Ad Hoc Routing Using Directional Antennas

Technical Report
August 2002

Romit Roy Choudhury and Nitin H. Vaidya

Dept. of Electrical and Computer Engineering, and
Coordinated Science Laboratory
University of Illinois at Urbana Champaign
Email: croy, nhv@crhc.uiuc.edu

Abstract—**This report addresses the problem of using directional antennas in ad hoc routing. Although the problems with utilizing directional antennas have been visited in the past, the research has been confined mostly to issues related to medium access control. To determine whether directional antennas are beneficial to ad hoc networks, it is necessary to evaluate the impact of directional antennas on the performance of routing protocols as well. In this report, we evaluate the performance of a reactive routing protocol for different scenarios using directional antennas. Using insights gained from simulations, we propose routing strategies that adapt the routing protocol to directional communication. Our analysis shows that by using directional antennas, ad hoc networks may achieve better performance. However, scenarios exist in which using omnidirectional antennas may be more appropriate.**

I. Introduction

Mobile ad hoc networks are typically assumed to be equipped with omnidirectional antennas. However, it may be possible to improve the network capacity by using directional antennas effectively. Although recent research activity has addressed some of the problems related to directional medium access, the impact of directional antennas on the performance of higher layer protocols has not been adequately explored. As evident from subsequent sections of the report, the impact of directional communication on the performance of an ad hoc network is often counter-intuitive. While fewer hop routes may be discovered due to the higher transmission range of directional antennas, performing a simple neighborhood broadcast may now require the antenna system to sweep its transmitting beam sequentially over multiple directions. As a result, neighbors of a node receive broadcast packets at different points of time (unlike with omnidirectional antennas). Also, sweeping incurs greater delay and can incur higher control overhead, partially negating the advantages derived from reducing the hop-count of discovered routes. To study this trade-off, we evaluate the performance of ad hoc routing using directional antennas. Specifically, we evaluate the *Dynamic Source Routing* protocol (DSR) [4], and also suggest optimizations to adapt to directional antennas.

The rest of the report is organized as follows. Section II discusses related work on utilizing directional antennas in ad hoc networks. Section III provides some background information. In Section IV we describe a simple medium access control (MAC) protocol, adapted to directional antennas – named DiMAC – and describe how DSR operates on DiMAC. Section V discusses the performance tradeoffs and evaluates the performance of the routing protocol. The optimizations on the routing protocol are also discussed in this section. Section VI discusses our plans for future work. Section VII concludes the report with a brief summary.

Work reported here is supported in part by National Science Foundation (NSF) under grant ANI 01-96410

II. Related Work

There is a growing interest towards utilizing directional antennas in ad hoc networks. While using directional antennas has been extensively studied in the context of cellular networks, research in ad hoc networks has been confined mostly to medium access control [1][2][6][7][12][14][17]. Work on routing protocols using directional antennas is limited [1][8] [10]. In [8], Nasipuri *et al.* have utilized directional antennas for the purpose of on-demand routing. Their primary aim is to minimize routing overhead by intelligently using directional antenna elements for propagating routing information. Their results highlight the large routing overhead incurred by using omnidirectional antennas and show that it is possible to perform better using switched beam antenna systems. However, the authors have equalized the transmission range of directional and omnidirectional antennas, thereby underestimating the potential of directional beamforming. Put differently, the impact of discovering shorter routes using directional antennas has not been studied. In our work, we evaluate the effect of various transmission ranges of directional antennas. [8] also assumes a conservative MAC model. We believe that medium access control could be more aggresive with directional communication and accordingly use a suitable MAC protocol.

Ramanathan [10] presents an interesting discussion on several issues arising from directional communication. [10] discusses the possibilities for taking advantage of higher transmission range of beamforming antennas. However, the discussion emphasizes more on the MAC layer, and does not focus on issues related to the performance of the routing protocol. In this report, we focus primarily on routing using directional antennas. [1] proposes a proactive routing protocol (originally proposed for omnidirectional antennas [13]) over ESPAR antennas. The routing protocol performs well and gives reason to believe that directional antennas may be acceptable for ad hoc networking. However, the MAC protocol is complex, requiring several exchanges of control packets before actual data can be transmitted. Since reactive protocols have been shown to perform well in ad hoc networks, in this report we consider a reactive routing protocol. Also, our MAC protocol is simpler than [1].

In [12], we present directional MAC protocols, identify several problems with directional medium access and evaluate their impact on network performance. However, [12] does not address the impact on routing. Some of our initial work on using directional antennas in ad hoc routing has been reported in [11]. In this report, we present a more complete version of our research.

III. Preliminaries

A. IEEE 802.11 Distributed Coordinated Function (DCF)

In the IEEE 802.11 MAC protocol [3] for omnidirectional antennas, an exchange of RTS/CTS precedes DATA communication. Both RTS and CTS packets contain the proposed duration of data transmission. Nodes located in the vicinity of communicating nodes, that overhear either (or both) of these control packets, must themselves defer transmission for this proposed duration. This is called *Virtual Carrier Sensing* and is implemented by using a variable called the *Network Allocation Vector* (NAV). A node updates the value of the NAV with the duration field specified in the RTS or CTS. Thus the area covered by the transmission range of the sender and receiver is reserved for data transfer, to overcome the hidden terminal problem [5].

The IEEE 802.11 MAC protocol uses a backoff interval to resolve channel contention. Before initiating transmission, a node S chooses a random *backoff interval* from a range [0, CW], where CW is called the *Contention Window*. Node S then decrements the backoff counter once every "slot time". When the backoff counter reaches 0, node S transmits the packet. If the transmission from S collides with some other transmission (collision is detected by the absence of a CTS), S doubles its CW, chooses a new backoff interval and attempts retransmission. The *Contention Window* is doubled on each collision until it reaches a maximum threshold, called CW_{max}. While in the backoff stage, if a node senses the channel as busy, it freezes its backoff counter. When the channel is once again idle for a duration called DIFS (DCF interframe spacing), the node continues counting down from its previous (frozen) value.

802.11 invokes the backoff procedure only after a channel has been sensed to be idle for DIFS duration. A shorter interframe space, SIFS, is used to separate transmissions belonging to a single dialog (i.e., a node performs physical carrier-senses for SIFS duration before transmitting CTS, Data and ACK frames).

B. Antenna Model

We have incorporated the antenna model at the radio layer of the Qualnet simulator [9]. The antenna model is comprised of N antenna elements. Each element has a conical radiation pattern, spanning an angle of $\frac{2\pi}{N}$ radians. The elements are fixed with non-overlapping beam directions, so that the N elements may collectively cover the entire plane as shown in Figure 1.

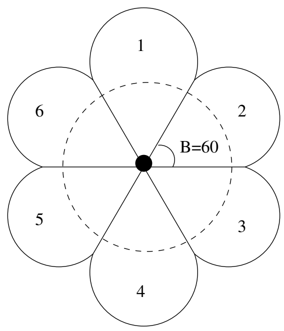

Fig. 1. The radiation pattern of a 6 element antenna system, with beamwidth B = 60 degrees. The dashed circle outlines the omnidirectional radiation pattern, indicating that the omni transmission-range is shorter than the directional range

The elements numbered 1 to N, maintain this configuration with respect to the earth's meridian even if the node changes its orientation. As an aside, this is achieved in modern antennas with the aid of a magnetic needle that remains colinear to the earth's magnetic field. This ensures that a particular element, is always pointing towards north for instance. In addition to the conical main lobe, a directional radiation pattern also includes side lobes, each having a low gain. The side lobes assume a complex pattern. In our implementation of the antenna model, motivated by the model in [10], we have approximated all the side lobes into a single sphere, with the node at its center. Thus, a directional radiation pattern assumes the shape of a bulb, similar to the shape in Figure 2 [10].

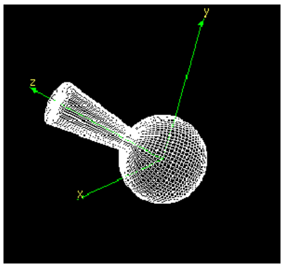

Fig. 2. A bulb shaped radiation pattern of a directional antenna, showing a spherical side lobe and a conical main lobe [8].

The antenna system offers two modes of operation: *Omni* and *Directional*. We assume that a node can operate in any one mode at a given time, but can toggle modes with negligible latency. In *Omni* mode, after a signal is detected, the antenna determines the element on which the received signal power is maximum. Rest of the packet is then received by using this element. We assume that in omni mode, signals are received with a gain G^o. An idle node stays in the *Omni* mode. In *Directional* mode, a node can select *only one* of its antenna elements and beamform using it, with a gain of G^d, $G^d >= G^o$. Since only a single antenna element may be used for directional transmission, an omnidirectional broadcast using this antenna may be achieved by sequentially using each of the elements to transmit the same packet. For example, if the node shown in Figure 1 wishes to broadcast a packet to all its surrounding neighbors, it must transmit the same packet 6 times, once with each antenna element.

The MAC layer we present in this report is aware of the antenna element that must be used to communicate to a given neighbor. Thus, when a packet needs to be unicast to a neighboring node, the MAC layer specifies the antenna element to the physical layer. Due to higher gain, nodes in directional mode can transmit/receive over a longer distance, in comparison to the omni mode, at a given transmit power level. We discuss the implications of higher transmission range in Section IV.

C. Dynamic Source Routing (DSR) Protocol

We briefly describe the operation of a reactive routing protocol called Dynamic Source Routing (DSR)[4]. We later evaluate DSR over directional antennas, and also suggest some modifications to DSR to adapt to directional antennas.

The key features of DSR are *reactive route discovery* and *source routing* [4]. When a node in an ad hoc network attempts to send data packets to a destination for which it does not already know the route, it uses a route discovery process to dynamically determine such a route. Route discovery works by flooding the network with *route request* (RREQ) packets. Each node receiving a RREQ rebroadcasts it, unless it is the destination node, or the RREQ received is a duplicate of an already received RREQ (duplicate RREQs are detected using sequence numbers). As the route request is propagated to various nodes, the path followed by the route request is included in the RREQ packet. On receiving the route request, the destination node responds by sending a *route reply* (RREP) message to the sender – the route reply follows a path that is obtained by reversing the route followed by the route request received by the destination. The route carried back by the RREP is cached at the source node for subsequent communication. If any link on a source route is broken, the source node is notified using a *route error* (RERR) packet. The source removes any route that includes this link from its cache. DSR also incorporates strategies which allow the nodes to learn (and cache) routes by overhearing RREPs or data packets containing source routes. Routes from caches may be used to avoid route discoveries, thereby reducing the overhead of the flooding process.

IV. MAC and Routing: Protocol Description

In this section, we describe the MAC and routing protocols used in our evaluation. The MAC protocol, named DiMAC, is similar in many respects to the protocols described in [12][16], although not identical. For routing, we use original DSR for initial evaluation, but later optimize it to suit directional antennas.

A. DiMAC Protocol

We present DiMAC, a directional MAC protocol, for the purpose of evaluating routing over directional antennas. The design of DiMAC is based on the notion of reserving the wireless channel before actual data is transmitted.

Channel reservation is performed using a RTS/CTS handshake (between the sender S and receiver R), both being transmitted directionally. More specifically, DiMAC specifies the antenna element to be used for RTS transmission. To do this, DiMAC maintains a look-up table containing two fields; *NeighborId, AntennaElement*. The look-up table is a cache that is updated reactively to minimize overhead, as discussed in the next subsections. Using the specified antenna element, the antenna system beamforms in the direction of the intended destination and senses the channel. Note that even though S was listening to the channel in the omni mode while idle, carrier sensing has to be performed again, after it has beamformed in a particular direction. This is necessary because the directional beamforms have a higher gain. If S senses the carrier busy by using its directional beam, it postpones RTS transmission until the medium is once again idle. If the medium is idle, S proceeds through the steps of waiting for a DIFS period and backing off for a random interval (similar to the steps in IEEE 802.11) before it transmits the RTS.

Node R, while idle, listens to the carrier in omni mode and receives the RTS meant for it. The RTS is received with omnidirectional gain G^o. DiMAC at node R now uses the same antenna element with which it received the RTS, to send back a CTS. The CTS is, however, transmitted in the directional mode of operation. Sending back the CTS is also preceded by carrier sensing and waiting for SIFS duration, as in 802.11. S in the meantime remains beamformed towards R and receives the CTS directionally, with directional gain G^d. Once the RTS/CTS handshake is accomplished, node S sends the data packet directionally to R. R receives the data packet using the same element that it used for sending the CTS.

Nodes in the neighborhood of nodes S and R, that overhear RTS, CTS or both, defer transmission for the proposed duration of transfer. However, these nodes defer transmission only for the antenna element with which RTS or CTS is overheard. DiMAC maintains a directional NAV (DiNAV) table (similar to [12][16]) comprising of the fields *AntennaElement* and *WhetherBlocked*. On overhearing a RTS or CTS on *AntennaElement i*, a node sets the *WhetherBlocked* flag for element i, for the proposed duration. Subsequently, this node defers transmissions which requires use of element *i*. After the proposed duration is over, the *WhetherBlocked* flag is reset. For example, when node A communicates with node C in Figure 4, node B's antenna elements 2 and 4 are blocked (in Figure 4, note that node B overhears A's RTS on antenna element 4, and C's CTS on antenna element 2). However, node B may initiate communication (or reply to a communication request) using the other antenna elements.

A conspicuous problem with a directional MAC protocol such as DiMAC is *"deafness"*, identified and discussed in [12]. Although *deafness* is a MAC layer issue, it affects the performance of routing protocols. Therefore, we briefly outline the problem of *deafness* here. A node that uses directional antenna beamforms towards a node with which it communicates. As a result, this node cannot receive signals originating from other nodes, that lie outside the directional pattern of the antenna element it is currently using. Consider Figure 3. In Figure 3, if node A is communicating with node B and node C transmits a packet intended for A, A fails to receive the packet. When using omnidirectional antennas, this problem may not arise, because node C would be aware of A's communication to B, by receiving a RTS or CTS from A, or by physically carrier-sensing A's transmission. As seen later, the problem of *deafness* affects routing performance.

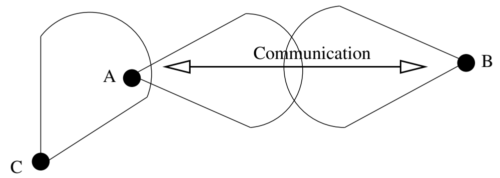

Fig. 3. An example scenario showing the possibility of deafness

B. DSR over DiMAC

We begin our evaluation of ad hoc routing using DSR and later suggest some improvements. A node attempting to send a data packet to another node initiates the route discovery process by broadcasting a RREQ to all its neighbors. When DSR is executed over 802.11 using omnidirectional antennas, a single RREQ transmission is received by all the neighbors of the transmitting node. Using DiMAC, broadcasting a packet to all neighbors of a node is achieved through *"sweeping"*. In other words, when a RREQ packet is handed down to DiMAC by DSR, DiMAC sequentially transmits the same packet over all the N antenna elements. A *sweep* emulates omnidirectionality by reaching neighbors located in all directions around a node. Observe that such sweeping incurs additional delay, almost N times greater (N being the number of antenna elements), in comparison to an omnidirectional broadcast.

While performing a *sweep*, packet transmissions are not preceded by *backing off*. If the antenna element being used for transmission senses the channel idle, it transmits

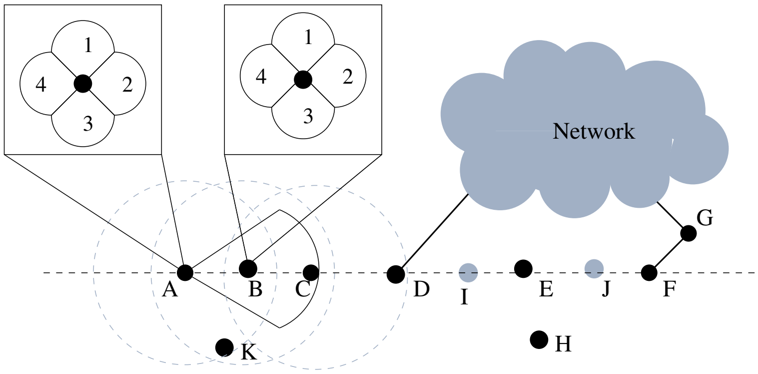

Fig. 4. An example scenario. A 4 element antenna system is shown in the inset. Node A uses antenna element 2 to communicate with node C. Node C uses element 4 to receive signals from mode A. The dashed circles outline the omnidirectional transmission range

the RREQ. Antenna elements that sense the channel busy are marked during the first round of transmission. Once the round is completed, DiMAC attempts to transmit the RREQ only over the marked antenna elements (sequentially). If some of the marked antenna elements now sense the channel idle, the RREQ is transmitted using them. For antenna elements that still sense the channel busy, DiMAC drops the RREQ. This entire procedure constitutes a single *sweep*. Note that the probability of RREQ collision does not increase much due to the elimination of *backing off*. This is because the receiver always receives with the antenna that experiences highest signal strength. Thus as long two different signals arrive on different antenna elements, the receiver would always receive at least one of them. [1]

Although the sweeping delay is high, by using directional transmission a RREQ can be delivered to a far away node over a single hop. For example, in Figure 4, if an RREQ is transmitted directionally by A, link A-C can be on the route. If the RREQ is transmitted omnidirectionally by node A in Figure 4, then node C would not receive A's transmission, and link A-C cannot be on the chosen route.

To unicast packets to a particular neighbor, a node must use the appropriate antenna element. This information is cached at each node in a look-up table. A node that receives a packet updates its look-up table with the *sender's identifier* and the *antenna element* that was used to receive the packet. This information is used in the near future for unicast transmission. For example, if node A receives a packet from node B using antenna element i, then to initiate communication to node B in the near future, node A uses the same element i. In mobile scenarios, it is possible that information cached in the look-up table gets stale. This could happen if neighbors of a node move to a location where using the antenna element specified by the look-up table is no more appropriate. To communicate to the relocated node, a different antenna element may now be appropriate. Thus to keep the look-up table updated, we incorporate a *scanning* mechanism within DiMAC. A *scan* is essentially "hello" packets, transmitted sequentially over all antenna elements (i.e., swept), whenever necessary. If node B receives a "hello" from node A, B replies using the same antenna element with which it received the "hello". The reply from B is delayed to allow node A to first finish transmitting all the hellos (otherwise, A would not receive B's reply since A would be busy transmitting hellos in other directions). Node A then receives the reply from B, say, using the antenna element j. Then DiMAC at node A updates the look-up table entry for node B to record antenna index j. DiMAC at A initiates future communication with B using antenna element j. In summary, *neighbor discovery* becomes more complex using directional antennas since neighbors are now associated to specific antenna elements at a given instance of time. *Scanning* is one way of reactively performing directional neighbor discovery. (In the next section, we suggest a partial scanning optimization to reduce the overhead of scanning.)

V. Performance Evaluation

We use the Qualnet simulator [9] version 2.6.1, for simulating our protocols. We have modeled the antenna system in the physical layer of Qualnet. The omnidirectional and directional gains (G_o and G_d respectively) have been

[1] Alternatively, we could introduce *backing off* before each transmission. Although it may reduce the probability of RREQ collisions, it would increase the sweeping delay at each node. Presently, we are looking into this possibility as well.

calculated using expressions from [10]. Using these antenna gains, the relation between the transmission range and the beamwidth of the beam's main lobe is shown in Figure 5. The transmission range indicates the maximum distance between a transmitter and a receiver at which communication is possible using DiMAC. Put differently, if a node (in directional mode) wishes to transmit an RREQ to an idle receiver (in the omni mode), it must be within *transmission range* distance from the receiver. The curves are plotted for different levels of transmission power. For example, in Figure 5, at power = 10 dBm, the transmission range for beamwidth = 360 degrees (omnidirectional antennas) is 250 meters. At the same power level, the transmission range for beamwidth = 60 degrees (corresponding to N = 6) is around 450 meters. Power levels of directional and omnidirectional transmissions are equal in all our simulations, both being at 10 dBm. We simulate our scenarios in a bounded region of 1500 X 1500 meters, with *Constant Bit Rate* (CBR) traffic for data communication. The topology, flow pattern and traffic distribution is described separately for each individual scenario. Initially we consider static networks for our simulations. Later we evaluate the impact of directional routing in mobile scenarios.

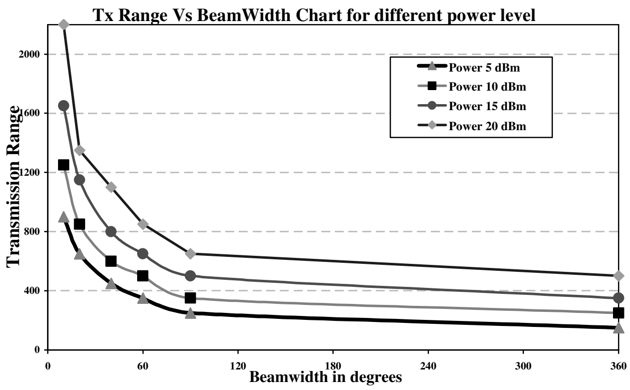

Fig. 5. The transmission ranges for increasing gain, for different power levels

The tradeoffs we discuss in this section arise from the specific characteristics of our protocol pair - Directional DSR (DDSR) and DiMAC. However, we believe that our broad observations will apply to other protocols as well.

An interesting tradeoff arises from the counteracting effects of large sweeping delay and higher transmission range (thus fewer hops on routes). On one hand, RREQ propagation gets delayed due to sweeping, while on the other hand, shorter routes could be discovered due to a larger transmission range of directional antennas. Figure 4 may be used to illustrate this. Observe that a RREQ from A can reach D in two hops via node C, if directional antennas are used. In contrast, by using omnidirectional antennas the number of intermediate hops will be greater (through nodes B and C).

In addition to this, spatial reuse affects delay in performing route discovery. Referring to Figure 4, note that when node K forwards the RREQ using omnidirectional antennas, node B must defer transmission of RREQ. This is not necessary for directional antennas as nodes B and J can transmit simultaneously in different directions. Thus the drawback of large sweeping delay may be compensated by the ability to perform simultaneous transmissions at nearby nodes. To evaluate these tradeoffs, we simulate several static and mobile scenarios, with different flow patterns. We discuss the insights gained from simulation results and propose optimizations that improve the performance of the DDSR protocol. We begin by simulating a grid topology with a single flow in the network.

The simulated topology is a rectangular grid as shown in Figure 6. Adjacent nodes in the grid are placed 200 meters apart. To evaluate network behavior, we measured *route discovery latency* (RDL) and *throughput*. Route discovery latency is the time duration calculated from the point a RREQ is transmitted by the sender till the point a RREP is received by the sender. We measure *route discovery latency* and *throughput* for DSR and DDSR.

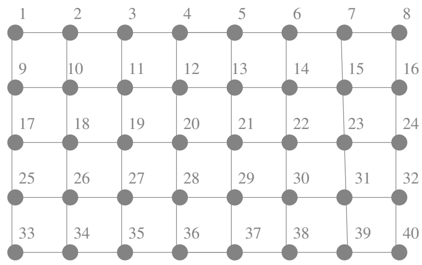

Fig. 6. The rectangular grid used for simulations

A. Evaluating Route Discovery Latency

Figure 7 shows the variation of *route discovery latency* (RDL) versus the physical distance separating source and destination nodes. For this simulation, we have used node 17 as the source node (Figure 6). The destination node is chosen according to the desired distance of separation. For example, to set up a distance-of-separation of 600 meters, node 20 is chosen as the destination. Figure 7 plots RDL for different values of N, N being the number of antenna elements. For example, DDSR6 refers to Directional

DSR simulated over DiMAC with 6 antenna elements. The beamwidth with N antenna elements is $(2\pi/N)$. Hereafter, we would denote Directional DSR simulated over DiMAC with N antenna elements, as "DDSR N". By "DSR", we would indicate DSR executed over omnidirectional antennas, using IEEE 802.11 as the MAC protocol (this is equivalent to DSR1).

From Figure 7, we observe that the directional and omnidirectional curves cross over frequently when the distance of separation between the source and destination is small. At larger separation, directional antennas (except when N = 4) exhibit lower route discovery latency in comparison to omnidirectional antennas. We now analyse the issues and discuss the insights gained from the results.

Behavior of *route discovery latency* (RDL) using directional antennas may be intuitively explained as follows. Several factors influence RDL: *higher transmission range* and *spatial reuse* reduce RDL, while *sweeping delay*, *deafness* and *higher directional interference* increase RDL.[2] When the separation between the source and destination node is small, the gain due to higher transmission range is only marginal. This is because both directional and omnidirectional antennas are capable of reaching the destination over a few hops. As an example, consider Figure 8. Reaching node D from node S in Figure 8 requires three hops for both DDSR4 and DSR (DDSR6 or DDSR9 would require at least two hops in this example). However, in such scenarios, the *sweeping* delay of DDSR is high, easily offsetting the slight (if any) gain of a shorter directional route. On the whole, route discovery latency for DDSR is higher than DSR, when separation between the source and destination node is small.

When source-destination are separated by larger distances, the advantage of higher transmission range begins to dominate. The gains due to a shorter route in DDSR, now offsets the additional delay incurred in sweeping. However, we observe that the route discovery latency for DDSR4 (meaning DDSR with 4 antenna elements, i.e., beamwidth = 90 degrees) is almost identical to DSR, even at higher source-destination separation. This happens because the transmission range of DDSR4 is 350 meters and the grid distance is 200 meters. Observe that even with a transmission range of 350 meters, DDSR4 can at best communicate directly to an adjacent node in the grid (Figure 8). DSR on the other hand, with a transmission range of 250 meters, will also be capable of communicating directly to adjacent nodes. Thus DDSR does not show any benefit over DSR. In fact, due to sweeping delay, the performance of DDSR4 is somewhat worse than DSR, although not much worse because of the advantages of spatial reuse. We now evaluate the impact of reducing grid-distance (distance between adjacent nodes in a grid) on the performance of DSR and DDSR.

[2] *Deafness and interference may prevent several nodes from receiving a broadcast RREQ. This may cause the RREQ to reach the destination node through a suboptimal path, increasing route discovery latency as a result.*

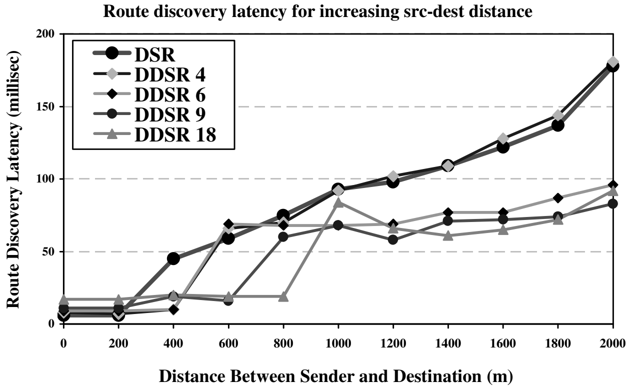

Fig. 7. Route Discovery latency for varying number of antenna elements at 200m grid-distance

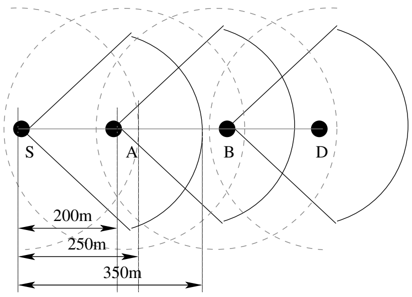

Fig. 8. A section of a grid, illustrating the possibility of DSR and DDSR4 discovering equal-hop routes

It is relevant to point out that the behavior of route discovery latency depends on node density as well. Consider the scenario discussed above with DDSR4, DSR and a grid distance of 200 meters. For a source destination pair separated by, say, 800 meters, both DSR and DDSR4 discover 4 hop routes, leaving DDSR4 with no advantage for higher transmission range. However, if the grid distance was 50 meters, it could be possible for DDSR4 to discover a 3 hop route, while DSR must still find a 4 hop route. This motivated us to evaluate the performance of DDSR in a high density network with grid distance 50 meters. Figure 9 shows the results of simulation. Counter to intuition, the performance of DDSR degrades in denser networks. As evident from Figure 9, route discovery latency using DSR is comparable to DDSR even at higher transmission ranges and large source-destination separation.

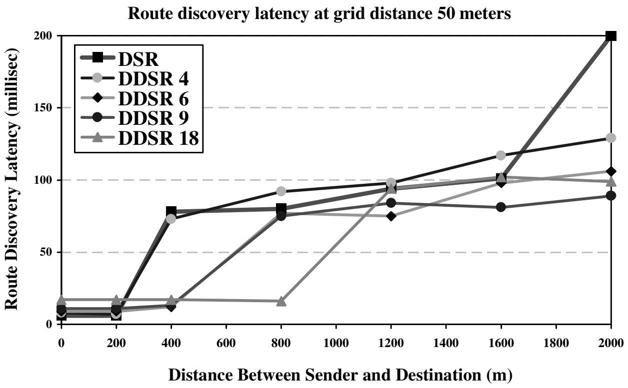

Fig. 9. Route Discovery latency for varying number of antenna elements at 50m node-spacing

Careful study reveals that in presence of high node density, the negative impacts of interference is significant. Due to closely packed nodes, signals received from unwanted directions using antenna side-lobes, increase the probability of collisions. As a result, RREQs take longer to reach the destination, resulting in higher route discovery latency. Thus expected gains of higher transmission range in dense networks gets offset by these counteracting factors. The net effect is that DDSR performs only marginally better than DSR. The advantage of high transmission range becomes conspicuous only when the source-destination separation is very large.

B. *Evaluating Throughput*

We now focus on throughput when using DDSR. Figure 10 shows the throughput of a single CBR flow for different source-destination separations. The traffic generated is large enough to keep the source backlogged at all times. The grid distance for this graph is 200 meters. We observe that DSR performance is comparable to DDSR4's performance. For smaller beamwidths (i.e., higher transmission range), the throughput is not much greater, although the expected number of hops using directional antennas could be far lesser than omnidirectional antennas. The reason for poor performance of DDSR may be attributed to two reasons, discussed below.

1. In the simulation of a single flow in the grid topology, we observed that DDSR often chooses sub-optimal routes. This happens because *sweeping* causes neighbors of a node to get the same RREQ at different points of time. Consequently, neighbors that receive the RREQs earlier have a better chance of delivering the RREQs earlier to the destination. Thus replying to the first arriving RREQ may not be suitable when sweeping is performed. Choosing sub-optimal routes may affect throughput significantly. Interestingly, replying to all the arriving RREQs (as optionally performed by DSR) is also not suitable when sweeping is performed. In the next subsection, we discuss the reasons for this and evaluate an optimization *"Delayed Route Reply"*, to handle this problem.

2. In addition to sub-optimal routing, MAC layer studies indicate that directional antennas perform poorly in "linear" topologies [12][10]. This is because higher directional interference in a particular direction reduces spatial reuse if many nodes are placed in the same direction. A grid topology is such a linear topology. Performance in random scenarios may be significantly different. We consider random topologies later in this subsection.

Delayed Route Reply Optimization: We now evaluate our first optimization to DSR. The optimization is motivated by the fact that the earliest RREQ received by the destination may not have traversed the shortest or the least congested path. This optimization requires the destination node to delay sending the RREP by a time duration T calculated from the time it received the first RREQ. This allows the destination node to choose the best among all the routes that arrive within this time T. We specify T as $T = \rho \times T_{sweep}$, where T_{sweep} is the time taken to complete one full sweep, and ρ is a configuration parameter. We emphasize that replying to all the RREQs that arrive at the destination is not equivalent to the optimization we propose. This is because, if the destination replies to the first RREQ, it would need to beamform towards one of its neighbors, to whom it must unicast the RREP. Now, while it is beamformed towards one of its neighbors, other RREQs may be transmitted by its neighbors. Due to the effect of deafness, the destination node may fail to receive these RREQs of which one could possibly have traversed the shortest path. Waiting for a sufficiently long time duration T minimizes this possibility. We have observed that with this simple optimization, routes chosen by DDSR are often shorter than what was chosen without it. Hereafter, all simulations incorporate the *delayed route reply* optimization with $\rho = 1$. Although hop-count reduces, route discovery latency of DDSR increases by T time units, which may be tolerable if route discovery is not performed too frequently.

The throughput when using directional antennas is better than omnidirectional antennas for a single flow. How-

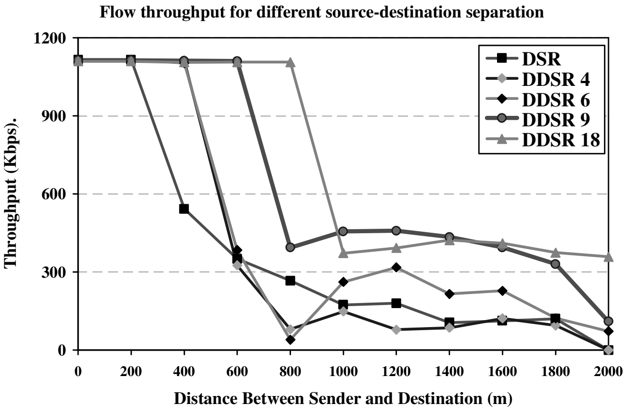

Fig. 10. Throughput versus distance-of-separation between source and destination, for varying number of antenna elements, at 200m grid-distance

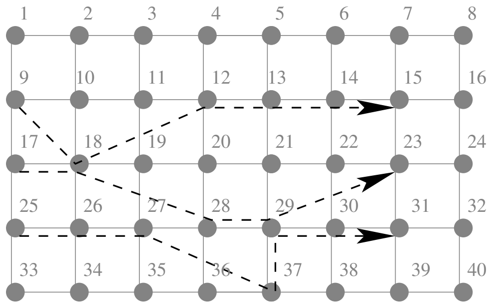

Fig. 11. A rectangular grid topology with 3 flows (node 9 to 15, 17 to 23 and 25 to 31)

ever, in presence of multiple flows, one may expect DDSR to further outperform DSR, due to spatial reuse of the channel. We simulated a rectangular grid topology with 3 flows as shown in Figure 11. The dashed lines show an example of routes discovered by DDSR. Adjacent nodes in the grid are separated by 200 meters. Figure 12 shows that for multiple flows the performance of directional antennas almost falls below the omnidirectional performance (note that the omnidirectional performance corresponds to the beamwidth of 360 degrees). This happens because in the presence of multiple flows, several MAC layer issues start affecting performance – specifically, the problems of topological linearities and deafness become dominating. We now discuss the MAC issues in brief to gain better insight into the performance degradation of DDSR. Routes that share common intermediate nodes suffer from unnecessary backing off due to deafness. To illustrate this we refer to the scenario in Figure 11. Assume that both nodes 9 and 17 discover routes to their respective destinations, through a common intermediate node, node 18. When 18 forwards 17's packets to 28, 9 transmits RTSs to 18 but receives no CTS in reply, since 18 is deaf to 9. Assuming congestion, 9 keeps backing off for increasingly greater durations and retransmits RTSs. This results in wastage of channel capacity affecting aggregate throughput of the network. Also, since routes exhibit linearity in the arrangement of its intermediate nodes, directional interference affects spatial reuse in such grid topologies. On the whole, the performance of DDSR is poor in terms of throughput.

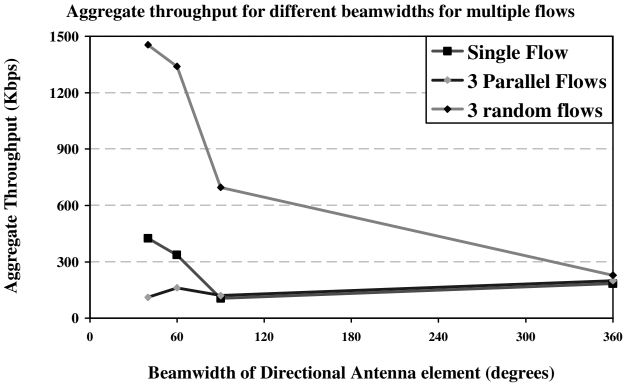

Fig. 12. Performance of single and multiple flows in linear and random topologies

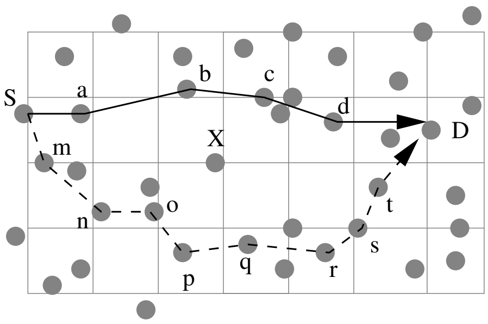

Fig. 13. A random topology showing the possibility of better connectivity due to higher transmission

To confirm that topological linearities and parallel flow patterns are the cause of degradation in performance, we displaced each node in the grid by a distance randomly selected from the interval [0, 200] meters. We simulated 5 random topologies, with 3 flows in each (the source destination nodes are identical to the three parallel flows used for the grid topology). Figure 13 shows one such random topology and indicates the routes taken by DDSR and DSR for one of the three flows. Figure 12 shows the average throughput of DSR and DDSR, simulated in random topologies (recall that DSR corresponds to a beamwidth of 360 degrees on the horizontal axis). Clearly, DDSR performs better than DSR. Closer examination reveals that, in several instances, higher transmission range of directional antennas is effective in bridging "voids" in the topology. Put differently, where omnidirectional antennas require long routes, directional antennas can form shorter routes. To illustrate this with an example, assume the scenario in Figure 13. The solid line shows the DDSR route, while the dashed line indicates the DSR route. Clearly, the DSR route is much longer than DDSR. This happens because the DSR transmission range is not sufficiently large to form a link between nodes a and b. This forces DSR to choose a longer route through nodes m, n, o, etc. On the other hand, higher transmission range of DDSR allows a link between nodes a and b, bridging, what may be called a "void region" for omnidirectional transmissions.

In random topologies not characterized by such "void regions", DDSR could still potentially find shorter routes in comparison to DSR. In some instances, the hop count of DDSR may equal DSR but rarely would DDSR discover a longer route. The advantage of smaller hop-count (due to higher transmission range), combined with spatial reuse of directional communication, can lead to higher aggregate throughput with DDSR as compared to DSR. This is evident in Figure 14. The graph shows the average of aggregate throughput over 25 random topologies for increasing beamwidth (thus decreasing transmission range). While achieved aggregate throughput clearly decreases with decrease in transmission range, at very high transmission ranges the trend is reversed. This may be explained as follows. To achieve higher transmission range (narrower beamwidth), the number of antenna elements required increases proportionally. This not only increases *sweeping delay* but worsens the effects of *deafness* as well. In other words, nodes spend more time in the directional mode while sweeping, increasing the possibilities of being deaf to arriving RREQs. As a result nodes that are capable of forming the shortest route between the source and destination often do not receive the RREQ. Sub-optimal routes are discovered in this scenario (note that sub-optimality due to such deafness is not handled by the *Delayed Route Reply* optimization). As a consequence the performance can degrade when using narrow beamwidths.

C. Routing Overhead

The overhead associated with sweeping degrades performance of DDSR, especially when the network is dense or route failures are frequent. Sweeping requires N sequential transmission of the same packet in N different directions. While a single directional transmission may interfere only in a particular direction, multiple such transmissions make up for a large overhead in comparison to omnidirectional antennas. We compare control overhead (of DSR and DDSR) using the following overhead metric α.

$$\alpha = \frac{\Sigma Number of Control Packets \times Area Blocked by Each Packet}{\Sigma No. of Data Packets}$$

where *Area Blocked by Each Packet* is the area covered by the radiation pattern of a single antenna element. Intuitively, network capacity consumed due to the transmission of each control packet is proportional to the region silenced by that transmission. The total control overhead is thus equivalent to *Number of Control Packets* $\times$ *Area Blocked by each packet*. This overhead is in exchange of the total *Number of Data Packets* delivered to destination nodes in the network. We use the ratio of these two quantities as the metric α.

The routing overhead of directional antennas could be much greater than omnidirectional antennas due to higher transmission range of directional antennas. To reduce this overhead in DDSR, we propose the following optimization.

Selective Forwarding Optimization: In Figure 15, when node Y receives a RREQ from node X, it may not be necessary to transmit that packet back in the direction of node X. This is because, nodes lying in the coverage of node X's antenna element 3, can receive the packet from X and need not receive a copy of the same packet from Y again. In omnidirectional scenarios, the problem of redundant reception of RREQs has been referred to as the *broadcast storm* problem [15], since control packets are unnecessarily broadcast in the channel. Using directional antennas, we consider an optimization whereby a node forwards a control packet with only n $(n <= N)$ antenna elements.

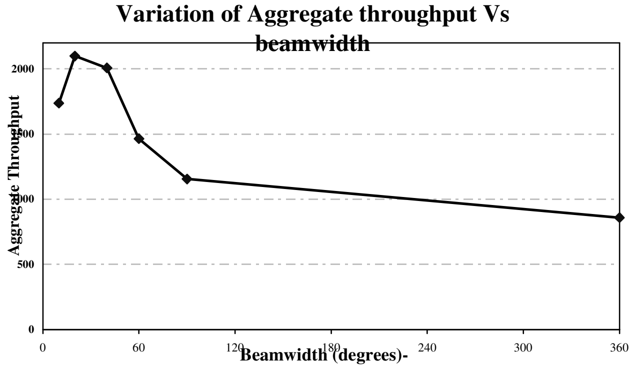

Fig. 14. Variation of performance for different beamwidths in a static network

A node initiating the control packet, however, forwards it using all elements.

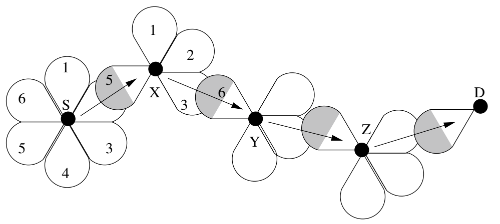

Fig. 15. An example scenario showing the propagation of RREQ using the *selective forwarding* optimization. The beams with which nodes receive RREQs are shaded.

We define $\gamma = \lceil \frac{n}{N} \rceil$. Parameter γ may be treated as a configuration parameter. We have used $\gamma = \frac{1}{2}$ for our simulations. Also, in our simulations, the *n* antenna elements used to forward control packets are the ones that are diagonally opposite to the element with which the control packet was received. More specifically, if $\gamma = \frac{1}{2}$, N = 6 (i.e., 6 antenna elements numbered from 1 to 6), and if a node receives a RREQ using antenna element 5, then antenna elements 1, 2 and 3 are used to forward the RREQ. Figure 15 shows how a RREQ may propagate using $\gamma = \frac{1}{2}$ and N = 6. Observe that node X receives the RREQ using antenna element 5 and forwards it using elements 1, 2 and 3.

Figure 16 plots routing overhead α against beamwidth of directional antennas. α was calculated from 25 simulations of random, static topologies. Clearly, DDSR routing overhead reduces due to this optimization, although it remains greater than DSR (recall that DSR corresponds to a beamwidth of 360 degrees). Thus, by selectively using directional antennas, it is possible to avoid control packet transmissions in unwanted directions, thereby enhancing the performance of the wireless system.

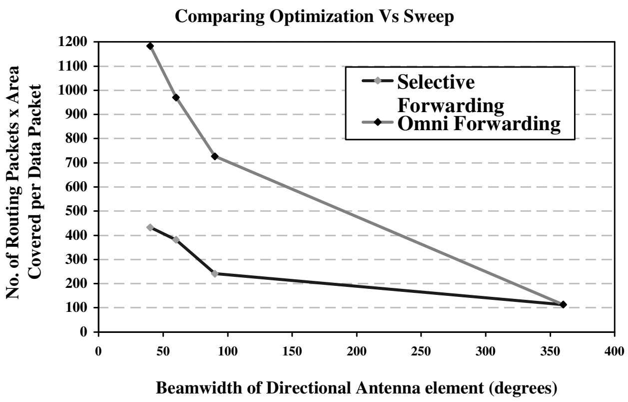

Fig. 16. Routing overhead α for optimized and unoptimized DDSR

D. Mobility and Route Error

The observations until this point may be summarized as follows. Directional antennas perform better than omnidirectional antennas when the network is random and sparse, and when the directional transmission range is substantially greater than the omnidirectional range. However, these conclusions are all under the absence of node mobility. Also, of interest is the behavior of DDSR under frequent route failures, since frequent route discovery may increase control overhead. In this section we discuss the impact of mobility on the performance of directional routing.

In the presence of node mobility, link failures may occur frequently due to disconnection of adjacent nodes in a route. When using directional antennas, a link failure is also possible when a node moves out of the radiation pattern of one antenna element to the radiation pattern of another antenna element. We call this the problem of *handoff*. Look-up tables thus may get stale due to node mobility and may require updating. In contrast, an advantage of using directional antennas in presence of mobility is greater link lifetime. In other words, due to higher transmission range, links comprising a route may persist for a longer duration when the end nodes of a link are moving away from each other. This reduces the frequency of route errors.

The problem of *handoff*, as mentioned earlier, arises when a receiver-node moves out of the transmitting beamform of the transmitter. As an example, a node S transmitting to node R using antenna element i, may require to transmit with antenna element j if node R moves into the radiation pattern of antenna element j. In our implementation, we have handled this problem of handoff at the MAC layer and DDSR remain unaware of it. DiMAC handles handoffs as follows. If DiMAC does not receive a reply from a neighboring node after a small number of attempts, it assumes handoff has occurred, and attempts to find the appropriate antenna element that could re-establish communication. To find such an antenna element, the node performs a partial *scan*, whereby the transmitter node sends "hello" packets using K antenna elements that are adjacent to the antenna element that was in use previously. As an example, if K = 2 and antenna element i was used previously, then "hello" packets are transmitted over elements $i-1$ and $i+1$, sequentially. Nodes that have moved to the coverage area of an adjacent antenna element ($i-1$ or $i+1$ in this example), can now re-establish communication, by replying to the "hello" packet. The look-up table at the transmitter is correspondingly updated with this new antenna element. If the link cannot be established using any of the K elements, DiMAC informs DDSR of a link failure and a route error packet (RERR) is generated. Thus, handoffs do not always trigger RERRs. We have assigned K = 2 in our simulations, implying that 2 antenna elements are used for *scanning*. A higher value of K may not be very useful since it is unlikely that nodes would move across the coverage area of multiple antenna elements within the time frame of a few packet transmissions.

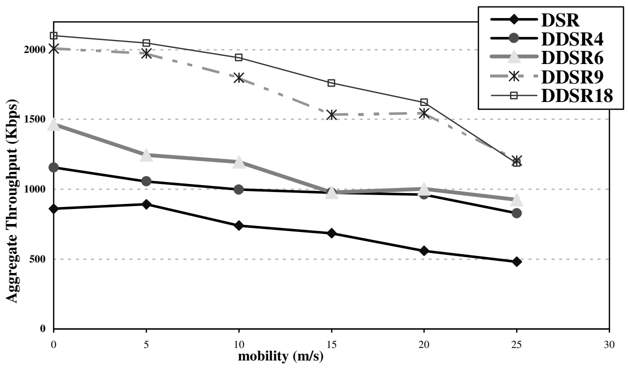

Fig. 17. Aggregate throughput measured for mobile scenarios, for varying transmission ranges

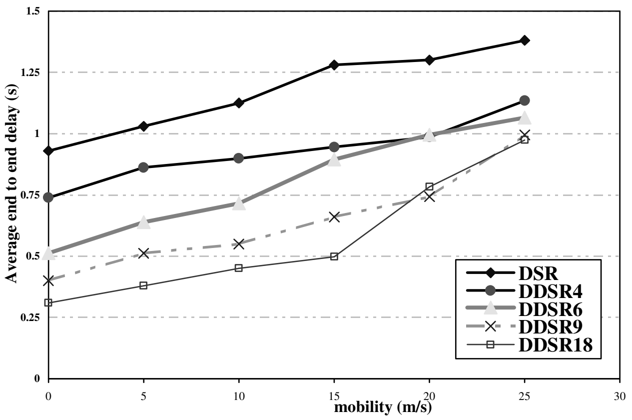

Fig. 18. Average end to end delay for mobile scenarios, for varying transmission ranges

Figure 17 shows a graph that measures aggregate throughput for several mobile scenarios and different beamwidths. A *random waypoint* mobility model [9][4] is assumed, with speed varying between 5 to 25 m/s. Caching strategies for DSR and DDSR was turned off for the simulations. For the purpose of uniformity, the time of simulation is proportionally reduced when mobility is increased. For example, the time of simulation for $mobility = 20m/s$ is half the time of simulation for $mobility = 10m/s$. As evident from Figure 17, aggregate throughput of DDSR degrades with mobility. This happens because the time for route discovery is almost a constant. To illustrate this consider an example where a simulation is performed for 100 seconds using $mobility = 10m/s$. Also assume that over the entire simulation, exactly one route error occurs at time = 20 seconds. Assume that a fresh route is discovered and data transmission resumes at time = 25 seconds. Now consider the scenario using $mobility = 20m/s$ and corresponding simulation time = 50 seconds. Observe that the route error in this scenario occurs at 10 seconds through the simulation and ends at 15 seconds. Consequently, the throughput at higher mobility is expected to decrease since a smaller fraction of the time (45/50) is now utilized for useful data communication (observe that at $mobility = 10m/s$, the fraction of time used for transmitting useful data is 95/100).

However, the difference in aggregate throughput achieved by DDSR and DSR is large, even at higher mobilities. This is an outcome of fewer route errors when using DDSR. Put differently, due to the higher transmission range of DDSR, links break less frequently. This reduces number of route discoveries and in turn the control overhead in the network. As a result a longer fraction of the simulation time is used for data communication when using DDSR. At very high transmission ranges however, the negative impact of each route discovery becomes excessive since almost all nodes in the network is flooded with control packets. This causes a sharp degradation in performance, partially offsetting the advantages due to relatively fewer route errors.

Figure 18 shows the measurement of average end-to-end delay for data packets in mobile scenarios. As seen in Figure 18, the average packet delivery delay is smaller with DDSR in comparison to DSR. However, between the DDSR curves, the performance is comparable. This happens because the advantage of higher transmission ranges gets compensated for, by larger sweeping delays. When route errors occur more often (at higher mobilities), frequent sweeping further reduces the differences in delay between the DDSR variations. On the whole, although DDSR incurs lower end-to-end delay in comparison to DSR, the improvement is not much especially at higher mobilities.

VI. Future Work

Optimization: Directional Route Repair

Using directional antennas, a node can have an idea about the direction of the route currently being used. If a route fails, restricting route recovery over a region enveloping the previous route direction could minimize control overhead. This may be done by selectively using antenna elements to propagate route recovery packets. Cases are possible in which such a restricted search could fail, in which case a network wide route discovery has to be performed subsequently. We plan to evaluate this optimization. Using this optimization, a node forwards route recovery control packets using K of its N antenna elements. We propose to choose the value of K as inversely proportional to the distance of the forwarding node from the destination.

Power Control
We have not discussed the effects of power control on the performance of directional antennas in ad hoc networks. In our future research, we plan to address power control with the objective of improving capacity as well as power conservation.

RTS/CTS and DATA/ACK at different powers:. A simple means of conserving power may be achieved by sending the RTS and CTS at the maximum power and exchanging DATA and ACKs at a lower power level. This may be achieved by exchanging information about the received power levels of the RTS and the CTS packets. Using this information, the minimum transmission power for successful DATA/ACK communication may be calculated, and used subsequently.

Adaptive power levels:. Reducing hop-count is clearly beneficial to the performance of an ad hoc network. We plan to evaluate the possibility of adaptively controlling transmit-power levels to reduce the hop-count of discovered routes. For example, consider a scenario in which nodes A, B and C are placed in a straight line. We have observed that the DDSR transmission range is often slightly shorter than the distance between A and C. As a result node A must route packets through node B. However, if the information about the distance between nodes A and C could be made available to node A *a priori*, A could increase its transmission range to ensure connectivity with C. Of course, increasing transmission range not only increases control overhead, but requires additional power consumption. The tradeoff seems interesting and merits investigation.

MAC layer packet scheduling.
In addition to power control, we also plan to maximize the gains due to spatial reuse. More specifically, a node intending to transmit a packet in a particular direction may have to wait for a specific interval of time if the DNAV table requires so. However, the node may have packets (from different flows) that needs to be transmitted in a different direction. Since MAC layers typically maintain a single queue of packets, it suffers the problem of *Head of the Line Blocking (HLB)*. We propose to implement multiple MAC layer queues on a per antenna-element basis. Scheduling packets over all the antenna elements may achieve higher spatial reuse of the medium.

Reduced route alignment
Another way of increasing spatial reuse would be by choosing routes characterized by lesser degree of alignment. We illustrate this using Figure 19. Observe that in Figure 19, two routes are possible between nodes S and D (as shown by a solid line and a dashed line). Although both routes are equal in hop-count, the route through node C may achieve higher spatial reuse. This is because simultaneous communication over links S-A and C-D may be possible. Using route S-A-B-D, instances are possible when transmissions from S to A may interfere (at node D) with a simultaneous transmission from B to D. Therfore, the unaligned route (S-A-C-D) may be a better choice in such cases. We plan to evaluate optimizations by which a destination node selects routes based on the route's *hop-count* as well as the *degree of alignment*. The *degree of alignment* of a route could be defined as a fuction of the antenna-elements used by intermediate nodes in the route.

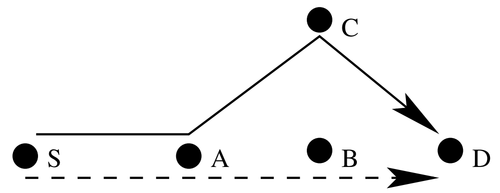

Fig. 19. An example scenario showing two alternative routes between nodes S and D

VII. Conclusion

We addressed the problems related to *ad hoc routing using directional antennas*, and investigated several issues that potentially play a role in determining network performance. We gained insights into some of the key tradeoffs in directional routing. We observed that using directional antennas may not be suitable when the network is dense or linear. However, directional antennas show improvement in throughput for random, sparse networks. Under mobile conditions, the performance improvement is encouraging and antenna handoffs do not pose a serious problem. Although by no means a complete evaluation, our optimizations have produced results that are encouraging and support the incorporation of directional antennas in mobile ad hoc networking.

VIII. Acknowledgments

We thank E. Jung, P. Kyasanur, W. List, M. Miller and X. Yang for their many useful comments and suggestions.

References

[1] S. Bandyopadhyay, K. Hausike, S. Horisawa and S. Tawara, *"An Adaptive MAC and Directional Routing Protocol for Ad Hoc Wireless Networks Using ESPAR Antenna"*, Proc. of the ACM/SIGMOBILE MobiHoc October 2001.

[2] Ravi Govindarajula and M. Capone, *"Enhancing the Capacity of Ad Hoc Networks with Smart Antennas"*, Ongoing Research at Arizona State University.

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

[4] D. Johnson, D. A. Maltz and J. Broch, *"The Dynamic Source Routing for Mobile Ad Hoc Networks"*, Mobile Ad Hoc Network (MANET) Working Group, IETF, March 1998.

[5] Phil Karn, *"MACA - A New Channel Access Method for Packet Radio"*, Appeared in the proceedings of the 9th ARRL Computer Networking Conference, London, Ontario, Canada, 1999

[6] Y. B. Ko, V. Shankarkumar and N. H. Vaidya, *"Medium access control protocols using directional antennas in ad hoc networks"*, Proc. of the IEEE INFOCOM 2000, March 2000.

[7] A. Nasipuri, S. Ye, J. You and R. Hiromoto, *"A MAC Protocol for Mobile Ad Hoc Networks using Directional Antennas"*, Proc. of the IEEE WCNC 2000.

[8] A. Nasipuri et al. *"On Demand Routing Using Directional Antennas in Mobile Ad Hoc Networks"*, Proc. of the IEEE WCNC 2000.

[9] *"Qualnet simulator version 2.6.1"*, Scalable Network Technologies, www.scalable-networks.com.

[10] Ram Ramanathan, *"On the Performance of Beamforming Antennas in Ad Hoc Networks"*, Proc. of the ACM/SIGMOBILE MobiHoc 2001, October 2001.

[11] R. Roy Choudhury and N. Vaidya *"On Ad Hoc Routing Using Directional Antennas"*, Presented at the Illinois Computer Systems Symposium (iCSS), UIUC, May 2002.

[12] R. Roy Choudhury, X. Yang, R. Ramanathan and N. Vaidya, *"Using Directional Antennas for Medium Access Control for Ad Hoc Networks"*, To appear in the Proc. of the ACM MOBICOM 2002, September 2002.

[13] R. Roy Choudhry, Somprakash Bandyopadhyay and K. Paul, *"A Distributed Mechanism for Topology Discovery in Ad Hoc Wireless Networks Using Mobile Agents"*, Proc. of the ACM/SIGMOBILE MobiHoc 2000, August 2000.

[14] M. Sanchez, T. Giles, and J. Zander, *"CSMA/CA with Beam Forming Antennas in Multi-hop Packet Radio,"*, Proc. Swedish Workshop on Wireless Ad Hoc Networks, Mar. 5-6, 2001.

[15] Sze-Yao Ni, Yu-Chee Tseng, Yuh-Shyan Chen and Jang-Ping Sheu, *"The Broadcast Storm Problem in a Mobile Ad Hoc Network"*, Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM 99), WA, USA.

[16] Mineo Takai, Jay Martin, Rajiv Bagrodia and Aifeng Ren, *"Directional Virtual Carrier Sensing for Directional Antennas in Mobile Ad Hoc Networks"*, To appear in the Proc. of the ACM/SIGMOBILE MobiHoc June 2002.

[17] T. S. Yum and K. W. Hung, *"Design Algorithms for Multihop Packet Radio Networks with Multiple Directonal Antennas"*, IEEE Transactions on Communications, vol 40, no. 11 pp 1716-1724, 1992.