97, 1998.

[30] E. Royer and C. Toh, "A review of current routing protocols for ad-hoc mobile wireless networks," 1999.

[31] D. Johnson and D. Maltz, "Dynamic source routing in ad hoc wireless networks," in *Mobile Computing, Kluwer Academic Publishers*, 1996.

[32] C. E. Perkins and P. R. Bhagwat, "Highly dynamic destination-sequenced distance vector routing (DSDV) for mobile computers," in *Proceedings of ACM SIGCOMM*, pp. 234–244, 1994.

[33] G. Pei, M. Gerla, and T. Chen, "Fisheye State Routing in Mobile Ad Hoc Networks," in *ICDCS Workshop on Wireless Networks and Mobile Computing*, (Taipei, Taiwan), April 10 2000.

[34] I. Kang, R. Poovendran, and R. Ladner, "Power-efficient broadcast routing in adhoc networks using directional antennas: Technology dependence and convergence issues," in *Proceedings of Globecom*, 2003.

[35] J. Cartigny, D. Simplot, and I. Stojmenovic, "Localized energy efficient broadcast for wireless networks with directional antennas," tech. rep., Univerity of Sciences and Technologies de Lille, July 2002.

[36] C. Hu, Y. Hong, and J. Hou, "On mitigating the broadcast storm problem with directional antennas," in *Proceedings of ICC*, 2003.

[37] C. Shen and Z. Huang, "Topology control with directional power intensity for ad hoc networks," *Proceedings of IEEE Computer Communications and Networks*, 2002.

[38] S. N. Technologies, "Qualnet simulator version 2.6.1," www.scalable-networks.com.

[39] S. Douglas, J. D. Couto, D. Aguayo, B. Chambers, and R. Morris, "Effects of loss rate on ad hoc wireless routing," in *MIT Laboratory of Computer Science Technical Report, MIT-LCS-TR-836*, March 2002.

[40] Y. Wang and J. Garcia-Luna-Aceves, "Broadcast traffic in ad hoc networks with directional antennas," in *Proceedings of Glbocom*, 2003.

[41] J. Deng and Z. Haas, "Dual Busy Tone Multiple Access (DBTMA): A New Medium Access Control for Packet Radio Networks," *IEEE Transactions on Communications*, 1998.

[42] Z. Huang, C. Shen, C. Srisathapornphat, and C. Jaikaeo, "A busy-tone based directional MAC protocol for ad hoc networks," in *Processdings of Milcom*, 2002.

[43] T. Korakis, G. Jakllari, and L. Tassiulas, "A MAC protocol for full exploitation of directional antennas in ad hoc wireless networks," in *Proceedings of Mobihoc*, 2003.

[44] M. Neufeld and D. Grunwald, "Deafness and virtual carrier sensing with directional antennas in 802.11 networks," tech. rep., Dept. of Computer Science, University of Colorado, Boulder, May 2004.

[45] H. Singh and S. Singh, "A mac protocol based on adaptive beamforming for ad hoc networks," in *Proceedings of PIMRC*, 2003.

[46] R. R. Choudhury, X. Yang, N. H. Vaidya, and R. Ramanathan, "Using directional antennas for medium access control in ad hoc networks," in *Proceedings of ACM MOBICOM*, September 2002.

[47] R. R. Choudhury and N. H. Vaidya, "Deafness: A MAC problem in ad hoc networks when using directional antennas," tech. rep., Proceedings of the IEEE International Conference on Network Protocols (ICNP), October 2004.

[48] R. R. Choudhury and N. H. Vaidya, "Impact of directional antennas on ad hoc routing," in *Proceedings of Personal and Wireless Communication (PWC)*, September 2003.

[49] R. Schlub, "Practical realization of switched and adaptive parasitic monopole radiating structures," tech. rep., PhD Thesis, Griffith University, January 2004.

[50] M. Inc, "Motia's javelin enhances wifi," tech. rep., http://www.motia.com/javelin.htm, 2004.

[51] P. M. Inc., "Paratek's smart scanning drwin antennas," tech. rep., http://www.paratek.com/products.htm, 2005.

[52] K. Sundaresan and R. Sivakumar, "A unified mac layer framework for ad-hoc networks with smart antennas," in *Proceesinngs of ACM International Symposium on Mobile Ad hoc Networking and Computing (MOBIHOC)*, May 2004.

[53] S. Vasudevan, J. Kurose, and D. Towsley, "On neighbor discovery in wireless networks with directional antennas," in *Proceedings of INFOCOM*, March 2005.

[54] A. Spyropoulos and C. Raghavendra, "Energy efficient communications in ad hoc networks using directional antennas," in *Proceedings of Infocom*, 2002.

[55] J. Tang, G. Xue, C. Chandler, and W. Zhang, "Interference-aware routing in multihop wireless networks using directional antennas," in *Proceedings of INFOCOM*, March 2005.

[56] L. Bao and J. Garcia-Luna-Aceves, "Transmission scheduling in ad hoc networks with directional antennas," in *Proceedings of Mobicom*, 2002.

[57] R. Ramanathan, J. Redi, C. Santivanez, D. Wiggins, and S. Polit, "Ad hoc networking with directional antennas: A complete system solution," in *To Appear in Journal of Selected Areas in Communications*, Jan 2005.

[58] R. R. Choudhury, T. Ueda, and J. Bordim, "Beamnet: Ad hoc networking testbed using beamforming antennas," in *Proceedings of IEEE Vehicular Technology Conference (VTC)*, September 2005.

[59] A. Kochut, A. Vasan, A. UdayaShankar, and A. Agrawala, "Sniffing out the correct physical layer capture model in 802.11b," in *Proceedings of the IEEE International Conference on Network Protocols (ICNP)*, October 2004.

[60] J. Boer, "Wireless lan with enhanced capture provision," in *U.S.Patent 5987033, Technical Report*, 1999.

[61] S. Das, C. Perkins, and E. Royer, "Ad hoc on demand distance vector (AODV) routing," Internet-Draft Version 4, IETF, October 1999.

[62] T. Ohira and K. Gyoda, "Electronically steerable passive array radiator (espar) antennas for low-cost adaptive beam forming," in *IEEE International Conference on Phased Array Systems*, May 2000.

[63] D. Maltz, J. Broch, and D. Johnson, "Quantitative lessons from a full-scale multi-hop wireless ad hoc network testbed," in *Proceedings of Wireless Communications and Networking Conference (WCNC)*, September 2000.

[64] K. Chin, J. Judge, A. Williams, and R. Kermode, "Implementation experience with manet routing protocols," in *Proceedings of ACM SigComm Computer Communications Review*, 2002 November.

[65] E. Royer and C. Perkins, "An implementation study of aodv routing protocol," in *Proceedings of Wireless Communications and Networking Conference (WCNC)*, September 2000.

[66] H. Lundergren, D. Lundberg, J. Nielsen, E. Nordstrom, and C. Tschudin, "A large-scale testbed for reproducible ad hoc protocol evaluations," in *Proceedings of Wireless Communications and Networking Conference (WCNC)*, March 2002.

[67] Y. Hu and D. Johnson, "Design and demonstration of live audio and video over multihop wireless ad hoc networks," in *Proceedings of IEEE Military Commuications Conference (MILCOM)*, October 2002.

[68] C. Toh, R. Chen, M. Delwar, and D. Allen, "Experimenting with ad hoc wireless network on campus: Insights and experiences," in *ACM Sigmetrics Performance Evaluation Review*, December 2000.

[69] D. Niculescu and B. Nath, "Vor base stations for indoor 802.11 positioning," in *Proceedings of Mobicom*, 2004.

Author's Biography

Romit Roy Choudhury received the B.Tech degree in Computer Science from Haldia Institute of Technology, India, in 2000, and the M.S degree in Electrical and Computer Engineering from the University of Illinois at Urbana-Champaign, in 2003. He joined the Computer Science department in 2003 to continue with a Ph.D program in the area of wireless networking and systems. Since 2001, he has been working as a member of the Wireless Networking Group, led by Professor Nitin Vaidya. He was awarded the Motorola Center for Communications Fellowship for the years 2004 and 2005. He was also awarded the Vodafone-US Foundations Fellowship in the College of Engineering in UIUC for the year 2006.