Performance of Ad Hoc Routing using Directional Antennas

Romit Roy Choudhury[+] Nitin H. Vaidya[*]

[+]Dept. of Computer Science,
[*]Dept. of Electrical and Computer Engineering, and
Coordinated Science Laboratory
University of Illinois at Urbana Champaign
{croy, nhv}@uiuc.edu

Abstract

This paper evaluates the tradeoffs involved in using directional antennas in ad hoc routing. Although problems with utilizing directional antennas have been visited in the past, the research has been confined mostly 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 paper, we evaluate the performance of DSR (Dynamic Source Routing) using directional antennas. We identify several issues that emerge from executing DSR (originally designed for omnidirectional antennas) over 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.

Keywords: *Directional Antennas*, *Ad Hoc Routing*, *Dynamic Source Routing*, *Directional Medium Access*

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 paper, 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. More subtle tradeoffs arise from directional medium access

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

control. A routing protocol that remains unaware of the underlying changes can degrade system performance significantly. To study such trade-offs, we evaluate the performance of an omnidirectional routing protocol (DSR) using directional antennas. In addition to these discussed issues, energy consumption of directional antennas may also be of interest. In this paper, we do not evaluate the tradeoffs associated to energy consumption – this is a part of our future work. The contributions of this paper includes identifying connectivity, delay, and throughput issues associated to directional routing, optimizations to address some of these issues, and a new metric to evaluate the control overhead of routing protocols.

The rest of the paper 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 medium access control (MAC) protocol, designed for 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 to the routing protocol are also discussed in this section. Section VI discusses our plans for future work. Section VII concludes the paper with a brief summary.

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 [2][4][10][11] [15][17][20]. Work on routing protocols using directional antennas is limited [2][12][14][7] [3]. In [12], 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 under-estimating the potential of directional beamforming. Put differently, the impact of discovering longer links (and therefore fewer-hop routes) using directional antennas has not been studied. In our work, we evaluate the effect of various transmission ranges of directional antennas. [12] 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. In [5], Haas *et al.* proposes *gossip based routing* in which broadcast packets are forwarded by each node with a specified probability p. The paper shows that using such a probabilistic approach, the routing overhead can be reduced by nearly 35 percent. However, the protocol implicitly assumes a dense network with large number of nodes. In this paper we intend to reduce the overhead of broadcasting control messages through selective use of directional antennas.

Ramanathan [14] presents an interesting discussion on several issues arising from directional communication. [14] 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 paper, we focus primarily on issues related to routing using directional antennas. [2] proposes a proactive routing protocol (originally proposed for omnidirectional antennas [16]) 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. In this paper we consider a reactive routing protocol. Also, our MAC protocol is simpler than [2].

In [3], the authors present the notion of *maximally disjoint routes*. Routes that do not share the same geographical area with other routes are preferred, thereby minimizing contention at the MAC layer. In [8], authors propose to use directional antennas to prevent networks from getting partitioned. The authors show that efficient power consumption may be achieved by combining the use of omnidirectional and directional modes of antenna systems. All the above studies highlight the encouraging performance of directional antennas in ad hoc networks. In this paper we identify additional issues that require the routing protocol to be directional-antenna-aware, failing which the system performance may degrade (as shown later).

Recently, [1] has proposed adaptive range control in which the ability of directional reception has been exploited. The authors observe that directional reception can be beneficial since the signal of interest can be received with higher reliability without adding to the total interference on the channel. By adaptively controlling the range of communications based on the neighborhood traffic and activity, the authors illustrate how directional antennas may be beneficial to ad hoc networking. In [15], the authors present directional MAC protocols, identify several problems with directional medium access and evaluate their impact on network performance. However, [15] and [1] both do not address the impact on routing.

III. Preliminaries

A. IEEE 802.11 Distributed Coordinated Function (DCF)

In the IEEE 802.11 MAC protocol [6] 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 [9].

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 [13]. The antenna model is comprised of N beam patterns, sometimes also referred to as radiation patterns. Each beam has a conical radiation pattern, spanning an angle of $\frac{2\pi}{N}$ radians. We assume that the beams are fixed with non-overlapping beam directions, so that the N beams may collectively cover the entire plane as shown in Figure 1(a). We are aware that real antenna systems may not allow strict non-overlapping beams. However, for simplicity of illustration, we depict the beam-patterns as in Figure 1(a). In our simulations, the beams are indeed overlapping.

The beams 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 beam 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 [14], 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 1(b) [14].

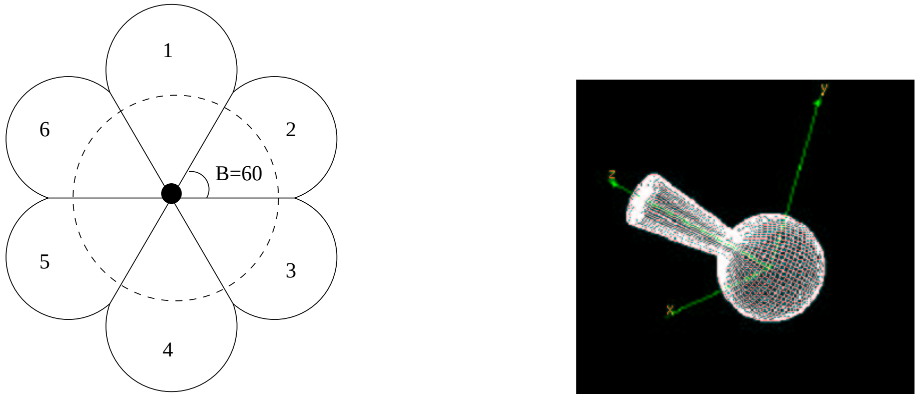

Fig. 1. (a) Left: The radiation pattern of a 6 beam antenna system, with beamwidth B = 60 degrees. The dashed circle outlines the omnidirectional radiation pattern, indicating that the omni transmission-range is shorter than the directional range. (b) Right: A bulb shaped radiation pattern of a directional antenna, showing a spherical side lobe and a conical main lobe [11].

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 beam on which the received signal power is maximum. Rest of the dialog is then carried out using this beam. 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 beams and beamform with a gain of G^d, $G^d >= G^o$. Note that the geographical distance over which two nodes may be able to communicate is proportional to the product of the transmission and the reception gain. As a result, the link-length between directional transmitters and omnidirectional receivers can be longer than that between omnidirectional transmitters and omnidirectional receivers (link-length being defined as the distance between the nodes forming the link). This motivates the notion of directional neighborhood – *Direction-Omni (DO) Neighbors* and *Omni-Omni (OO) Neighbors*. In other words, a node B is a DO-neighbor of a node A if node B can receive a directional transmission from A even if B is in the *Omni* mode. A node B is a OO-neighbor of a node A if node B can receive a omnidirectional transmission from A even if B is in the *Omni* mode. Observe that all OO-neighbors are also DO-neighbors, but not necessarily the vice versa.

Clearly, the notion of *broadcast* changes when directional antennas are used. Since only one beam may be used at a time for directional transmission, an omnidirectional broadcast using directional antennas may be achieved by sequentially using each of the beams to transmit the same packet. For example, if the node shown in Figure 1(a) wishes to broadcast a packet to all its surrounding neighbors, it must transmit the same packet 6 times, once with each conical beam. We call this entire operation *sweeping*. Observe that the coverage of a single *sweep* is larger than the coverage of an omnidirectional broadcast (if both are carried out at the same transmit power level). Put differently, while a single *sweep* reaches all the DO-neighbors of a node, an omnidirectional broadcast reaches only the OO-neighbors. [1]

C. Dynamic Source Routing (DSR) Protocol

We briefly describe the operation of a reactive routing protocol called Dynamic Source Routing (DSR)[7]. We later evaluate the performance of DSR over directional antennas, identify problems arising from directional routing and suggest modifications to DSR to adapt to directional antennas.

The key features of DSR are *reactive route discovery* and *source routing* [7]. 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 in DSR 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 [15][19], 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

[1] One possible way of reaching DO-neighbors without sweeping, could be through higher-power omnidirectional broadcast. However, a high power omnidirectional broadcast not only introduces power assymetry (and thus problems related to it), but may also not be feasible if a communication is in progress in the vicinity. Sweeping on the other hand can avoid transmitting in directions in which it might cause interference and sweep over the remaining directions. This would be explained in the discussion on selective forwarding.

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 beam to be used for RTS transmission. To do this, DiMAC maintains a look-up table containing two fields; *NeighborId*, *AntennaBeam*. The look-up table is a cache that is updated reactively to minimize overhead, as discussed in the next subsections. Using the specified *AntennaBeam*, 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. While the RTS is being received in the *Omni* mode, node R is susceptible to collision due to signals arriving from other directions. DiMAC at node R now determines the antenna beam on which the received signal power of the RTS was maximum, and uses that same beam 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 beam 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. DiMAC maintains a directional NAV (DiNAV) table (similar to [15][19]) comprising of the fields *AntennaBeam* and *WhetherBlocked*. On overhearing a RTS or CTS on *AntennaBeam i*, a node sets the *WhetherBlocked* flag for beam i, for the proposed duration. Subsequently, this node defers transmissions that require use of beam *i*. After the proposed duration is over, the *WhetherBlocked* flag is reset. For example, when node A communicates with node C in Figure 3, node B's beams 2 and 4 are blocked (in Figure 3, note that node B overhears A's RTS on beam 4, and C's CTS on beam 2). However, node B may initiate communication (or reply to a communication request) using the other beams.

A conspicuous problem with a directional MAC protocol such as DiMAC is *"deafness"*, identified and discussed in [15]. Although *deafness* is a MAC layer issue, it affects the performance of routing protocols.

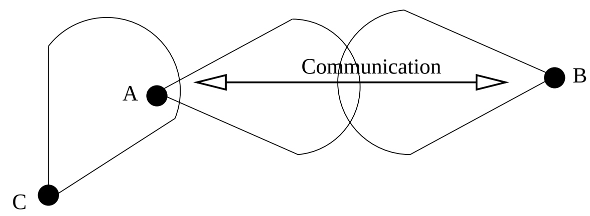

Fig. 2. An example scenario showing the possibility of deafness.

To aid better understanding of the problems in directional routing, 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 beam it is currently using. Consider Figure 2. In Figure 2, 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.

B. DSR over DiMAC

We begin our evaluation of ad hoc routing using DSR and later suggest 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 using DiMAC, an omnidirectional broadcast is emulated through *sweeping*. When a RREQ packet is handed down to DiMAC by DSR, DiMAC sequentially transmits the same packet over all the N beams. Observe that sweeping incurs additional delay, almost N times greater than an omnidirectional broadcast. However, as discussed earlier, *sweeping* allows RREQs to propagate to far away nodes (i.e., to DO-neighbors) in comparison to an omnidirectional broadcast that can reach only the OO-neighbors. For example, in Figure 3, 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 3, then node C would not receive A's transmission, and link A-C cannot be on the chosen route.

While performing a *sweep*, packet transmissions are not preceded by *backing off*. If the beam being used for transmission senses the channel idle, it transmits the RREQ. Antenna beams 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 beams (sequentially). If some of the marked beams now sense the channel idle, the RREQ is transmitted using them. For beams 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

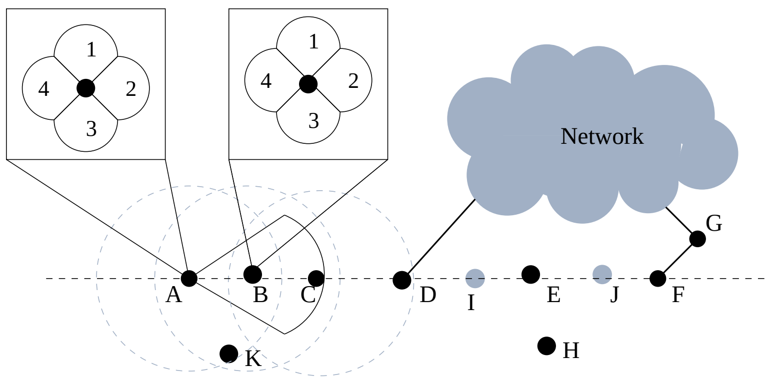

Fig. 3. An example scenario. A 4 beam antenna system is shown in the inset. Node A uses beam 2 to communicate with node C. Node C uses beam 4 to receive signals from mode A. The dashed circles outline the omnidirectional transmission range

antenna that experiences highest signal strength. Thus as long two different signals arrive on different beams, the receiver would always receive at least one of them. [2]

To unicast packets to a particular neighbor, a node must use the appropriate beam. 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 beam* 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 beam i, then to initiate communication to node B in the near future, node A uses the same beam 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 beam specified by the look-up table is no more appropriate. To communicate to the relocated node, a different antenna beam 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 beams (i.e., swept), whenever necessary. If node B receives a "hello" from node A, B replies using the same antenna beam with which it received the "hello". The reply from B is delayed to allow node A to first finish transmitting all the hellos (otherwise, A would not receive B's reply since A would be busy transmitting hellos in other directions). Node A then receives the reply from B, say, using the antenna beam 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 beam j. In summary, *neighbor discovery* becomes more complex using directional antennas since neighbors are now associated to specific antenna beams 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.)

[2] 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.

V. Performance Evaluation

We use the Qualnet simulator [13] 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 calculated using expressions from [14]. Using these antenna gains, the relation between the transmission range and the beamwidth of the beam's main lobe is shown in Figure 4(a). 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 4(a), 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.

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 3 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 3, 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 K 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

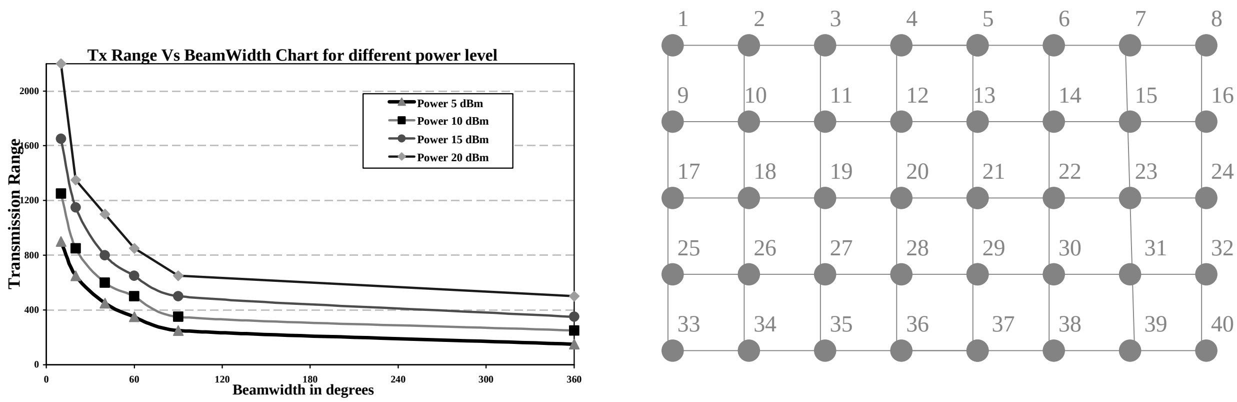

Fig. 4. (a) Left: The transmission ranges for increasing gain, for different power levels. (b) Right: The rectangular grid used for simulations.

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 4(b). 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.

A. Evaluating Route Discovery Latency

Figure 5(b) 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 4(b)). 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 5(b) plots RDL for different values of N, N being the number of beams. For example, DDSR6 refers to Directional DSR simulated over DiMAC with 6 beams. The beamwidth with N beams is $(2\pi/N)$. Hereafter, we would denote Directional DSR simulated over DiMAC with N antenna beams, 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 5(b), 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 analyze the issues and discuss the insights gained from the results.

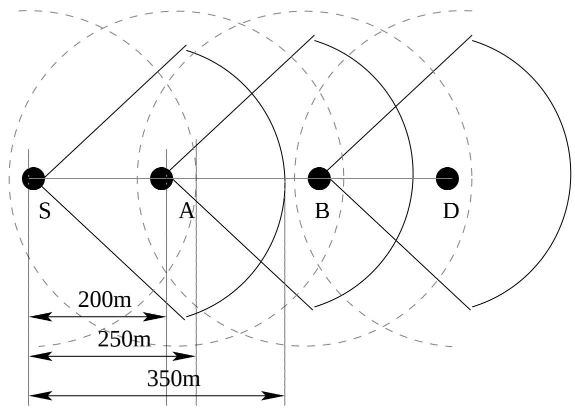

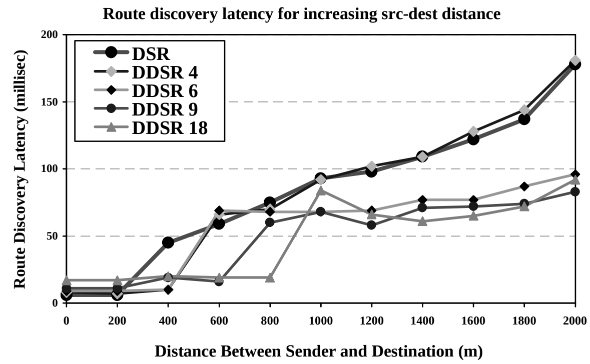

Fig. 5. (a) Left: A section of a grid, illustrating the possibility of DSR and DDSR4 discovering equal-hop routes. (b) Right: Route Discovery latency for varying number of beams at 200m grid-distance.

Behavior of *route discovery latency* (RDL) using directional antennas may be intuitively explained as follows. Several factors influence RDL: *fewer hop routes* and *higher spatial reuse* reduce RDL, while *sweeping delay*, *deafness* and *higher directional interference* increase RDL.[3] 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 5(a). Reaching node D from node S in Figure 5(a) 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-destinations 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 beams, 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 5(a)). 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

[3] *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.*

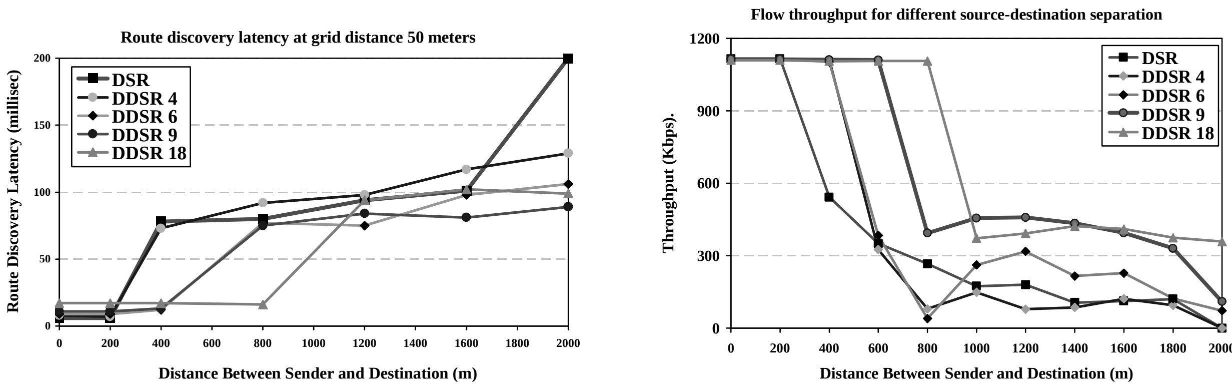

Fig. 6. (a) Left: Route Discovery latency for varying number of beams at 50m node-spacing. (b) Right: Throughput versus distance-of-separation between source and destination, for varying number of beams, at 200m grid-distance

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.

The behavior of route discovery latency depends on node density as well. To illustrate this, consider the DDSR4 and DSR on a grid distance of 200 meters, as discussed above. For a source destination pair separated by 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, DDSR4 could discover a 3 hop route, while DSR must still use a 4 hop route. This motivated us to evaluate the performance of DDSR in a high density network with grid distance of 50 meters. Figure 6(a) shows the results of simulation. Counter to intuition, the performance of DDSR degrades in denser networks. As evident from Figure 6(a), route discovery latency using DSR is comparable to DDSR even at higher transmission ranges and large source-destination separation.

The degraded performance of DDSR in Figure 6(a) is a result of interference. 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. 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 extremely large.

B. Evaluating Throughput

We now focus on throughput when using DDSR. Figure 6(b) shows the throughput of a single CBR flow for different source-destination separations. The traffic generated is large enough to keep the source back-logged 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 hop-count for DDSR can be far lesser than DSR. The reasons are 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 (DO-neighbors) of a node to get the same RREQ at different points of time. Consequently, neighbors that receive the RREQs earlier, have a higher probability of delivering the RREQs earlier to the destination. Observe that earlier delivery in this case may not indicate a shorter-hop route (omnidirectional routing protocols assume that the earliest arriving RREQ traversed the shortest-hop route). If the routing protocol requires the destination to reply only to the earliest arriving RREQ, a potential suboptimal route can be established. Prolonged use of such a suboptimal route may result in serious performance degradation. Replying to the first arriving RREQ, therefore, may not be suitable when sweeping is performed. Interestingly, replying to all RREQs (as optionally performed by DSR) is also not suitable when using DDSR. 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 [15][14]. 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 observation that the earliest RREQ received by the destination may not have traversed the optimal path when *sweeping* is used. This optimization requires the destination node to delay sending the route reply (RREP) by a time duration T, calculated from the time it received the first RREQ. This allows the destination node to choose the best among all the routes that arrive within this time T. We specify T as $T = \rho \times T_{sweep}$, where T_{sweep} is the time taken to complete one full sweep, and ρ is a configuration parameter. We argue that replying to all the RREQs that arrive at the destination, as performed optionally by DSR, 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 optimal 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. To maintain a low route discovery latency for this optimization, we have chosen the value of ρ to be 1 henceforth. Of course, route discovery latency of DDSR increases by T time units. However, as shown later, ensuring a shortest hop route at the expense of somewhat higher route discovery latency favors the performance tradeoff, especially when route discovery is not too frequent.

The throughput when using directional antennas is better than omnidirectional antennas for a single flow. However, in presence of multiple flows, one may expect DDSR to further outperform DSR, due to spatial reuse of the channel. To verify this, we simulated a rectangular grid topology with 3 flows as shown in Figure 7(a). The dashed lines show an example of routes discovered by DDSR. Adjacent nodes in the grid are separated by 200 meters. Surprisingly, Figure 7(b) 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 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 7(a). 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. In summary, the performance of DDSR degrades with multiple flows in a grid topology.

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 such random topologies, with 3 identical flows used previously for the grid topology. Figure 8(a) shows one such random topology and outlines the routes taken by DDSR and DSR for one of the three flows. Figure 7(b) 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

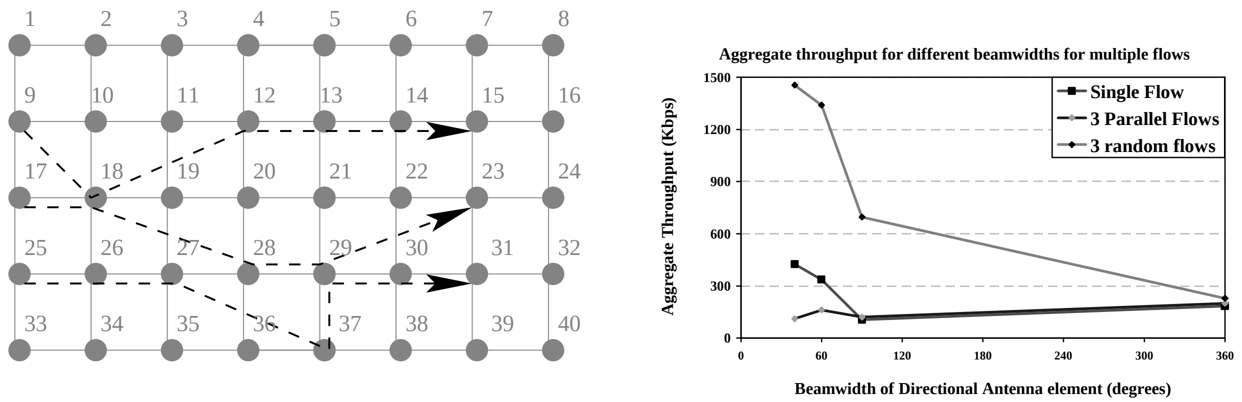

Fig. 7. (a) Left: A rectangular grid topology with 3 flows (node 9 to 15, 17 to 23 and 25 to 31. (b)Right: Performance of single and multiple flows in linear and random topologies.

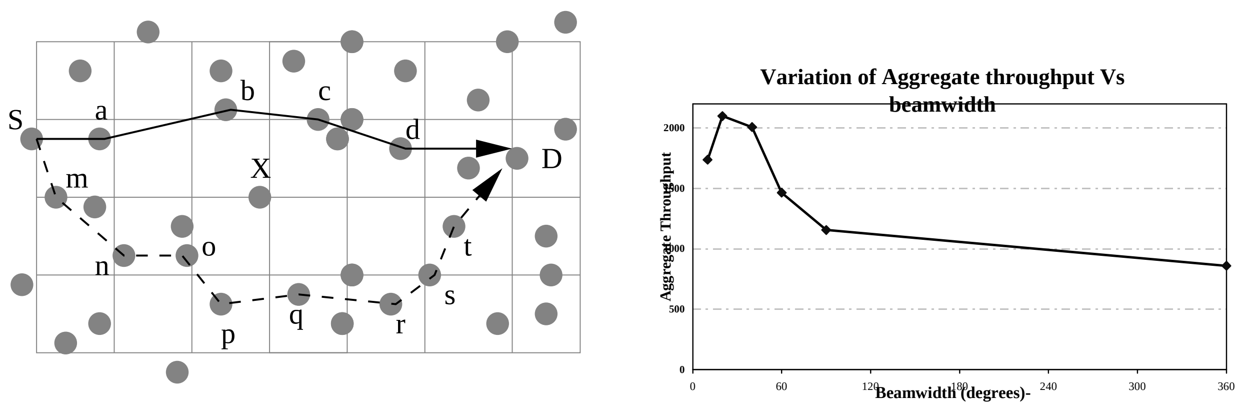

Fig. 8. (a) Left: A random topology showing the possibility of better connectivity due to higher transmission. (b) Right: Variation of performance for different beamwidths in a static network.

antennas is effective in bridging "voids" in the topology. To illustrate this with an example, assume the scenario in Figure 8(a). 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 8(b). 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 beams 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 with extremely 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 add up to significant overhead in comparison to a single omnidirectional broadcast. We compare control overhead (of DSR and DDSR) using the following overhead metric α.

$$\alpha = \frac{\Sigma NumberofControlPackets \times AreaBlockedbyEachPacket}{\Sigma No.ofDataPackets}$$

where *Area Blocked by Each Packet* is the approximately conical area covered by the radiation pattern of a single beam. Intuitively, network capacity consumed due to the transmission of each control packet is proportional to the region silenced by that transmission. The total control overhead is thus equivalent to *Number of Control Packets* $\times$ *Area Blocked by each packet*. Put differently, we argue that using a metric like "*total number of transmitted control packets*" as a metric of control overhead is not appropriate, because, when compared to omnidirectional antennas, directional beamforming can achieve greater number of simultaneous transmissions. This indicates that the consumed capacity due to each packet is only proportional to its wireless footprint, as captured by the metric above. Since the control overhead is in exchange of the total *Number of Data Packets* transmitted 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.

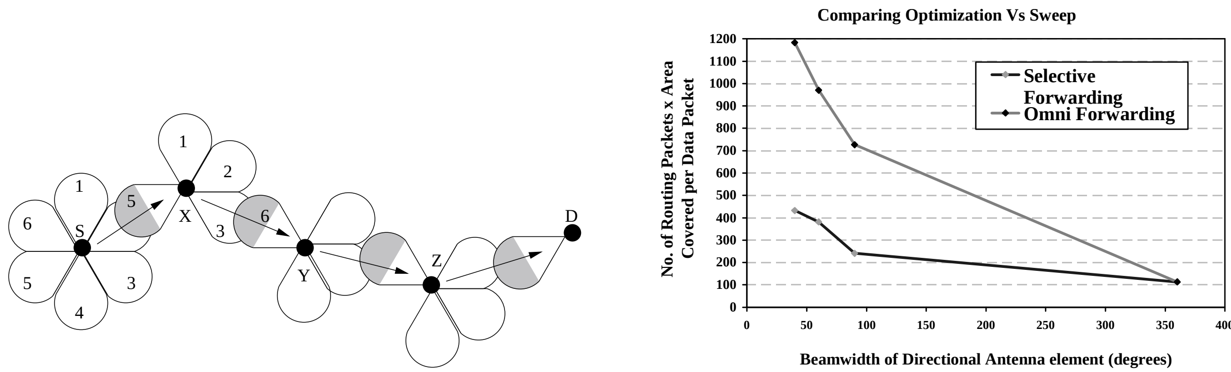

Fig. 9. (a) Left: An example scenario showing the propagation of RREQ using the *selective forwarding* optimization. The beams with which nodes receive RREQs are shaded. (b) Right: Routing overhead α for optimized and unoptimized DDSR.

Selective Forwarding Optimization:

In Figure 9(a), when node Y receives a RREQ from node X, it may not be necessary to transmit that packet back in the direction of node X. This is because, nodes lying in the coverage of node X's beam 3, can receive the packet from X and need not receive a copy of the same packet from Y again. In omnidirectional scenarios, the problem of redundant reception of RREQs has been referred to as the *broadcast storm* problem [18], 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)* beams. A node initiating the control packet, however, forwards it using all beams.

We define $\gamma = \lceil \frac{n}{N} \rceil$. Parameter γ may be treated as a configuration parameter. We have used $\gamma = \frac{1}{2}$ for our simulations. Also, in our simulations, the *n* beams used to forward control packets are the ones that are diagonally opposite to the beam with which the control packet was received. More specifically, if $\gamma = \frac{1}{2}$, N = 6 (i.e., 6 beams numbered from 1 to 6), and if a node receives a RREQ using beam 5, then antenna beams 1, 2 and 3 are used to forward the RREQ. The choice of the value of γ and the use of diagonally opposite beams are motivated by fact that the RREQs can be forwarded outward from the source node. Figure 9(a) shows how a RREQ may propagate using $\gamma = \frac{1}{2}$ and N = 6. Observe that node X receives the RREQ using beam 5 and forwards it using beams 1, 2 and 3.

Figure 9(b) 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.

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

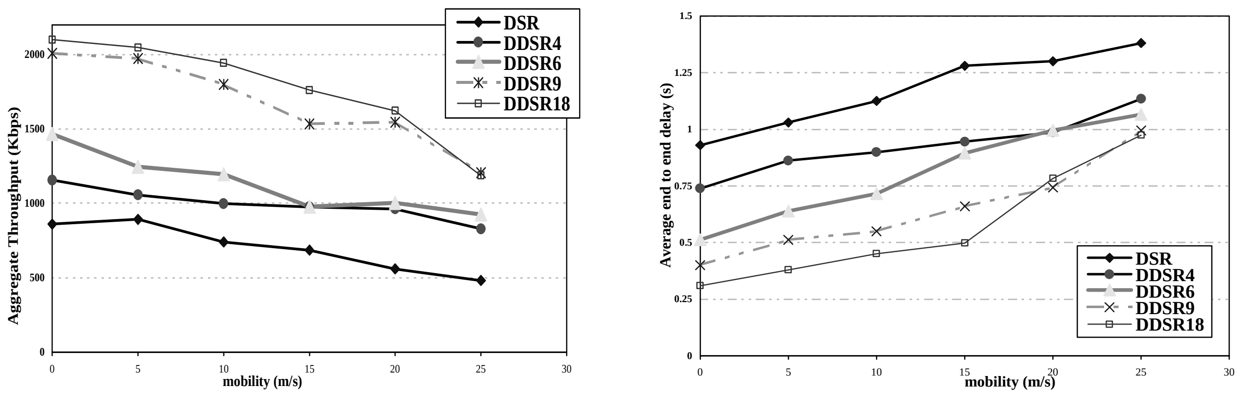

Fig. 10. (a) Left: Aggregate throughput measured for mobile scenarios, for varying transmission ranges. (b) Right: Average end to end delay for mobile scenarios, for varying transmission ranges.

Figure 10(a) shows a graph that measures aggregate throughput for several mobile scenarios and different beamwidths. A *random waypoint* mobility model [13][7] is assumed, with speed varying between 5 to 25 m/s. 10 pairs of source and destination nodes are chosen for each of the scenarios. The results are an average of 25 runs for each specified mobility. 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 10(a), 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 begins to offset the benefits of using directional antennas. The overhead of frequent sweeping increases, leading to degrading throughput of DDSR. However, below a mobility of $25m/s$, DDSR always outperforms DSR.

Figure 10(b) shows the measurement of average end-to-end delay for data packets in mobile scenarios. As seen in Figure 10(b), 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.

VI. Conclusion

This paper evaluates the impact of directional antennas on ad hoc routing. We identify several problems that emerge from executing DSR (originally designed for omnidirectional antennas) over directional antennas. Using insights from simulation, we proposed strategies to adapt DSR over directional antennas - leading to Directional DSR (DDSR). Performance evaluation suggests that using directional antennas may not be suitable when the network is dense or linear. However, the improvement in performance is encouraging for networks with sparse and random topologies. Under mobile conditions, DDSR outperforms DSR and the problem of handoff does not pose a serious problem. Although by no means a complete work on directional routing, our evaluation and optimizations have produced results that are encouraging, suggesting the need to use directional antennas more efficiently in ad hoc networking.

VII. Acknowledgments

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

References

[1] M. Takai, J. Zhou and R. Bagrodia, *”Adaptive Range Control using Directional Antennas in Mobile Ad Hoc Networks”*, Proc. of the 6th international workshop on Modeling analysis and simulation of wireless and mobile systems, 2003.

[2] 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.

[3] S. Bandyopadhyay, D. Saha, S. Roy and T. Euda *”A Network Aware MAC and Routing Protocol for Effective Load Balancing in Ad Hoc Networks with Directional Antennas”*, Accepted for publication in ACM/SIGMOBILE MobiHoc 2003.

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

[5] Zygmunt J. Haas, Joseph Y. Halpern and Li Li, *”Gossip-Based Ad Hoc Routing”*, Proc. of the IEEE INFOCOM 2002.

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

[7] 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.

[8] A. saha and D. Johnson, *”Energy Saving and Partition Bridging using Directional Antennas in mobile ad hoc networks”*, Submitted Master’s Thesis, Rice University.

[9] 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

[10] 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.

[11] 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.

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

[13] *"Qualnet simulator version 2.6.1"*, Scalable Network Technologies, www.scalable-networks.com.
[14] Ram Ramanathan, *"On the Performance of Beamforming Antennas in Ad Hoc Networks"*, Proc. of the ACM/SIGMOBILE MobiHoc 2001, October 2001.
[15] 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.
[16] R. Roy Choudhury, 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.
[17] 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.
[18] 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.
[19] 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.
[20] 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.