Routing Improvement using Directional Antennas in Mobile Ad Hoc Networks

Amit Kumar Saha, David B. Johnson
Department of Computer Science,
Rice University
07/03/2003

Abstract—
In this paper, we present the design and evaluation of two techniques for routing improvement using directional antennas in mobile ad hoc networks. First, we use directional antennas to bridge network partitions by adaptively transmitting selected packets over a longer distance, using the capabilities of the directional antenna, yet still transmitting most packets shorter distance in order to reduce power consumption and interference. Through simulations, we demonstrate the effectiveness of our design by modifying the Dynamic Source Routing (DSR) protocol, an on-demand ad hoc network routing protocol. Our simulations show that the modified protocol is able to effectively bridge network partitions, and that the protocol is otherwise equivalent to the original protocol when no partitions are present. Second, we propose a method to use directional antennas to repair routes in use when an intermediate node moves out of wireless transmission range along the route. By using the capability of a directional antenna to selectively transmit packets over a longer distance, we bridge the break in the route caused by the intermediate node's movement, thus reducing packet delivery latency and avoiding dropped packets and additional routing overhead. We present the results of simulations giving a preliminary performance evaluation of this technique demonstrating its effectiveness.

I. 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, the creation of network partitions due to the change in relative distance between nodes is of primary concern in a mobile wireless environment. The only way to bridge permanent partitions in a wireless network is to increase the transmission range. However, increasing the transmission range of omnidirectional antennas directly translates to super-linear increase in transmission power. However, directional antennas use substantially less energy to transmit directionally over similar distances. In this work we use this ability of a directional antenna to bridge network partitions. Even in the absence of permanent partitions in the network, the ability to transmit over longer distance seems attractive, and we propose a method to reduce the routing overhead associated with a routing protocol.

Relevance of Directional Antennas: 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]. However, I reiterate here, why directional antennas are feasible even for nodes which typically constitute a mobile ad hoc network. Constituent nodes in a mobile ad hoc network are typically small in size. Arguably, the biggest hurdle in using directional antennas in devices like PDAs and laptops is the size of such antennas. We argue that though presently the size of a directional antenna is a hurdle yet the future of directional antennas seems bright. At 2.4 GHz and a half-wavelength element spacing, the size of an eight element cylindrical array would be or radius 8 cm. which, of course, is too large. However, with the foreseeable increase in the operating frequency, the antenna size will shrink. The IEEE 802.11a is already working on 5GHz band, in which the size of a 8 - element antenna would have a radius of 3.3 cm. At 24GHz ISM band, which is not too far out in the future, a similar antenna will have a radius of 0.8 cm. Even with a radius of 3.3 cm it is not unthinkable that directional antennas cannot be used in a laptop. With the radius as small as 0.8 cm, PDAs are well within the scope of directional antennas.

Advantages of Directional Antennas: A directional antenna has a lot of advantages over a traditional antenna. 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 inter-

ference 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 paper we 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 paper we use this property to achieve routing improvement in a mobile ad hoc network.

To enable us to do realistic simulations we 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 (Section II) of a directional antenna, neither does it model the energy used in changing these attributes.

The rest of the paper is organized as follows. In Section II, we briefly describe the mathematical model of a directional antenna that we have used. In Section III, we explain the Dynamic Source Routing (DSR) protocol, the routing protocol that we have used for evaluating our design. In Section IV, we present our design to bridge network partitions and also evaluate our design. Then, in Section V, we present and evaluate our design to repair broken routes in the routing protocol, using directional antennas. We survey related work in this field in Section VI. Finally, in Section VII, we conclude and describe areas for further research.

II. 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 our design decisions, we revisit the concepts in an intuitive manner. Readers interested in going to the depths of this field can refer to Rappaport [3] and Ramanathan [14].

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 $\vec{d} = (\theta, \phi)$ is given by

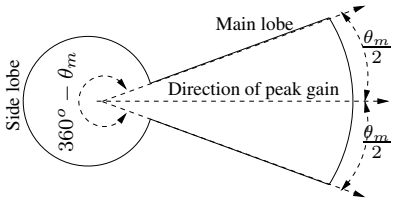

Figure 1. Approximate Hypothetical 2D Directional Antenna Pattern

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

where

$U(\vec{d})$	is the power density in direction $\vec{d}$
U_{avg}	is the average power density over all directions, and
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. For the purpose of analysis all the side lobes are collectively approximated by a single side lobe, as shown in Figure 1.

The *beamwidth* of a directional antenna (θ_m in Figure 1) 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 1. For simplicity we do not model the 3 dB loss in gain on either side of the primary direction but consider the entire main lobe to have the peak gain. 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.

III. The Dynamic Source Routing Protocol

In this section we give a brief overview of the basic functionality of the Dynamic Source Routing (DSR) protocol [5, 7, 6], the routing protocol upon which we have based our design and implementation. We use DSR in our 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.

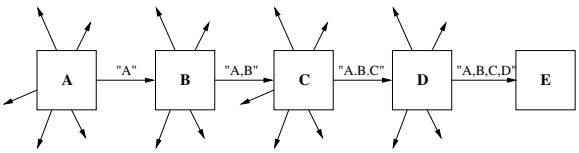

Figure 2. Route Discovery example : Node **A** is the initiator and node **E** is the target.

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, 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 2. 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), 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.

IV. BRIDGING NETWORK PARTITIONS

A. 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 4. Figure 5 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. However, the ideas presented are simply tuned towards being implemented with DSR as the base routing protocol, and can be easily applied to any other on demand routing protocol.

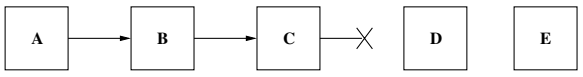

Figure 3. Route Maintenance example : Node **C** is unable to forward a packet from **A** to **E** over its link to next hop **D**.

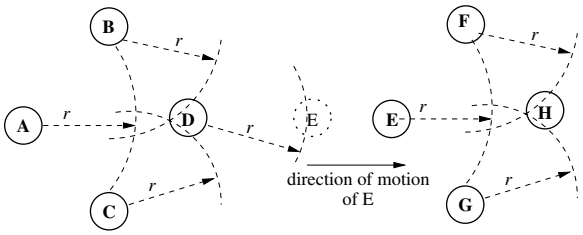

Figure 4. Example of network partition: Node **E** moves out of the range of node **D** (*r* is the omnidirectional transmission range

The basic idea behind our technique is to use the capability of a directional antenna to transmit over longer distances, but to use this capability *only* when necessary for *selected* packets. Each node in the routing protocol learns when this capability may be necessary for any given routing or data packet; most packets are transmitted using the directional antenna over a normal, shorter distance, thus reducing power consumption and interference and generally increasing the capacity of the ad hoc network.

Data Structure Modifications*:* In the modified protocol, however, each node additionally maintains a Passive Acknowledgment Table recording information about ROUTE REQUESTS (with the *trigger partition bridging*

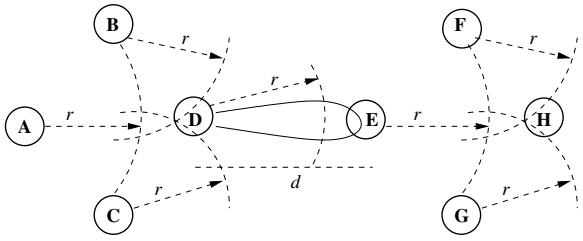

Figure 5. Example of bridging network partition: Node **D** transmits directionally to reach node **E** (*r* is the omnidirectional transmission range, and *d* is the directional transmission range)

Target Address	When inserted	Angular Ranges	ROUTE REQUESTs
5	2	List of angular ranges, initialized to 240^0 in the direction opposite to the direction of arrival of ROUTE REQUEST	List of Route Requests, from different sources
...	...	...	...

TABLE I
PASSIVE ACKNOWLEDGEMENT TABLE ENTRY

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.
- *When inserted*: Depending upon whether the difference between current time and the time in this entry has exceeded a threshold, the owner of this Passive Acknowledgement Table decides whether to initiate partition bridging in order to find a route to the destination.
- *A list of angular ranges* 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.

Schematically, an entry in the table looks as shown in Table I.
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.

We also add two flags to the source route header: the *trigger partition bridging* flag and the *long hop* flag. The protocol handles ROUTE REQUESTs differently if the *trigger partition bridging* flag is set. If the *long hop* flag is set in the source route header of a packet it indicates that the packet had been sent with a greater transmission power than normal in order to transmit the packet over a distance which is larger than the normal omnidirectional transmission range.

Modifications to Route Discovery*:* We first explain how, by modifying Route Discovery in DSR, we are able to find a source route containing one or more *long* hop. The hop between two nodes is called a long hop if the distance between these two nodes is greater than the normal omnidirectional transmission range. Otherwise, we consider the hop to be a *normal* hop. In this algorithm, it is the responsibility of the MAC layer to use different transmission power and beamwidth when a packet is sent over a long hop versus when a packet is sent over 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 *trigger partition bridging* flag cleared. The flag is cleared so that an intermediate node treats the packet normally (as in base DSR).
    - However, if 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 *trigger partition bridging* 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 *trigger partition bridging* flag in the ROUTE REQUEST is not set then the node locally broadcasts the ROUTE REQUEST packet.
    - If, however, the *trigger partition bridging* 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 enters the present time in the *When inserted* field. and also adds a copy of the ROUTE REQUEST to the list of ROUTE REQUEST packets present for 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. Additionally, the receiving node clears the *long hop* flag in the ROUTE REQUEST, and re-transmits the packet omnidirection-

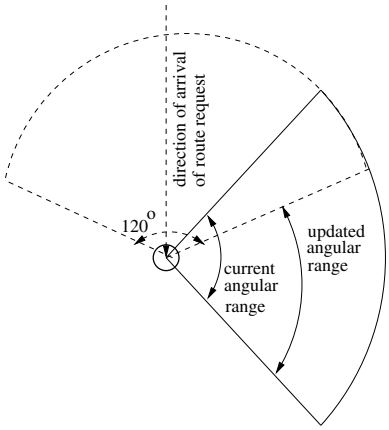

Figure 6. Example of updating of angular range on arrival of a ROUTE REQUEST

ally. The packet cannot be sent out directionally since the argument that the receiving node can leave out 120^0 from the direction of arrival of the packet does not hold if the packet is received over a long hop.

- 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 difference between the present time and the *When entered* field in the entry has not is not greater than a threshold, the list of directions 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 6.
  - However, if the difference is greater than a threshold, 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 difference between the present time and the *When entered* field 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. Each of these packets are sent considering the next hop to be a long hop. The More than one ROUTE REQUESTs may be necessary if the angular range is wider than the beamwidth of the ROUTE REQUEST packets. However, these ROUTE REQUESTs are sent with a higher 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 for each listed ROUTE REQUEST. 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.

Modifications to Route Maintenance: 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 packet for a source finds the next hop to be a long hop then the MAC layer is responsible to send the packet accordingly.
- 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).
- 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) and considering the next hop to be a long hop.
  - 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.

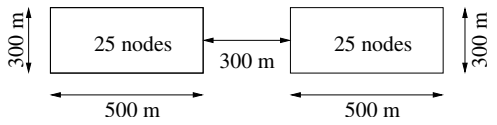

Figure 7. Scenario used for evaluation of partition bridging protocol

  - 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 over a long hop, 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. A receiver successfully receives a packet if the bit error rate is below a threshold. This indirectly implies a lower bound on the signal-to-interference-and-noise-ratio (SINR). Hence, in order to model whether a node is near enough to consider it a normal hop away rather than a long hop away, we need to model the SINR at the receiver. However, for implementation we take a much simpler approach which considers just the received signal strength. Whenever a node receives a packet with the *long hop* flag set, the node calculates the difference in the received signal strength and the minimum threshold required to correctly sense a packet. If this difference is greater than a prefixed threshold 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 prefixed threshold is key to avoiding oscillations between treating the next hop as a long hop and treating it as a normal hop. Upon receiving this indication, the sender treats the next hop to be a normal hop.

B. Evaluation Methodology

We use 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, we 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 such that the MAC protocol sends RTS and CTS omnidirectionally and sends DATA and ACK directionally. In Section VI we justify why our simple MAC protocol is similar to an existing MAC protocol [9].

We evaluate the partition bridging mechanism in two ways. First, we report results for how this protocol performs in scenarios that do not have any permanent network partitions. This is important because we consider that our protocol would be of lesser value if the protocol is able to bridge partitions but is not able to match the base DSR protocol when there are no permanent partitions in the network. Second, we evaluate the protocol in scenarios which have a permanent partition. We use a beamwidth of 60^o, since any 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 scenarios where there are no permanent partitions we present simulation results based on 20 randomly generated scenarios (for each pause time), each involving 50 mobile nodes moving about in a rectangular area 1500 m $\times$ 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.

We 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.
- *90th Percentile Packet Latency*: Computed as the 90th 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 offline omniscient algorithm that assumes a maximum distance of 250 m per hop.

For evaluating the protocol in the presence of network partitions we use a scenario in which there are a total of 50 mobile nodes, 25 of which are moving about in an area 500 m $\times$ 300 m, while another 25 are moving about in a similar 500 m $\times$ 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 7. Here as well, the nodes move according to the Random Waypoint model. For these experiments we use the pause time values of 0 s, 30 s, 60 s, 120 s, 300 s, and 600 s.

The communication model contains 5 different flows from nodes in one rectangular area to nodes in the other rectangular area (instead of the 10 flows used for the previous experiment). Since all the flows have to cross the partition, having too many flows can seriously hamper the performance by choking the nodes near the partition. This is unlike the 10 flows in the previous simulation in which the flows are randomly distributed.

Another concern in using directional antennas is the accuracy with which the antenna can identify the angle of arrival of a packet (and hence can estimate the direction of the sender), and how this accuracy affects the performance of the protocol. Hence, for the scenarios having permanent partitions we evaluate our protocol using three different accuracies in estimating the arrival of a packet, 10^o, 20^o, and 40^o.

C. Results

Figure 8 summarizes the basic operation of the routing protocol including the modifications for bridging partitions, for scenarios in which no permanent partitions occur. Figure 8(a) shows the packet delivery ratio for the protocol with directional transmission only and for the version also including modifications for partition bridging, Figure 8(b) shows the corresponding path optimality, Figure 8(c) shows the packet overhead, and Figure 8(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 8(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 but not noticeably (Figure 8(b)). This is due to some routes consisting of a longer number of hops even with a long hop, whereas if another route discovery were to be carried out then the optimum path could have been found out.

Packet overhead does not change noticeably when using partition bridging modifications (Figure 8(c)). However, as shown in higher mobility scenarios, the ability to avoid a new Route Discovery when the next-hop node on a route moves out of range of the transmitting node, also contributes to some decrease in packet overhead.

Packet latency increases slightly when using the partition bridging modifications (Figure 8(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 a higher power level after the initial transmission times out. However, as the packet delivery ratios of the two protocols approach each other the difference in packet latency become less significant.

The performance of the protocol, as shown in , though much better than a protocol using an omnidirectional antenna (which would give packet delivery ratio of zero), was not satisfactory. Upon analysis of the traces of the simulation runs we found that there were quite a large number of packets which underwent very high latency. This is due to the fact that as the nodes move about, the links which bridge the partition have a greater chance of breaking since these are directional links. Once a route is broken Route Maintenance kicks in and this leads to a new Route Discovery which takes longer than the Route Discovery in base DSR. Also, if one Route Discovery fails, then the next Route Discover attempt is made after an exponentially increasing back off period. All these contribute to the delay that a packet (sitting in the Send Buffer of the source) experiences. Such frequent route breaks also affects packet delivery ratio and packet over-

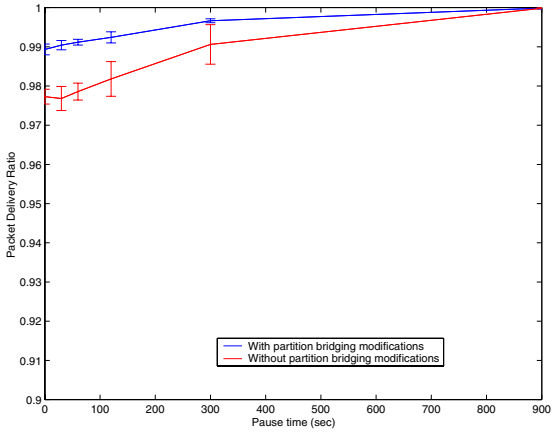

(a) Packet Delivery Ratio

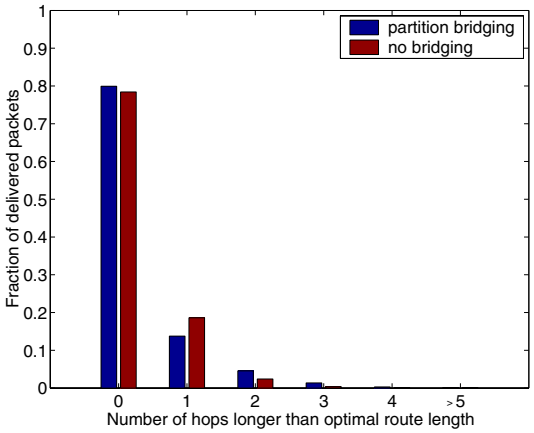

(b) Path Optimality

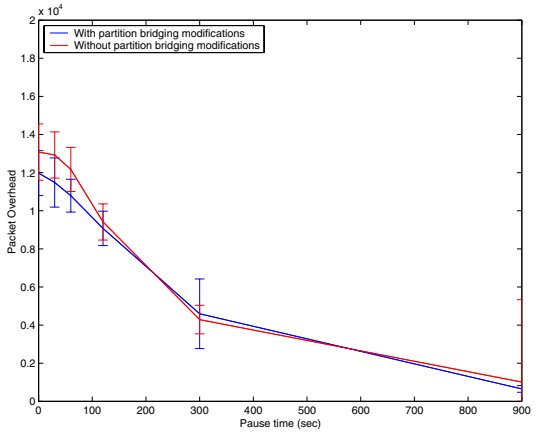

(c) Packet Overhead

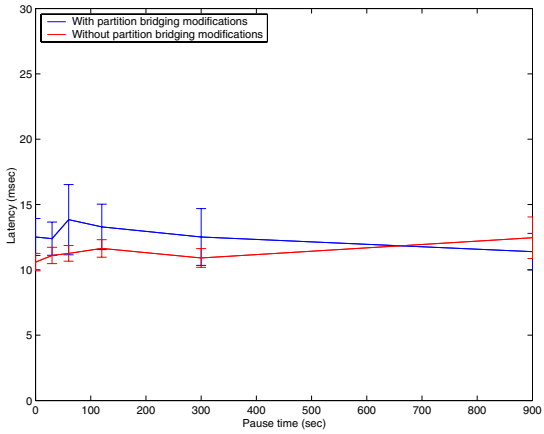

(d) 95th Percentile Packet Latency

Figure 8. Performance evaluation results for DSR using directional antennas for partition bridging when no permanent partitions are present.

head. The fact that the performance improves as the mobility in the network reduces, also supports the reasoning that we have just presented. In order to avoid this problem the routes at the edge of a network should have better methods of tracking each other (say through, GPS, etc) so that breaks in the links that bridge a partition can be made more reliable. In order to evaluate our protocol under the assumption that the nodes at the edge of the partition have some means to track each other we placed stationary nodes, at the edges of the partition, as shown in Figure 9. The rest of the nodes in the scenario exhibit the same mobility as in the previous scenario. With this scenario we were able to get much better performance and could also increase the number of flows in the communication model. We present results with 10 flows instead of 5.

As shown in Figure 10, Figure 11, and Figure 12 the protocol is not too sensitive to the accuracy to identify the angle of arrival of packets. There is noticeable difference

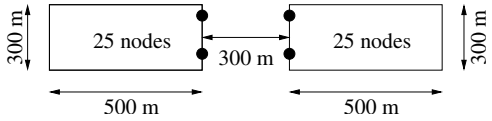

Figure 9. Modified scenario used for evaluation of partition bridging protocol (the filled black circles represent stationary nodes)

in the performance when the accuracy decreases to 20^o from 10^o. However, when the accuracy further decreases to 40^o the performance degradation is barely noticeable.

Figures 13, 14, ,and 15 summarize the basic operation of the routing protocol in the scenarios in which a permanent partition is present (along with stationary nodes at the edge of the partition), as illustrated in Figure 9. As expected adding the stationary nodes at the edge of the partition (which actually means that the nodes at the edge are better able to track each other) these results are similar to those with no permanent partitions (Figure 8), even though here all flows must bridge the 300 m separation

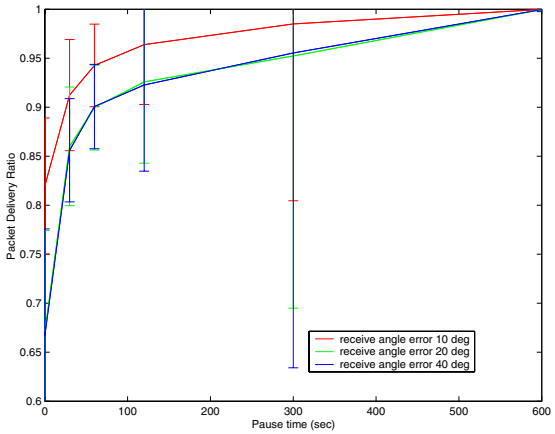

Figure 10. Packet Delivery Ratio for 5 flows transmitting across a partition.

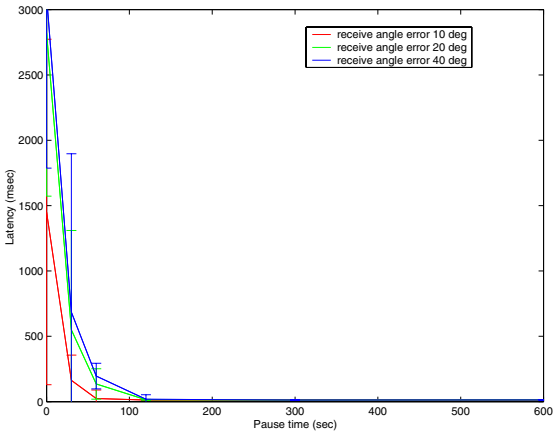

Figure 11. Mean latency for 5 flows transmitting across a partition.

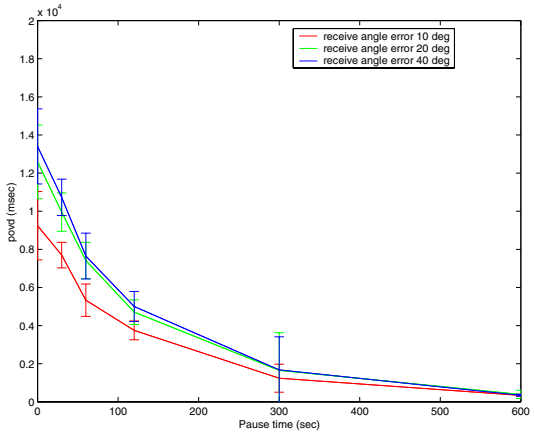

Figure 12. Packet overhead for 5 flows transmitting across a partition.

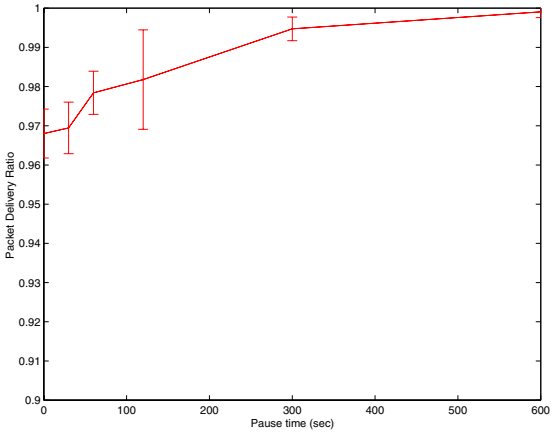

Figure 13. Packet Delivery Ratio with all 10 flows transmitting across a partition.

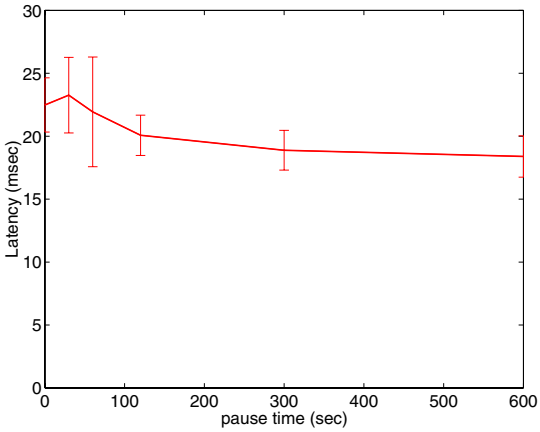

Figure 14. Mean Latency with all 10 flows transmitting across a partition.

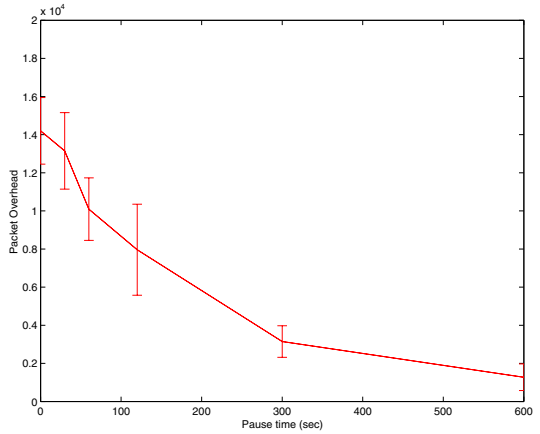

Figure 15. Packet Overhead with all 10 flows transmitting across a partition.

between the two rectangular areas. Specifically, our protocol achieves high packet delivery ratio, has low mean packet delivery latency, and modest packet overhead. The packet delivery latency is higher than the base version of DSR because the protocol first attempts to find a source route to the destination using base DSR and only when it fails does the protocol try to bridge the permanent partition.

We do not report path optimality results since the computation of path optimality considers only links up to length 250 m and is thus not defined in these scenarios.

Scenarios where Partition Bridging Would Not Work: There are certain conditions under which the partition bridging protocol as proposed would not work. Some of them are as follows:

- If the directional antenna at the other end of the partition is unable to estimate the angle of arrival of a packet (say, because of interference with other packets), then the partition cannot be bridged using directional transmission. However, instead of the receiver being unable to estimate the angle of arrival of a packet, if the antenna is able to estimate but with a large error then the problem is still solvable using directional antennas. The problem is also less severe because for long hops the area sweeped by the main lobe of the directional area is quite large and hence even if the beam is misdirected, the receiver may still lie in the main lobe of the antenna.
  However, to bridge partitions in the case that the antenna is unable to give any estimate of the angle of arrival of a packet, we would have to increase the transmission power and try to bridge the partition with omnidirectional transmission. This is similar to the idea proposed by Ramanathan and Hain [15].
- From the results we can deduce that if node mobility is very high then the protocol is less able to bridge partitions. Even if the protocol can bridge the partition, the delay incurred by the packets is prohibitively high.

V. Repairing Broken Routes using Directional Antennas

Having demonstrated the effectiveness of directional antennas in networks having permanent partitions, we now propose a new technique to also use directional antennas to augment the operation of Route Maintenance, to repair broken routes in use.

A. Protocol Description

Here, we describe how we can use directional antennas in a network without any permanent partitions and by doing so can improve the repairing of broken routes. In particular, we propose how to modify DSR (or any other on-demand routing protocol) to use directional antennas to improve Route Maintenance. As with the partition bridging protocol, the protocol description here is tuned towards using DSR as the base routing protocol, but the modifications can be applied to any routing protocol that uses source routes (actually, a node forwarding a packet need only know which is the expected next-to-next hop node for that packet).

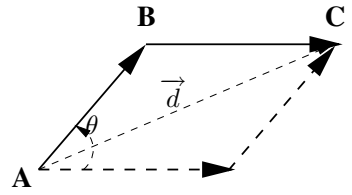

Figure 16. Estimate of direction and beamwidth from **A** to next to next hop **C**

As with our partition bridging technique, the basic idea behind this technique is to use the capability of a directional antenna to transmit over longer distances, but to use this capability *only* when necessary for *selected* packets. When forwarding a packet along some route, an intermediate node detects that the next-hop node is unreachable and may then attempt to bridge over that node to reach the following next-hop node after it. Most packets are still transmitted using the directional antenna over a normal, shorter distance; the use of directional antennas provides a mechanism for avoiding (or lessening) the routing disruption caused by the original next-hop node's movement.

Data Structure Modifications*:* In addition to maintaining the estimated direction of all next hop neighbors, a node also maintains the direction of what the next hop considers to be the direction of the next-to-next hop. For example, in Figure 16 *A* not only keeps the estimated direction of *B* but also maintains the estimated direction of *C* from *B*. *B* piggy backs the estimated direction of *C* in the ACK packet that *B* sends to *A*. This implies that *B* must have received at least one packet from *C* in the recent past.

Each node also maintains a counter associated with each next hop node. We name this the *skip counter* and it identifies the number of times that this next hop was skipped consecutively.

The source route data structure is also modified to include an additional field which identifies how many *long hop*s the packet has already encountered. The size of the field would hence depend upon the number of long hops that we want to allow for a particular route.

Modifications to Route Maintenance*:* Here we describe how we modify Route Maintenance in order to decrease the overhead of repairing broken routes.

- When an intermediate node receives a packet to forward, it first checks how many times the next hop

node has been skipped recently (i.e. checks the value of the skip counter).

- If *skip counter* = 0 (the next hop has not been skipped at all) then it means that the intermediate node has been successfully forwarding packets to the next hop. Hence, as in base DSR, the intermediate node forwards the data packet to the next hop.
- If *skip counter* $<$ THRESHOLD and *skip counter* $>$ 0 (the next hop has been skipped consecutively but not more than THRESHOLD number of times) then the data packet is forwarded to the MAC layer as in base DSR, but the MAC layer handles this packet differently. Instead of retrying multiple times to send the data packet to the next hop, the MAC layer tries to send this packet to the next hop just once. If the transmission succeeds then the MAC layer resets the *skip counter* associated with the next hop to zero. This THRESHOLD is important because deciding to skip a next hop just because one data packet could not be sent to the next hop is an extreme measure. A single data packet might be lost due to many reasons (such as sudden interference, sudden channel degradation, etc), most of which are transient.
- If *skip counter* $\geq$ THRESHOLD (the next hop HAS been skipped consecutively for more than THRESHOLD number of times) then the intermediate node modifies the source route in the packet to point to the next to next hop (thus skipping the next hop present in the original source route). Additionally, the *skip counter* associated with that next hop is incremented. The intuition behind this is that since the next hop is unreachable the next hop has probably moved away topologically. However, a timer is associated with each neighbor node and when the timer expires and the source route still has the skipped node as the next hop, the next hop is again tried at the MAC layer before the intermediate node decides to skip it. This ensures that transient failures do not isolate a node.

- An intermediate 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 using a link-layer acknowledgement (such as provided in IEEE 802.11 [2]). A passive acknowledgement does not work in this case since the data packet is sent directionally. Even after multiple MAC layer retries, if the intermediate node cannot verify that the packet reached the next hop as mentioned in the source route, the node instead of sending back a ROUTE ERROR (as in base DSR) tries to estimate the direction of the next-to-next hop in the source route. The idea here is to skip the next hop and reach the next-to-next hop directly. In order to do this the intermediate node estimates the direction of the next to next hop (from the estimated direction of the next hop and the estimated direction of the next to next hop from the next hop). The intermediate node also utilizes these estimates to estimate the beamwidth to be used. Later we explain how the intermediate node makes these estimates. Once the estimate direction and beamwidth is known the intermediate knows uses the ability of a directional antenna to transmit over longer distance to transmit the packet. The transmission power is set such that the packet can travel (in the direction of transmission) at least twice the omnidirectional transmission range (since if we assume that the next to next hop has not moved, then the next to next hop can be at most at a distance of twice the omnidirectional transmission range.

  Also, once the intermediate node decides to skip the next hop the intermediate node increments the *skip counter* associated with the next hop.

As explained above, the protocol will try to abstain from generating ROUTE ERRORs as long as possible by skipping broken links. However, this procedure introduces long hops into the source route and thus makes the route more fragile. Hence, we limit the number of long hops to be encountered by a packet. In other words only a limited number of hops are allowed to be skipped.

Estimation of Direction and Beamwidth to Reach Next to Next Hop*:* As shown in Figure 16, the estimated region in which the next to next hop C (from A) is expected to lie is given by the parallelogram. This is true because B could lie at any point on the line joining A and B. Similarly, C could lie anywhere on the line between B and C. The expected direction of C is given by the direction of the diagonal from A to C.

B. Evaluation

We do preliminary evaluation of our protocol in the scenario shown in Figure 17. In the scenario, A sends packets to E. After the initial Route Discovery has succeeded in finding a route to E and some packets have been delivered to E (so that the next to next hop directions can all be set up correctly), C goes dead after every 5 seconds and comes back up after another 5 seconds. This simulates some periodic problem with C , or C periodically moving away.

When base DSR is used for sending packets in such a scenario, a ROUTE ERROR is generated by B shortly after C dies. This is followed by a ROUTE REQUEST from A which does not succeed till C comes up again (the partition bridging modifications are not present in base DSR).

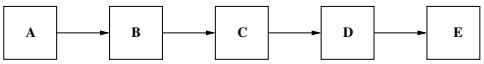

Figure 17. Scenario used for preliminary evaluation of route repair protocol

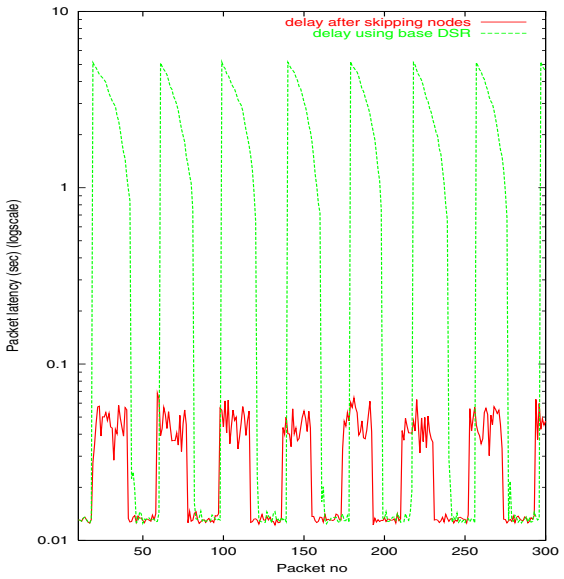

Figure 18. Latency incurred by packets before getting delivered

However, when we add our route repairing modifications to DSR and use a directional antenna to go directly from *B* to *D* thus skipping *C*, we are able to deliver packets without any ROUTE ERRORS (and hence no new ROUTE REQUESTS) being generated. This directly translates to the delays that packets incur before being delivered, as shown in Figure 18. The packets incur more delay when using base DSR because the packets spend a long time in the Send Buffer of the sender (*A*) waiting for the source to find out a source route. However, the modified scheme tries to skip of *C* when *C* starts dropping packets, and the little delay that is incurred is because of the retries that *B* has to do before it can decide to skip *C*. Once *C* has been skipped, the node decides to retry for *C* for some number of times and finally skips *C* for good. However, after a timeout *C* is again given a chance so that any transient problems with *C* does not isolate *C*.

VI. 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*. In contrast our work focuses on reducing the overhead of route maintenance. Moreover, 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. Wieselthier et al [17] 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.

Ramanathan and Hain [15] used antennas coupled with adjusting the transmission power to control the topology of multi hop wireless networks. Our work is complementary to this work. In this work the authors reduce the transmission power and still maintain connectivity. This increases throughput due to less interference. Our protocol starts when the limits of omnidirectional antennas (as far as transmit power is concerned) has been reached. Also, unlike their work, we do not require any link state algorithm. Instead our algorithm merges well with any on demand routing protocol.

Unlike in previous works, in this paper we target routing improvement in mobile ad hoc networks. Specifically, we achieve the following: bridge network partitions in the presence of permanent partitions *and* even in the absence of permanent partitions we improve route maintenance by decreasing ROUTE ERRORS (and hence ROUTE REQUESTS). For both of the above mentioned achievements we utilize the fact that a directional antenna is better able to channel its energy in the direction of transmission.

In our directional MAC protocol we send RTS and CTS packets omnidirectionally and send DATA and ACK packets directionally. Instead of using our simplistic directional MAC protocol we could have used one similar to the ones proposed by Ko et al [9]. However, we claim that our MAC protocol is a close approximation of the protocol suggested by Ko et al (without the optimizations) in the case when the nodes do not have any location information (through GPS, etc). For such a scenario, their protocol, just like ours, uses an omnidirectional RTS as well as an omnidirectional CTS. The only difference is that their protocol sends a directional RTS if the sending node has some estimate of the direction of the receiver. If however, the directional RTS fails then an omnidirectional RTS is sent. Our protocol always sends an omnidirectional RTS.

VII. Conclusion

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. In this paper, we have presented the design and evaluation of two techniques for routing improvement using directional antennas in mobile ad hoc networks.

First, we use directional antennas to bridge network partitions by adaptively transmitting selected packets over a longer distance, using the capabilities of the directional antenna, yet still transmitting most packets shorter distance in order to reduce power consumption and interference. Each node in the routing protocol learns when this capability of directional antennas may be necessary for any given routing or data packet; most packets are transmitted using the directional antenna over a normal, shorter distance, thus reducing power consumption and interference and generally increasing the capacity of the ad hoc network. Through simulations, we demonstrated the effectiveness of our design by modifying the Dynamic Source Routing (DSR) protocol, an on-demand ad hoc network routing protocol. Our simulations show that the modified protocol is able to effectively bridge network partitions, and that the protocol is otherwise equivalent to the original protocol when no partitions are present.

Second, we proposed a method to use directional antennas to repair routes in use when an intermediate node moves out of wireless transmission range along the route. By using the capability of a directional antenna to selectively transmit packets over a longer distance, we bridge the break in the route caused by the intermediate node's movement, thus reducing packet delivery latency and avoiding dropped packets and additional routing overhead. As with our partition bridging technique, the basic idea behind this technique is to use the capability of a directional antenna to transmit over longer distances, but to use this capability only when necessary for selected packets. We presented the results of simulations giving a preliminary performance evaluation of this technique demonstrating its effectiveness.

In future work, we plan to integrate the partition bridging protocol with the proposed method to repair routes in mobile ad hoc networks. Also, we plan to address the issues of energy consumed by a directional antenna as compared to an omnidirectional antenna.

References

[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, 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, December 1994.

[6] David B. Johnson and David A. Maltz. Dynamic Source Routing in Ad Hoc Wireless Networks. In *Mobile Computing*, edited by Tomasz Imielinski and Hank Korth, 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 *Ad Hoc Networking*, edited by Charles E. Perkins, 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] 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.