ENERGY SAVING AND PARTITION BRIDGING USING DIRECTIONAL ANTENNAS IN MOBILE AD HOC NETWORKS

Amit Kumar Saha

Thesis: Master of Science
Computer Science
Rice University, HOUSTON, TEXAS (MAY 2003)

RICE UNIVERSITY

Energy Saving and Partition Bridging using Directional Antennas in Mobile Ad Hoc Networks

by

Amit Kumar Saha

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

Master of Science

APPROVED, THESIS COMMITTEE:

Dr. David B. Johnson, Chair
Associate Professor of Computer Science

Dr. Willy Zwaenepoel
Karl F. Hasselman Professor of
Computer Science

Dr. Edward W. Knightly
Associate Professor of
Electrical and Computer Engineering

HOUSTON, TEXAS

MAY, 2003

Energy Saving and Partition Bridging using Directional Antennas in Mobile Ad Hoc Networks

Amit Kumar Saha

Abstract

In this thesis, I present the design and evaluation of new techniques for using directional antennas to save energy and to bridge network partitions in a mobile ad hoc network. This thesis advocates close but simple collaboration between the routing layer and the Medium Access Control (MAC) layer and shows through simulations, the effectiveness of this design by modifying the Dynamic Source Routing (DSR) protocol, an on-demand ad hoc network routing protocol based on source routing. First, in order to save energy, Route Requests and data packets are transmitted directionally. Extensive simulations show that without affecting the behavior of the routing protocol noticeably, energy savings of up to 75% is achieved. Second, in order to bridge network partitions, the routing protocol is modified to use the ability of a directional antenna to transmit directionally over longer distance as compared to an omnidirectional antenna, both antennas using the same power. Again, through simulations, the protocol is shown to be able to bridge network partitions. Also, when no partitions are present, the protocol is otherwise equivalent to the version without the partition bridging modifications.

Acknowledgments

I thank my advisor Dr. David B. Johnson for his guidance throughout the process of this thesis. The discussion sessions that we had were the most helpful steps in getting this work and this thesis to its present form.

I am grateful to the other members of my reseach group: Santashil, Yih-Chun and Jorjeta who helped me out in many small but important situations. I am also grateful to Rajnish who along with Santashil provided me with valuable reviews which helped in polishing this thesis.

Last but not the least, I thank the other members of my thesis committee: Dr. Willy Zwaenepoel and Dr. Edward W. Knightly for their reviews and suggestions which greatly improved the quality of this thesis.

Contents

Illustrations

Chapter 1

Introduction

An ad hoc network is a group of mobile wireless nodes that dynamically forms a network without the aid of any existing centralized administration or network infrastructure. Since a wireless node has limited transmission range, nodes need to cooperate to forward packets for each other so that a node can send packets to another node not in its direct transmission range. Among other issues, there are two primary concerns in a mobile wireless environment, the power consumed by the nodes and the creation of network partitions due to the change in relative distance between nodes, both of which can render a network useless.

1.1 Motivation

In this thesis I present the design and evaluation of the use of directional antennas to save power and to bridge network partitions in a mobile ad hoc network. In doing so I make the *routing layer* and the *Medium Access Control* (MAC) layer to closely collaborate with each other, and I show that there is much to be gained by such a cooperation. Traditionally, omnidirectional antennas have been used for wireless transmission. However, in recent years, directional antennas have become practical, and there has been considerable interest in harnessing the potential of using a directional antenna in a mobile ad hoc network. The relevance of directional antennas for ad hoc networks have been explained, for example, by Ramanathan [14]. Some of the advantages that a directional antenna has over an omnidirectional antenna are as follows:

- A directional antenna can transmit directionally and hence cause less interference to receivers that are not in the direction of transmission. This is unlike an omnidirectional antenna, which causes equal interference in all directions around itself. This property of a directional antenna has the potential to increase the effective throughput of the network. However, in this thesis I do not attempt to address this issue.

- For a given transmission power, a directional antenna can transmit over longer distance in a particular direction as compared to an omnidirectional antenna. This is because a directional antenna is able to use most of its power in the direction of transmission, whereas an omnidirectional antenna uses the power to transmit equally in all directions and hence transmits over shorter distances. This also implies that a directional antenna will use less power than an omnidirectional antenna to transmit over a given distance in a particular direction. In this thesis I use this property to achieve my goals of saving power and bridging network partitions in a mobile ad hoc network.

1.2 Related Work

Most of the effort towards using directional antennas in mobile ad hoc networks has been concentrated on and limited to the MAC layer and have been targeted towards increasing the throughput of the network. For example, Ko et al [9] have designed new MAC protocols for use in an ad hoc network using directional antennas and have shown throughput improvements for these protocols. Nasipuri et al [12] designed a MAC protocol to extract higher throughput from the network. Nasipuri et al [11] designed an on-demand routing protocol for use with directional antennas for reducing the number of routing packets transmitted during *route discovery*. However, their simulations do not model the MAC layer or the physical layer, and hence the effects of collisions and interference are not reflected in their results. Ramanathan and Hain [15] used directional antennas coupled with adjusting the transmission power to control the topology of multi hop wireless networks.

However, in this thesis I target two other uses of a directional antenna, *saving power* and *bridging network partitions*, which follow immediately from the fact that a directional antenna is better able to channel its energy in the direction of transmission. Some work has been done in this field as well and are as follows. Wieselthier et al [19] considered connection oriented *multicast* traffic and quantitatively analyzed the benefits obtained in saving power by using a directional antenna. Spyropoulos and Raghavendra [16] presented an energy efficient routing and scheduling algorithm in which they minimize the total time for all possible transmitter-receiver pairs to communicate with each other. However, their protocol unrealistically assumes predictability of end-to-end traffic pattern. Moreover, their protocol has not been tested under realistic mobile scenarios. My hypothesis is that in order to get the maximum benefits from using directional antennas in mobile ad hoc networks, we cannot rely on just the MAC layer taking intelligent decisions. Rather, we should utilize cooperation between the routing layer and the MAC layer, as will be proven by this thesis.

1.3 Contributions

I modified DSR to use directional antennas for energy savings and have shown that this modified protocol provides substantial energy savings (as high as 75%) without noticeably changing the overall behavior of the protocol. In order to bridge network partitions, I modified DSR to use directional antennas to transmit directionally over longer distance. I have shown, through simulations, that this modified protocol is able to deliver packets across network partitions. Also, in the absence of partitions in the network, this protocol is otherwise essentially equivalent to the version without the partition bridging modificiations.

To enable me to do realistic simulations I extend the *ns-2* network simulator [13] to include:

- the ability of a node to use a directional antenna,
- a realistic directional antenna model, and
- a propagation model that takes into consideration not only the transmission power but also the direction of transmission.

The simulator does not model the time required to change the attributes (Chapter 2) of a directional antenna. The simulator also does not model any energy loss that might occur when a directional antenna changes its attributes. Also, I assume an electronically steerable directional antenna and neglect both the energy spent to steer the antenna beam to a particular direction and the energy spent to change the transmission power of the antenna.

The results in this thesis are based upon simulations of various scenarios in an ad hoc network of 50 mobile wireless nodes communicating with each other.

1.4 Organization

The rest of the thesis is structured as follows. In Chapter 2, I describe the mathematical model of a directional antenna that I have used. In Chapter 3, I explain the Dynamic Source Routing (DSR) protocol, the routing protocol that I have used for evaluating my design. In Chapter 4, I present the design and evaluation of my protocol to save power in a mobile ad hoc network. Then, in Chapter 5, I present my design to bridge network partitions and also evaluate my design. Finally, in Chapter 6, I conclude and describe areas for further improvements.

Chapter 2

Background: Antenna Model for Directional Antenna

The concepts needed to better understand the differences between a directional antenna and an omnidirectional antenna are presented by Ram Ramanathan [14]. However, since the concepts are fundamental to justifying my design decisions, this chapter revisits the concepts in an intuitive manner. Readers interested in going to the depths of this field can refer to [3].

An ideal omnidirectional antenna transmits as well as receives energy equally well in all directions. An ideal directional antenna transmits and receives more energy in one direction called the *primary direction* of the antenna. The *gain* of a directional antenna [3] in a particular direction $\vec{d}$ is given by

$$G(\vec{d}) = \eta \frac{U(\vec{d})}{U_{avg}}$$

where,

$U(\vec{d})$ is the power density in direction $\vec{d}$

U_{avg} is the average power density over all directions

η is the efficiency of the antenna that accounts for energy losses.

The maximum gain taken over all directions is called the *peak* gain of the antenna. An *antenna pattern* is the specification of the different gain values in each direction in space. A directional antenna has a *main lobe* of peak gain and several *side lobes* of lesser gain, as shown in Figure 2.1. For the purpose of analysis all the side lobes are collectively

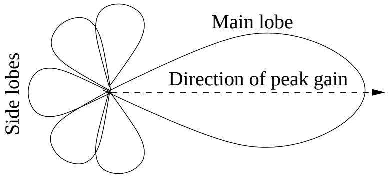

Figure 2.1: Hypothetical 2D Directional Antenna Pattern

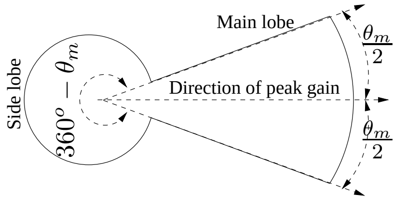

Figure 2.2: Approximate Hypothetical 2D Directional Antenna Pattern

approximated by a single side lobe, as shown in Figure 2.2. This is actually a simplification into 2 dimensions of the actual antenna pattern in 3 dimensions as shown in Figure 2.3.

The *beamwidth* of a directional antenna is the angle subtended by the two directions on either side of the direction of peak gain that are 3 dB lower in gain, as shown in Figure 2.2. An antenna is more directional if it has a higher gain and a smaller beamwidth. However, two antennas with different beamwidths can have the same gain.

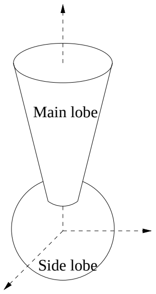

Figure 2.3: Approximate Hypothetical 3D Directional Antenna Pattern

For this work, an approximate model for a directional antenna. The antenna pattern for a given gain value g_m consists of a main lobe of beamwidth θ_m and a side lobe of beamwidth $360^o - \theta_m$, as shown in Figure 2.2. For simplicity I do not model the 3 dB loss in gain on either side of the primary direction but consider the gain to be g_m for the entire beamwidth. The maximum beamwidth [14] θ_{max} for a given gain g_m is given by

$$\theta_{max} = 2tan^{-1}\sqrt{\frac{4}{g_m}}$$

The gain for the side lobe (g_s) is given by

$$g_s = \frac{\eta\triangle - g_m}{\triangle - 1} \qquad where \;\; \triangle = \frac{4}{tan^2(\frac{\theta_m}{2})}$$

If θ_{max} is considered to be 180^o then the maximum value of g_m turns out to be 0 dB = 1.0 (absolute value). So, if the required gain be g with beamwidth $\theta > 180^o$ in direction d then the antenna has to be directed in a direction diametrically opposite to d and should use a main lobe gain $g_m < 1.0$ with beamwidth $360^o - \theta$. If g_m is chosen correctly, then the gain of the side lobe g_s (which has beamwidth $360^o - (360^o - \theta) = \theta$) matches the desired gain g.

Chapter 3

Background: The Dynamic Source Routing Protocol

In this chapter I give a brief overview of the basic functionality of the Dynamic Source Routing (DSR) protocol [5, 7, 6], the routing protocol upon which I have based my design and implementation. I use DSR in my experiments since the the protocol was shown to perform well in simulation studies carried out earlier [1, 4].

DSR is a totally *on-demand* (or *reactive*) routing protocol for mobile ad hoc networks. DSR is based on *source routing*, in which the originator of a packet decides the entire sequence of hops through which the packet is to be forwarded to the final destination. The protocol uses *Route Discovery* and *Route Maintenance* to maintain source routes to arbitrary destination nodes.

If a node does not have a source route to a destination node, then it initiates Route Discovery by locally broadcasting a ROUTE REQUEST packet containing the address of the destination (known as the *target* of the Route Discovery). The ROUTE REQUEST packet also contains a *request identifier* from the source (also known as the *initiator*) of the Route Discovery. The Route Discovery is uniquely identified by the request identifier, the source node address, and the target address. Apart from these fields, the ROUTE REQUEST also has a list of nodes, used to record the route through which the ROUTE REQUEST has been forwarded.

When a node receives a ROUTE REQUEST, it first checks to see if it has previously forwarded a ROUTE REQUEST from this Route Discovery, by examining the source address, the target address, and the request identifier. If the node has recently seen this identifier,

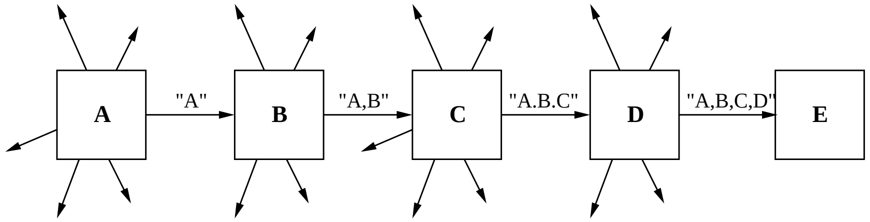

Figure 3.1: Route Discovery example : Node **A** is the initiator and node **E** is the target.

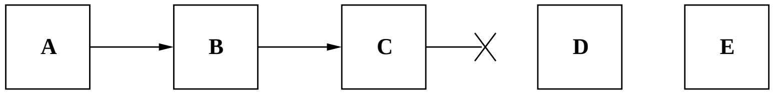

Figure 3.2: Route Maintenance example : Node **C** is unable to forward a packet from **A** to **E** over its link to next hop **D**.

or if its own address is already present in the list of nodes that this ROUTE REQUEST has traversed, the node silently drops the packet. Otherwise, the node appends its address to the list and locally rebroadcasts the ROUTE REQUEST, as shown in Figure 3.1. When a ROUTE REQUEST reaches the target node, the target node replies with a ROUTE REPLY packet destined to the source of the ROUTE REQUEST. This ROUTE REPLY packet contains a copy of the node list from the ROUTE REQUEST. When the initiator of the request receives the ROUTE REPLY, it adds the newly acquired route to its Route Cache.

In Route Maintenance, a node forwarding a packet for a source tries to verify that the packet successfully reached the next hop in the route. A node confirms this by either using a link-layer acknowledgment (such as is provided in IEEE 802.11 [2]), or a passive acknowledgment [8], or by means of a network-layer acknowledgment.

If a packet is not acknowledged, after a limited number of retransmission attempts, the forwarding node assumes that the next-hop destination is unreachable over this link (as shown in Figure 3.2), and sends a ROUTE ERROR packet indicating the broken link to the source of the packet. A node receiving a ROUTE ERROR removes that link from its Route Cache.

Chapter 4

Energy Saving using Directional Antenna

4.1 Protocol Modifications

In this Chapter, I describe how I modify the behavior of the protocol to achieve energy savings. Figure 3.1 illustrates a Route Discovery where node *A* is trying to find a source route to node *E*. The protocol modification is based on the fact that node *B* need not transmit the ROUTE REQUEST packet omnidirectionally, since there is essentially no benefit to be obtained from transmitting the ROUTE REQUEST in the direction from which it arrived at *B*. However, this is not entirely true and there are cases in which a node which is present in the overlap region of the *ideal* transmission range of *A* and *B* can be reached only by a transmission from *B* and not by a transmission from *A*. I explain later how the protocol handles such cases. With a directional antenna, *B* can transmit the packet in all directions apart from the direction of arrival of the packet. It is assumed that upon reception of a packet, a directional antenna can identify the direction of arrival of the packet with reasonable accuracy. Additionally, upon reception of any packet from a nearby node (irrespective of whether the packet was sent directly to the receiver or the receiver promiscuously received that packet), the receiver maintains, in some data structure, the estimated direction of that nearby node.

Specifically, in order to save energy using directional antennas, I made the following changes to the basic operation of DSR.

When a node *A* receives a ROUTE REQUEST from direction d_1, node *A* forwards the

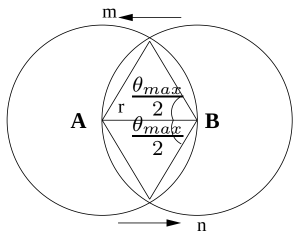

Figure 4.1: ROUTE REQUEST using directional antenna (*r* is the omnidirectional transmission range)

ROUTE REQUEST in a direction d_2 that is diametrically opposite to d_1. Also, *A* transmits the ROUTE REQUEST with beamwidth $360^o - \theta$, where $\frac{\theta}{2}$ is the angle on either side of d_1 where the ROUTE REQUEST *need not* be transmitted, as shown in Figure 4.1 (The figure shows the maximum value of $\frac{\theta}{2}$, $\frac{\theta_{max}}{2}$). However, there is the minute possibility that there exists a node *C* that lies in the region of overlap of the omnidirectional transmission ranges of *A* and *B* but *C* is not directly reachable from node *A* (maybe due to a physical obstruction) and is reachable from *B* only if *B* transmits omnidirectionally. The protocol handles such possibilities simply by reverting back to omnidirectional transmission. I now mathematically deduce the maximum value of θ (θ_{max}), required for correct behavior of the routing protocol when used with a directional antenna.

I assume that the omnidirectional antenna is ideal (i.e. has a spherical transmission range, circular in a two-dimensional plane) and that the transmission range of all the antennas used is the same, denoted by r in Figure 4.1. From Figure 4.1, it is clear that $\frac{\theta_{max}}{2} = \frac{180^o}{3} = 60^o$ since all the sides of the triangle are of length r, the transmission range of the antenna. Hence, the maximum value of $\theta_{max} = \frac{360^o}{3} = 120^o$. Thus, node *B* can safely leave out θ_{max} from the direction of arrival of the ROUTE REQUEST packet. This is because that part of space was already covered by the ROUTE REQUEST transmitted by *A*.

When a node *A* has a unicast packet to send to node *B* which is in the direction d_{AB}

relative to *A*, then *A* transmits the packet directionally towards *B* with a beamwidth θ_{AB}. However, if *A* does not have an estimate of the direction of *B* with respect to itself, then *A* behaves conservatively and transmits the packet omnidirectionally. θ_{AB} is a network-wide beamwidth used by any transmitting node when transmitting a unicast packet. (I experimented with several values of θ_{AB} in Chapter 4.2.) The channel access protocol is based on the Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) approach. However, when using a directional antenna as opposed to an omnidirectional antenna, the MAC layer uses information provided by the routing layer in order to decide whether to transmit the packet directionally or omnidirectionally. Since Request-To-Send (RTS) and Clear-To-Send (CTS) [2] packets are effectively *broadcast* packets, they are transmitted omnidirectionally, as is any packet that has to be retransmitted. Although, transmitting RTS and CTS packets omnidirectionally means that most of the throughput improvements of transmitting directionally would be lost, I consider this to be a separate MAC layer issue, which has been handled previously [9, 12]. Additionally, Address Resolution Protocol (ARP) packets, which are very few in number, are also sent omnidirectionally.

Intuitively, it follows that the smaller the beamwidth θ_{AB}, the higher the directionality of the antenna and the lower the energy required to transmit over a given distance. To illustrate the energy savings by using a directional antenna, I considered two commonly used propagation models, the *Free space model* and the *Two-ray ground reflection model* [13]. In both these models, the received signal power [18] at a distance $P_r(d)$ from the transmitter is given by

$$P_r(d) \propto P_t G_t G_r$$

where P_t is the transmitted signal power, and G_t and G_r are the antenna gains of the transmitter and the receiver, respectively. In this case, $G_r = 1.0$ for both omnidirectional anten-

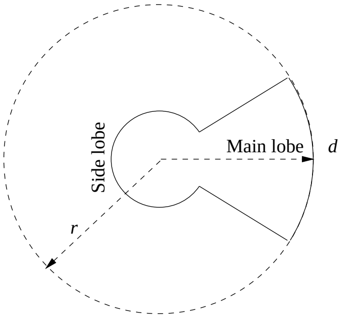

Figure 4.2: Power control to keep the transmission range same as omnidirectional antenna (*r* is the omnidirectional transmission range, and *d* is the directional transmission range, the transmission power is controlled to make *d* equal to *r*)

nas and directional antennas. To achieve a fixed received signal power level at a particular distance from the transmitter, $P_r(d)$ is a constant and thus:

$$P_t \propto \frac{P_r(d)}{G_t} \propto \frac{1}{G_t}$$

that is, required transmission power is inversely proportional to the gain of the transmitter.

In both of these modifications (transmitting a ROUTE REQUEST packet and transmitting a unicast packet), the transmitting node decreases its power such that the effective distance over which the packet can be received (in the direction of transmission) remains the same as in the case of omnidirectional transmission, as shown in Figure 4.2. Even in the presence of obstacles, the distance covered by a directional transmission in the direction of the obstacle will be the same as the distance covered by an omnidirectional transmission in the direction of the obstacle. This happens because the obstacle attenuates both the transmissions equally. A similar technique has been previously used by Takai et al [17] along with directional carrier sensing to increase the throughput of an ad hoc network rather than to save energy.

4.2 Evaluation Methodology

I based my evaluation of the use of directional antennas on the base version of DSR, which uses omnidirectional antennas, as described in Chapter 3. I used the *ns-2* network simulator, with mobility extensions from the Monarch project [10]; version 2.1b8a of *ns-2* was used, with the standard path cache data structure for the Route Cache [1]. Additionally, I extended the antenna model and the propagation model to simulate the antenna pattern of a directional antenna. This version of *ns-2* simulator models the physical layer and the MAC layer and includes modeling of contention, collisions, capture, backoff, and propagation, both for omnidirectional antennas and directional antennas. The network interface is modeled after the Lucent/Agere WaveLAN/ORiNOCO IEEE 802.11 product, which has a nominal transmission range of 250 m and a data rate of 2 Mbps; for the omnidirectional antenna case, the network interface uses the IEEE 802.11 Distributed Coordination Function (DCF) [2] MAC protocol, which employs physical and virtual carrier sensing for collision avoidance. When a directional antenna is used, the behavior of the simulator is changed to simulate the physical layer and the MAC layer, as explained in Chapter 4.

I present simulation results based on 40 randomly generated scenarios (for each pause time, explained below), each involving 50 mobile nodes moving about in a rectangular area 1500 m × 300 m for 900 simulated seconds. Nodes in the simulations move according to the Random Waypoint model [1]. In this model each node begins at a randomly chosen position, picks a new random position to which to move, and moves there in a straight line at a randomly chosen speed. This behavior is repeated independently by each node for the duration of the simulation run.

Before a node begins moving to its next chosen position, it remains stationary for a period called the Pause Time. In the simulations, the movement speed of each node is chosen with a maximum speed of 20 m/s, and Pause Time is varied between 0 s (a continuously moving network) and 900 s (a stationary network). Specifically, the following Pause Time values were used in the simulations: 0, 30, 60, 120, 300, 600, and 900 s.

The communication model between nodes used in the simulations is Constant Bit Rate (CBR) traffic. 10 different flows (each from a different source) send 4 packets per second each to a different destination chosen randomly for that flow. Each packet carries 512 bytes of data payload, making the minimum packet size including an IP header 532 bytes.

For DSR with directional antennas, I report results for beamwidths (for transmitting unicast packets) of 40^o, 60^o, 80^o, and 100^o. The beamwidth used for forwarding ROUTE REQUESTS is fixed at 240^o since that is the minimum angle required for correctness of the protocol, as proven in Chapter 4.

I computed five metrics for each simulation run:

- *Packet Delivery Ratio*: The total fraction of application-level data packets sent that are actually received at the intended destination node.

- *95th Percentile Packet Latency*: Computed as the 95th percentile of the packet delivery latency, which is the time elapsed from when a data packet is first sent to when it is first received at its destination.

- *Packet Overhead*: The number of transmissions of routing packets; for example, a ROUTE REPLY sent over three hops would count as three overhead packets in this metric.

- *Path Optimality*: Compares the length of routes used to the optimal hop length as determined by an off-line omniscient algorithm that assumes a maximum distance of 250 m per hop.

- *Total Energy Consumption*: The total energy, in Joules, consumed for all packet

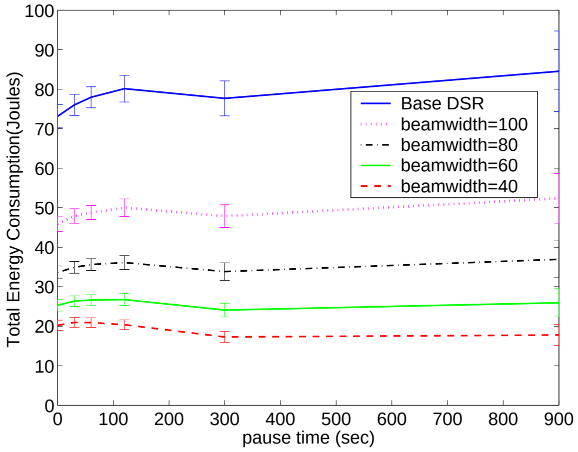

Figure 4.3: Total Energy Consumption for DSR using directional antennas for energy savings

transmissions (including overhead packets) in a scenario. The energy to transmit a single packet is calculated by multiplying the transmission power and the transmission time for that packet. I did not consider the energy spent in receiving packets.

4.3 Results

I present results averaged over 40 randomly generated scenarios. The error bars in the figures represent 99% confidence intervals of the averages shown.

Figure 4.3 shows the total energy consumption for the base DSR protocol and for four versions of the modified protocol that use directional antennas for energy saving; the four versions each use a different beamwidth for directional transmissions. Since the energy required to transmit over a fixed distance decreases with increase in directionality of a directional antenna (Chapter 4), the total energy consumed decreases substantially with the decrease in beamwidth used for directional transmissions. The energy savings over the base DSR protocol vary from 35% when using a beamwidth of 100^o, to about 75% when

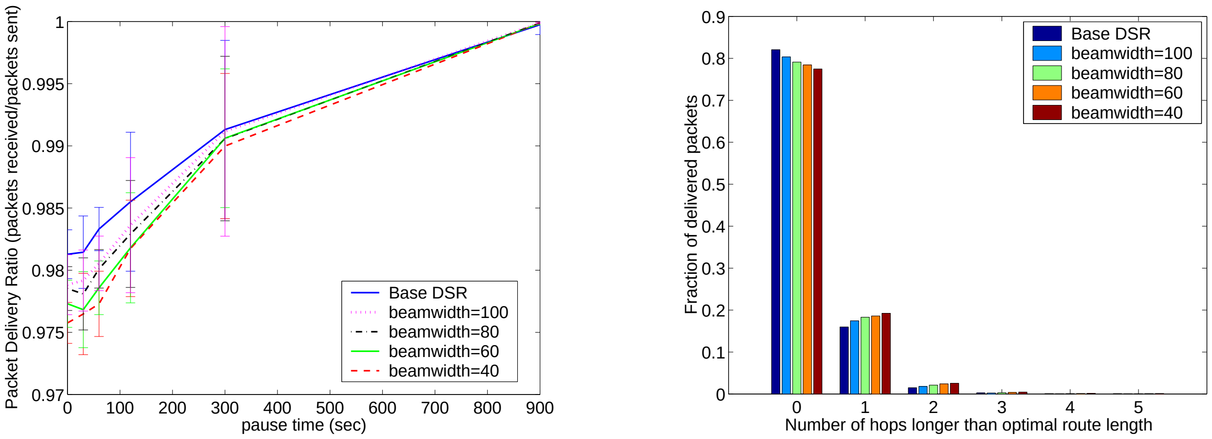

(a) Packet Delivery Ratio

(b) Path Optimality

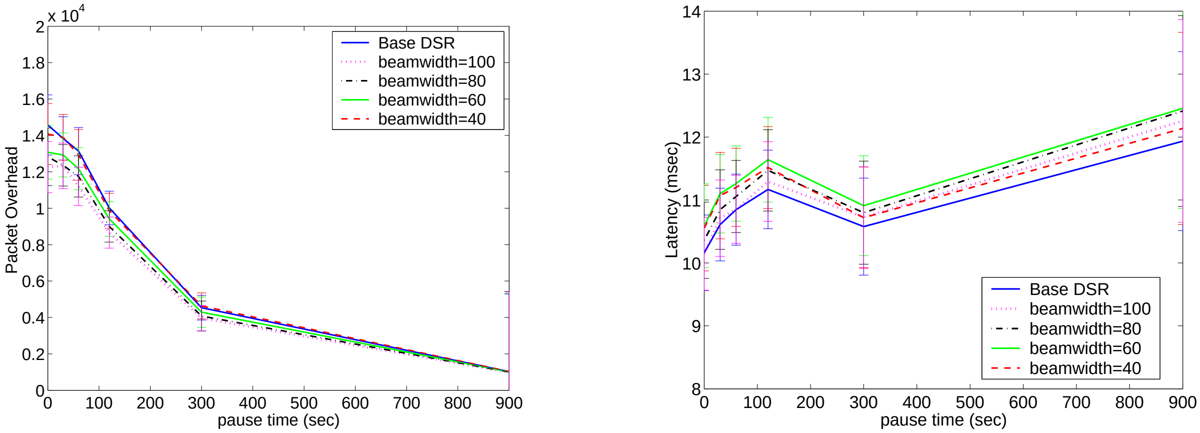

(c) Packet Overhead

(d) 95th Percentile Packet Latency

Figure 4.4: Performance evaluation results for DSR using directional antennas for energy savings

transmitting with a beamwidth of 40^o. With the advancement of technology, directional antennas will likely be able to transmit with even narrower beamwidths, thus achieving even more energy savings. The total energy savings depends primarily on the directionality of the antenna, not on the rate of mobility of the nodes. In all the runs, the protocol delivered similar number of packets.

Figure 4.4 summarizes the basic operation of the routing protocol while achieving these energy savings. Figure 4.4(a) shows the packet delivery ratio for the base DSR protocol and the four modified versions of the protocol, Figure 4.4(b) shows the corresponding path optimality, Figure 4.4(c) shows the packet overhead, and Figure 4.4(d) shows the 95th percentile packet latency. All results in Figure 4.3 and 4.4 are taken from the same set of simulation runs.

Directional transmission marginally reduces packet delivery ratio (Figure 4.4(a)). When a node transmits directionally, it uses an estimate of the direction of the next hop based on the last time the sender received (or overheard) any packet from that next hop. Between the time the direction was estimated and the time the packet is sent, the receiver node might have moved out of the area covered by the sender's beamwidth. As the beamwidth used for directional transmission decreases, the probability of the next-hop node moving out of the covered area increases.

The path length optimality decreases marginally for narrower beamwidths (Figure 4.4(b)). As the beamwidth used decreases, nodes are less likely to overhear packets from other nodes. Since DSR uses the headers of overheard packets to learn additional routing information, the protocol is less able to take advantage of such information for shortening routes in use. In addition, DSR also uses overheard packets to detect when two nodes in a route have moved close enough together such that an intermediate node between them is no longer needed in the route; for narrower beamwidths, the protocol is likewise less able to take advantage of this automatic route shortening optimization.

Packet overhead generally increases with narrower beamwidths (Figure 4.4(c)). As mentioned above, for narrower beamwidths, a next-hop node is more likely to move out of the area covered by the previous hop's transmission in the known direction toward that node, causing this node to return a ROUTE ERROR to the original sender. If the sender has no other route to the destination in its Route Cache, it must initiate a new Route Discovery, increasing the packet overhead. However, the base DSR protocol, which uses omnidirectional transmissions, has higher packet overhead than any of the modified versions that use directional transmissions. This higher overhead for the base version is likely due to the increased probability of packet collisions from the omnidirectional broadcast transmissions of ROUTE REQUEST packets, relative to the narrower directional transmission of these packets (and thus fewer nodes at which a collision could occur) with the versions of the modified protocol.

Packet latency is only slightly affected by the use of directional transmission or the choice of beamwidth (Figure 4.4(d)). Across the different versions of the protocol, there is generally a small (but not statistically significant) increase in latency with decreasing beamwidth, again due to the probability of the next-hop node moving out of the area covered by a directional transmission.

Chapter 5

Bridging Network Partitions

5.1 Protocol Modifications

The mobility of nodes in a mobile ad hoc network allows the relative distance between any two nodes to change. This mobility might lead to network partitioning, as shown in Figure 5.2. Figure 5.3 illustrates how we adaptively use the ability of a directional antenna to transmit over longer distances (as compared to the transmission range of an omnidirectional antenna using equal transmission power) to bridge such partitions when needed. To achieve these goals, we modify DSR while still maintaining the basic mechanisms of Route Discovery and Route Maintenance. In the modified protocol, however, each node additionally maintains a Passive Acknowledgment Table recording information about ROUTE REQUESTS (with the *force omnidirectional* flag set) it has received. Each entry in some nodes Passive Acknowledgement Table contains the following fields:

- Target address: The address of the node to which a source route is sought.
- Timer: Depending upon whether or not the timer has expired the owner node of this Passive Acknowledgement Table decides whether to send a ROUTE REQUEST with the *long hop* flag set.
- A list of *angular range*s around this node in which this node should search for the target address. An angular range is specified by a direction and equal angular widths on each side (clockwise and counterclockwise) of the direction.

- A list of ROUTE REQUEST packets having different source addresses but each targeted for the same target address.

Space for entries in the table is maintained in a Least Recently Used (LRU) fashion, and entries expire and are automatically deleted after a timeout.

The source route header present in a packet is also modified to have two additional flags, called *force omnidirection* and *long hop*.

We first explain how, by modifying Route Discovery in DSR, we are able to find a source route containing at least one *long* hop. The hop between two nodes is called a long hop if the protocol has determined that a directional transmission that goes over a greater distance is necessary to transmit from one node to the other. Otherwise, we consider the hop to be a *normal* hop.

A source node that has a packet to send to a destination node checks its own Route Cache for a source route to that destination. If present, the source node uses that source route for sending the data packet to the destination node. Otherwise, the source node initiates Route Discovery as follows.

If the source node does not have any pending ROUTE REQUESTS (i.e., awaiting response from the network) for the destination node, then the source node sends a ROUTE REQUEST packet omnidirectionally (as in the base DSR protocol) with the *force omnidirection* flag cleared. The flag is cleared so that an intermediate node can leave out 120^o from the direction of arrival of the ROUTE REQUEST, as explained in Chapter 4. If, however, the initiator node has already sent a ROUTE REQUEST for that destination and that ROUTE REQUEST has timed out, then the initiator node sends a new ROUTE REQUEST for the same destination with the *force omnidirection* flag set. A node receiving a ROUTE REQUEST packet behaves as follows.

If the receiving node's Passive Acknowledgment Table does not contain an entry for the target of the ROUTE REQUEST, it does the following. If the *force omnidirectional* flag in the ROUTE REQUEST is not set then the node transmits the ROUTE REQUEST packet in a direction opposite to the direction of arrival, and with a beamwidth of $360^o - 120^o = 240^o$. If, however, the *force omnidirectional* flag is set, the node forwards the ROUTE REQUEST omnidirectionally and enters the target address of the ROUTE REQUEST in its Passive Acknowledgement Table; simultaneously, the node starts the timer in the new entry and also adds a copy of the ROUTE REQUEST to the list of ROUTE REQUEST packets present in this new entry. Additionally, the node initializes the list of angular ranges for this entry to all directions within 240^o of the direction opposite to the direction of arrival of the ROUTE REQUEST. If the *long hop* flag in the ROUTE REQUEST is set, then the receiving node stores the address of the previous hop of the ROUTE REQUEST so that the node knows to send any packets to the previous hop (such as the ROUTE REPLY) with a transmission power level of *Long Hop Transmit Power*. This power level is chosen as a tradeoff between the power consumed transmitting over a long hop and the ability of this transmission to bridge large partitions. Additionally, the receiving node clears the *long hop* flag in the ROUTE REQUEST, and re-transmits the packet omnidirectionally.

Otherwise, (if this node's Passive Acknowledgement Table contains an entry for the target of the ROUTE REQUEST), this node adds the ROUTE REQUEST packet to the list of ROUTE REQUEST packets for that entry since this node might need to reply with a ROUTE REPLY to the source of the received ROUTE REQUEST. If the timer in the entry has not expired, the list of angular ranges for the entry is updated to leave out the overlap between the list of angular ranges currently present and 60^o on each side of the direction of arrival of the ROUTE REQUEST, as shown in Figure 5.1. If the timer has expired, it means that this node has transmitted one or more ROUTE REQUESTs for this Route Discovery with the *long hop* flag set has already been transmitted by this node and hence no further action is taken by this node.

If the timer in an entry in the Passive Acknowledgment Table expires, then the node (the owner of the Passive Acknowledgment Table) checks the list of angular ranges in that entry. For each angular range in the list, the node sends one or more ROUTE REQUESTs with the *long hop* flag set. More than one ROUTE REQUESTs may be necessary if the angular range is wider than the beamwidth of the ROUTE REQUEST packets. Each of these packets are sent with a beamwidth of *Long Hop Beamwidth* which, for the sake of uniformity, is kept equal to the beamwidth used for the directional transmission of unicast packets. However, these ROUTE REQUESTs are sent with a transmit power of *Long Hop Transmit Power*, so as to transmit each over a longer distance.

When a node receives a ROUTE REQUEST targeted to itself, it sends a ROUTE REPLY back to the source of the ROUTE REQUEST, as in the base DSR protocol. An intermediate node receiving a ROUTE REPLY checks its Passive Acknowledgment Table for all ROUTE REQUEST packets with a target of the originator of the ROUTE REPLY and creates a new ROUTE REPLY packet. Each of these ROUTE REPLY packets has the reply route as the concatenation of the route from the originator of the ROUTE REQUEST to the intermediate node and the route from the intermediate node to the intended target of the ROUTE REQUEST, leaving out loops if any. Additionally, the intermediate node deletes from its Passive Acknowledgement Table, the entry corresponding the source of the ROUTE REPLY. However, if there is no entry for the source of the ROUTE REPLY in the Passive Acknowledgment Table, then the intermediate node forwards the packet as in the base DSR protocol.

Once a source route has been found to the intended destination node, Route Maintenance takes over. The basic mechanism of Route Maintenance remains the same as in the base DSR protocol, with the following changes. If an intermediate node forwarding a

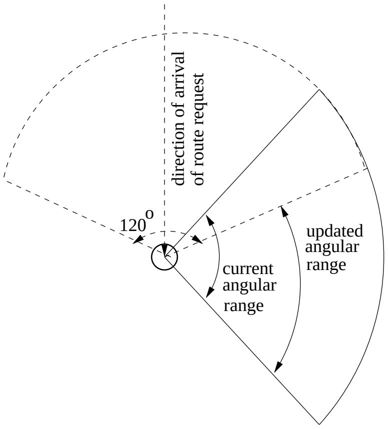

Figure 5.1: Example of updating of angular range on arrival of a ROUTE REQUEST

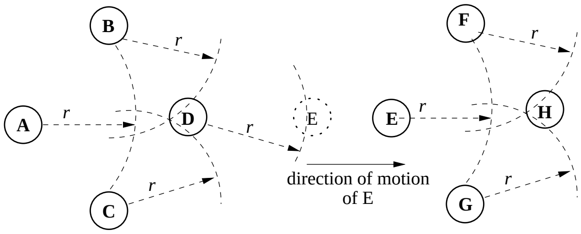

Figure 5.2: Example of network partition: Node **E** moves out of the range of node **D** (*r* is the omnidirectional transmission range

.

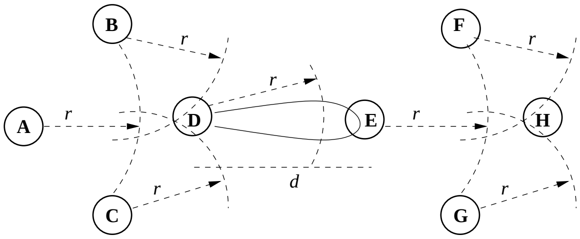

Figure 5.3: Example of bridging network partition: Node **D** transmits directionally to reach node **E** (*r* is the omnidirectional transmission range, and *d* is the directional transmission range)

.

packet for a source finds the next hop to be a long hop (since, during Route Discovery, the intermediate node stored the next hop as a long hop), then the node transmits the packet with the *long hop* flag set (it also sends the RTS and CTS directionally) and transmits the packet with the transmit power of *Long Hop Transmit Power*. If, however, the next hop is not a long hop, this node sends the packet directionally (omnidirectionally if this node does not have an estimate of the direction of the next hop), as explained in Chapter 4. If the next hop is not a long hop and the transmitted packet is not acknowledged (with a MAC acknowledgement), then the forwarding node retransmits the packet in the estimated direction of the next hop (if such an estimate is available) but with a beamwidth of *Long Hop Beamwidth* and with transmit power of *Long Hop Transmit Power*. If the intermediate node receives an acknowledgement, the intermediate node stores the next hop to be a long hop. If, however, the intermediate node still does not receive an acknowledgement, it returns a Route Error to the original sender of the packet. If the forwarding node does not have an estimated direction for the next hop, then the forwarding node returns a ROUTE ERROR after a limited number of omnidirectional retransmissions. If there is no acknowledgment to a packet sent with *Long Hop Transmit Power*, then after a limited number of retransmissions, the intermediate node assumes that the next-hop destination is unreachable and sends to the source of the packet a ROUTE ERROR indicating the broken link. As in the base DSR protocol, a node receiving a ROUTE ERROR removes the indicated link from its Route Cache.

As explained above, whenever a node considers the next hop to be a long hop, the node uses higher transmit power. Without a feedback mechanism to control this transmission power, the node would waste power when the receiving node moves nearer and would also needlessly interfere with other nearby nodes. To address this problem, whenever a node receives a packet with the *long hop* flag set, the node calculates the path loss for this

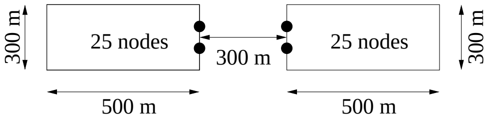

Figure 5.4: Scenario used for evaluation of partition bridging protocol (the filled black circles represent stationary nodes)

transmission as the difference between the *Long Hop Transmit Power* and the received signal strength. If this path loss is less, by at least a threshold, than the transmission power required for a directional transmission to transmit over the omnidirectional transmission range, then the receiving node piggybacks on the MAC acknowledgement packet an indication that the next hop need no longer be considered a long hop. This threshold is key to avoiding oscillations between treating the next hop as a long hop and treating it as a normal hop. Upon receiving indication, the sender treats the next hop to be a normal hop. However, we do not implement this modification in the simulation and hence do not report results for the energy usage of the protocol. In addition, we disabled DSR's automatic route shortening optimization, since packets transmitted at the *Long Hop Transmit Power* power level may falsely trigger this optimization.

5.2 Evaluation Methodology

We evaluate the partition bridging mechanism in two ways. First, we report results for how this protocol performs in the same scenarios used for evaluating our energy saving protocol in Chapter 4.2 (due to time constraints, we present results averaged over only 20 of these scenarios). These scenarios do not have any permanent network partitions. Second, we evaluate this mechanism for 20 scenarios that have a permanent partition, as explained below. Unlike in Chapter 4.2, we present results only for beamwidth of 60^o, since any

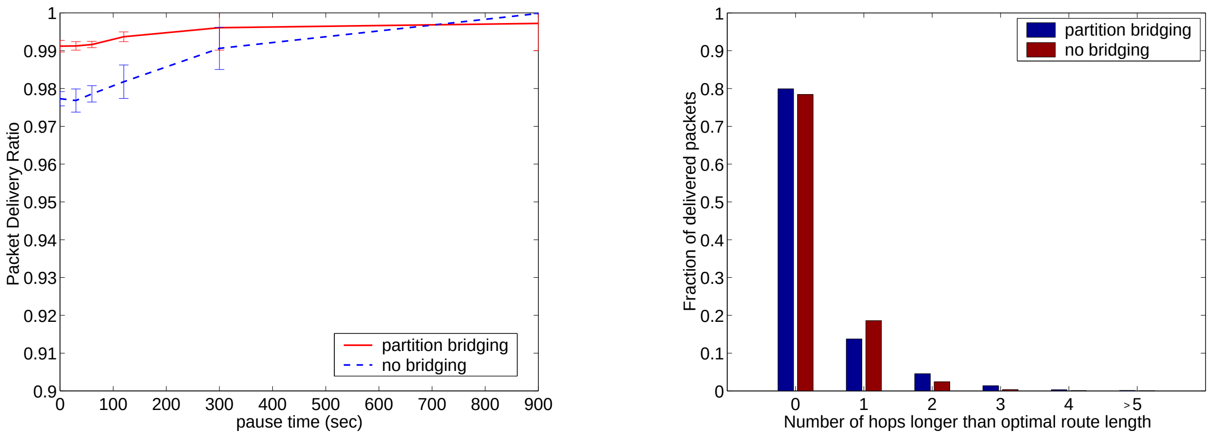

(a) Packet Delivery Ratio

(b) Path Optimality

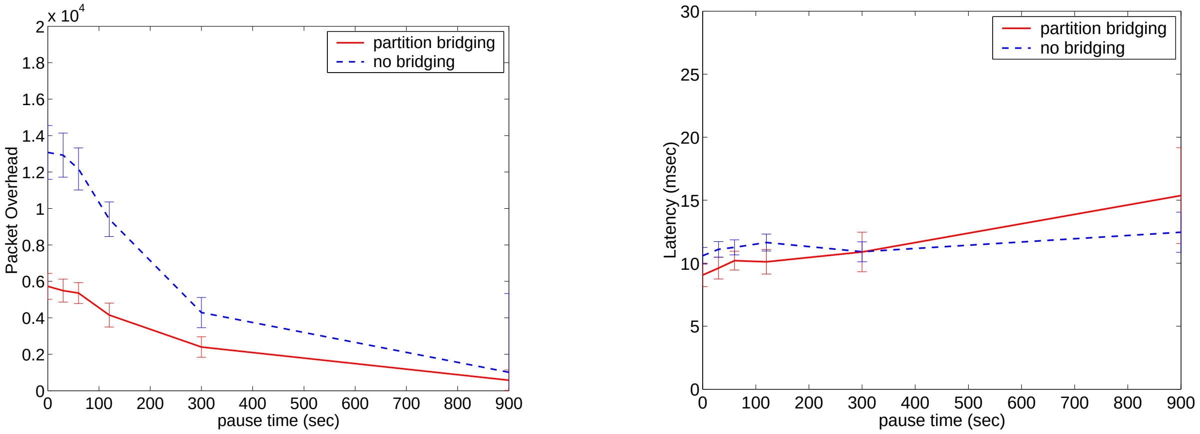

(c) Packet Overhead

(d) 95th Percentile Packet Latency

Figure 5.5: Performance evaluation results for DSR using directional antennas for partition bridging when no permanent partitions are present.

beamwidth, which would allow a directional transmission to transmit over a longer distance as compared to an omnidirectional transmission (both using the same transmit power), suffices for the purpose of bridging a network partition. Varying the beamwidth while keeping the transmission power fixed affects the distance that the directional transmission is able to reach.

For evaluating the protocol in the presence of network partitions we modify the scenario so that there are a total of 54 mobile nodes, 25 of which are moving about in an area 500 m × 300 m, while another 25 are moving about in a similar 500 m × 300 m area, with the short ends of the area separated by 300 m (the nominal omnidirectional transmission range is 250 m), as shown in Figure 5.4. Additionally, two stationary nodes are placed at each of the facing edges of the rectangular areas. These nodes are required since, in the absence of these stationary nodes, the Random Waypoint may not have any nodes close to the facing edges of the two areas. This can increase the effective distance between the two areas to a distance that could not be covered even with a directional transmission using the transmission power of *Long Hop Transmit Power*.

The communication model is also modified so that all the 10 different flows (each sending 4 packets per second, each packet having 512 bytes data payload) are from nodes in one rectangular area to nodes in the other rectangular area.

5.3 Results

Figure 5.5 summarizes the basic operation of the routing protocol including the modifications for bridging partitions, for scenarios in which no permanent partitions occur. Figure 5.5(a) shows the packet delivery ratio for the protocol with directional transmission only (Chapter 4) and for the version also including modifications for partition bridging, Figure 5.5(b) shows the corresponding path optimality, Figure 5.5(c) shows the packet overhead, and Figure 5.5(d) shows the 95th percentile packet latency.

The packet delivery ratio increase when using the partition bridging modifications, relative to the protocol version using only directional transmission without partition bridging (Figure 5.5(a)). Although in these scenarios, no permanent partitions occur, the next-hop on a route may move out of range of the transmitting node. The partition bridging code allows the transmitter to reach the next hop node without creating a broken link and requiring a new Route Discovery.

When using the partition bridging modifications, the path length optimality decreases since we do not currently use DSR's automatic route shortening optimization (Figure 5.5(b)). If the implementation were modified to include this optimization, we expect the path optimality to improve equal to the version without partition bridging.

Packet overhead decreases significantly when using the partition bridging modifications (Figure 5.5(c)). This decrease is due primarily to the absence of the automatic route shortening optimization since no gratuitous ROUTE REPLY packets are sent. In addition, however, the ability to avoid a new Route Discovery when the next-hop node on a route moves out of range of the transmitting node, as noted above, also contributes to this decrease in packet overhead.

Packet latency increases slightly when using the partition bridging modifications (Figure 5.5(d)). Although the partition bridging allows the transmitting nodes to reach the next-hop node when it moves out of normal transmission range, packets sent this way experience higher latency; such a packet is first transmitted at normal power level and then retransmitted at *Long Hop Transmit Power* after the initial transmission times out.

Figures 5.6, 5.7, and 5.8 summarize the basic operation of the routing protocol in the scenarios in which a permanent partition is present, as illustrated in Figure 5.4. These results are similar to those with no permanent partitions (Figure 5.5), even though here all

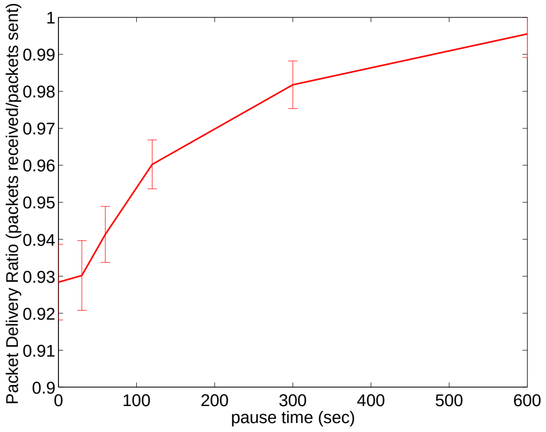

Figure 5.6: Packet Delivery Ratio with all 10 flows transmitting across a partition.

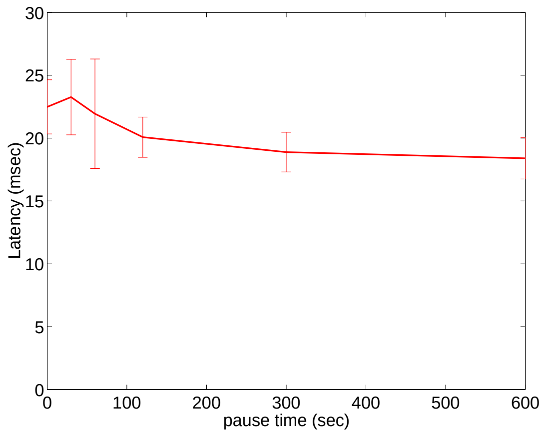

Figure 5.7: Mean Latency with all 10 flows transmitting across a partition.

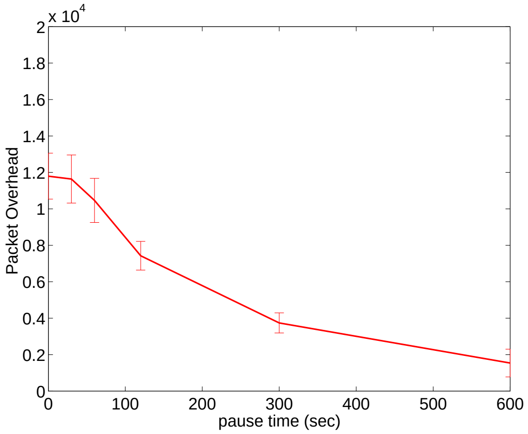

Figure 5.8: Packet Overhead amsaha with all 10 flows transmitting across a partition.

flows must bridge the 300 m separation between the two rectangular areas. We do not report path optimality results for the scenarios in which permanent partition is present, since the computation of path optimality considers only links up to length 250 m and is thus not defined in these scenarios.

Chapter 6

Conclusion

In this thesis I have presented the use of directional antennas for the purpose of energy savings and for partition bridging in mobile ad hoc networks. The first part of the thesis is based on the observation that, compared to omnidirectional transmission, directional transmission takes radically less energy to transmit a packet over a transmission distance in the intended direction of transmission. In order to observe the energy savings obtained by using directional transmission, I modified DSR to use directional antennas for energy savings and have shown that this modified protocol provides substantial energy savings (as high as 75%) without noticeably changing the overall behavior of the protocol.

The second part of the thesis proposes a novell protocol to bridge network partitions with the help of directional antennas. This protocol uses the fact that directional antennas can transmit directionally over longer distance as compared to the transmission range of an omnidirectional antenna. I implemented this protocol by modifying DSR to use directional antennas to transmit directionally over longer distance. I have shown, through simulations, that this modified protocol is able to deliver packets across network partitions. Also, in the absence of partitions in the network, this protocol is otherwise essentially equivalent to the version without the partition bridging modifications.

Bibliography

[1] Josh Broch, David A. Maltz, David B. Johnson, Yih-Chun Hu, and Jorjeta Jetcheva. A performance comparison of multi-hop wireless ad hoc network routing protocols. In *Proceedings of the Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking* (MobiCom'98), pages 85–97. ACM/IEEE, October 1998.

[2] IEEE Computer Society LAN MAN Standards Committee. *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Std 802.11-1997. The Institute of Electrical and Electronics Engineers, New York, New York, 1997.

[3] J.C.Liberti and T.S.Rappaport. *Smart Antennas for Wireless Communications*. Prentice-Hall PTR, 1999.

[4] Per Johansson, Tony Larsson, Nicklas Hedman, Bartosz Mielczarek, and Mikael Degermark. Scenario-based performance analysis of routing protocols for mobile ad-hoc networks. In *Proceedings of the Fifth Annual International Conference on Mobile Computing and Networking (MobiCom 1999)*, pages 195–206, August 1999.

[5] David B. Johnson. Routing in ad hoc networks of mobile hosts. In *Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications* (WMCSA'94), pages 158–163. IEEE Computer Society, December 1994.

[6] David B. Johnson and David A. Maltz. Dynamic source routing in ad hoc wireless

networks. In Tomasz Imielinski and Hank Korth, editors, *Mobile Computing*, chapter 5, pages 153–181. Kluwer Academic Publishers, 1996.

[7] David B. Johnson, David A. Maltz, and Josh Broch. The dynamic source routing protocol for multihop wireless ad hoc networks. In Charles E. Perkins, editor, *Ad Hoc Networking*, chapter 5, pages 139–172. Addison-Wesley, 2001.

[8] John Jubin and Janet D. Tornow. The DARPA packet radio network protocols. *Proceedings of the IEEE*, 75(1):21–32, January 1987.

[9] Y. Ko, V. Shankarkumar, and N.H. Vaidya. Medium access control protocols using directional antennas in ad hoc networks. In *Proceedings of IEEE INFOCOM 2000*, pages 13–21, Tel Aviv, Israel, March 2000.

[10] The Monarch Project. Rice monarch project: Mobile networking architectures, project home page. Available at http://www.monarch.cs.rice.edu/.

[11] A. Nasipuri, J. Mandava, H. Manchala, and R.E. Hiromoto. On-demand routing using directional antennas in mobile ad hoc networks. In *Proceedings of the IEEE International Conference on Computer Communication and Networks (ICCCN 2000)*, Las Vegas, Nevada, October 2000.

[12] A. Nasipuri, S. Ye, J. You, and R.E. Hiromoto. A mac protocol for mobile ad hoc networks using directional antennas. In *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC 2000)*, Chicago, Illinois, September 2000.

[13] The VINT Project. The *ns* manual (formerly *ns* notes and documentation). Available at http://www.isi.edu/nsnam/ns/, November 2001.

[14] R. Ramanathan. On the performance of ad hoc networks with beamforming antennas. In *Proceedings of the Second ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2001)*, pages 95–105, Long Beach, California, October 2001.

[15] R. Ramanathan and R. Hain. Toplogy control of multihop radio networks using transmit power adjustment. In *Proceedings of IEEE INFOCOM 2000*, pages 404–413, Tel Aviv, Israel, March 2000.

[16] A. Spyropoulos and C. Raghavendra. Energy efficient communications in ad hoc networks using directional antennas. In *Proceedings of IEEE INFOCOM 2002*, pages 221–228, New York, June 2002.

[17] M. Takai, J. Martin, R. Bagrodia, and A. Ren. Directional virtual carrier sensing for directional antennas in mobile ad hoc networks. In *Proceedings of the Third ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2002)*, pages 183–193, Lausanne, Switzerland, June 2002.

[18] T.S.Rappaport. *Wireless Communications, Principle and Practice*. Prentice-Hall, 1996.

[19] J. E. Wieselthier, G. Nguyen, and A. Ephremides. Energy-limited wireless networking with directional antennas: The case of session-based multicasting. In *Proceedings of IEEE INFOCOM 2002*, pages 190–199, New York, June 2002.