

Impact of Directional Antennas on Ad Hoc Routing ^{*}

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, USA

Abstract. Previous research on directional antennas has been confined mostly to medium access control. However, it is necessary to evaluate the impact of directional antennas on the performance of routing protocols as well. In this paper, we identify the issues and evaluate the performance of an omnidirectional routing protocol, DSR, when executed over directional antennas. Using insights gained from simulations, we propose routing strategies suitable for directional communication. Our analysis shows that by using directional antennas, ad hoc networks may achieve better performance. However, scenarios exist in which omnidirectional antennas may be suitable.

1 Introduction

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 higher transmission ranges 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

^{*} This work is supported in part by National Science Foundation (NSF) under grant ANI 01-96410 and 01-25859.

antennas). Also, sweeping incurs greater delay and higher control overhead, partially negating the potential advantages. More subtle tradeoffs also arise.

A routing protocol that remains unaware of such changes in the radio hardware can degrade system performance. To study the impact, we evaluate an omnidirectional routing protocol called DSR (Dynamic Source Routing), over directional antennas, and propose modifications. In [4], [7], [9], [11] authors addressed the problem of medium access control using beamforming antennas. However, work on directional routing protocols is limited [5],[7],[1]. In [5] directional antennas are utilized for the purpose of on-demand routing. In [1], the authors present the notion of *maximally disjoint routes*, in which routes that do not share the same geographical area are chosen. Our contribution in this paper includes identifying issues in directional ad hoc routing, optimizations to address some of the issues, and a new metric to evaluate the control overhead of routing protocols.

2 Antenna Model

The antenna model we use is comprised of N beam patterns. The main lobe of each beam has a conical radiation pattern, spanning an angle of $\frac{2\pi}{N}$ radians. In our implementation of the antenna, motivated by the model in [7], we have approximated all the side lobes into a single sphere, with the node at its center. 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 packet is then received by 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 , $G^d \geq G^o$. 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. This motivates the notion of directional neighborhood – *Direction-Omni (DO) Neighbors* and *Omni-Omni (OO) Neighbors*. 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 an 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.

The notion of *broadcast* changes when directional antennas are used. For example, if a node wishes to broadcast a packet to all its surrounding neighbors, it must transmit the same packet N times, once with each conical beam. We call this entire operation, *sweeping*. Observe that a single *sweep* reaches all the DO-neighbors of a node, while an omnidirectional broadcast reaches only the OO-neighbors.

3 DiMAC Protocol

We have designed 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. DiMAC at S consults a look-up table to determine the antenna beam that must be used to initiate RTS to R (discussed later). Once the beam is determined, 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[2]) before it transmits the RTS.

Node R , while idle, listens to the channel in omni mode and receives the RTS meant for it. The RTS is received with omnidirectional gain G^o . 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 directionally. 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, to which R replies with a directional ACK. 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 table (DiNAV), similar to [9],[11], that remembers the beams on which RTSs or CTSs were received. Transmission is deferred for the proposed duration, only on these beams.

A conspicuous problem with a directional MAC protocol such as DiMAC is “*deafness*”, identified and discussed in [9]. Although a MAC layer issue, it affects the performance of routing protocols. Briefly, deafness is caused when a node C attempts to initiate dialogue with a node A, while A is engaged in communication with another node, B. Node A fails to receive signals from C since it remains beamformed towards B over the duration of communication. Node C interprets the absence of a reply from A as indicative of a collision, and retransmits the packet. This can repeat multiple times, until node C is forced to give up and drop the packet.

DSR over DiMAC

We begin our evaluation of ad hoc routing using DSR and suggest improvements later. 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*. Recall that sweeping incurs additional delay (almost N times greater), but can reach the DO-neighbors of a node, unlike omnidirectional broadcasts. For example, in Figure 1, if an RREQ is transmitted directionally by A, link A-C can be on the route. However, if the RREQ is transmitted omnidirectionally, then node C would not receive A’s transmission. Now, 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. In mobile scenarios, it is possible that information cached in the look-up table gets stale. We incorporate a *scanning* mechanism to address this problem. A *scan* is essentially “hello” packets, transmitted se-




Fig. 1. An example scenario. A 4 beam antenna system is shown in the inset.

quentially 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”. Node A records the antenna beam with which it receives the reply, and initiates future communication with B using the same beam. 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. (In the next section, we suggest a partial scanning optimization to reduce the overhead of scanning.)

4 Performance Evaluation

We use the Qualnet simulator [6] version 2.6.1, for modeling our antenna system and simulating the protocols. We simulate our scenarios in a bounded region of 1500 X 1500 meters, with *Constant Bit Rate* (CBR) traffic for data communication. 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, shorter routes could be discovered due to a larger transmission range of directional antennas. Figure 1 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 1, 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. 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.

The simulated topology is a rectangular grid, with nodes placed 200 meters apart. To evaluate network behavior, we measured *route discovery latency* (RDL) and *throughput* for DSR and DDSR. 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.

Evaluating Route Discovery Latency

Figure 2(b) shows the variation of *route discovery latency* (RDL) versus the physical distance separating source and destination nodes. For a given source node, the destination node is chosen according to the desired distance of separation. Figure 2(b) plots RDL for different values of N , N being the number of beams. Hereafter, we would denote Directional DSR with N antenna beams, as “DDSR N ”. For example, DDSR6 refers to DSR simulated over DiMAC with 6 beams. The beamwidth with N beams is $(2\pi/N)$. By “DSR”, we would indicate DSR executed over omnidirectional antennas, using IEEE 802.11 (this is equivalent to DSR1). From Figure 2(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 for DDSR4) exhibit lower route discovery latency in comparison to omnidirectional antennas.

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 (discussed in detail in [8]). 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, reaching node D from node S in Figure 2(a) requires three hops for both DDSR4 and DSR (DDSR6 or DDSR9 would require at least two hops). However, in such scenarios, the *sweeping* delay of DDSR is high, easily offsetting the slight gain (if any) of a shorter directional route.

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 RDL for DDSR4 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 DDSR4 can at best communicate directly to an adjacent node in the grid (Figure 2(a)). DSR, with a transmission range of 250 meters, can also do the same. Thus DDSR4 does not show any benefit over DSR. In fact, due to sweeping delay, the performance of DDSR4 is somewhat worse than DSR.

The behavior of route discovery latency depends on node density as well. To illustrate this, consider 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 is expected to discover a 3 hop route, while DSR must still require 4 hops. This motivated us to evaluate the performance of DDSR in a high density network with grid distance of 50 meters. Figure 3(a) shows the results of simulation. Counter to intuition, the performance of DDSR degrades in denser networks, even at high transmission ranges and large source-destination separation.

The degraded performance of DDSR in Figure 3(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. 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.

Evaluating Throughput

Figure 3(b) 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. 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 2 key 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 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. 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 a "*Delayed Route Reply*" optimization, to handle this problem.
2. Furthermore, MAC layer studies indicate that directional antennas perform poorly in "linear" topologies [9][7]. Performance in random scenarios may be significantly different, as observed later.

Delayed Route Reply Optimization: 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.

For a single flow, directional antennas achieve higher throughput than omnidirectional antennas. 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 4(a). The dashed lines show an example of routes discovered by DDSR. Adjacent nodes in the grid are separated by 200 meters. Surprisingly, Figure 4(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. A detailed discussion on these issues is available in [8], [9]. To confirm that




Fig. 2. (a)Left: Section of a grid showing DSR and DDSR4 discover equal-hop routes. (b)Right: RDL at 200m grid-distance.




Fig. 3. (a)Left: RDL at 50m node-spacing. (b)Right: Throughput versus source-destination separation, at 200m grid-distance.




Fig. 4. (a)Left: Grid topology with 3 flows. (b)Right: Performance of single and multiple flows in linear and random topologies.




Fig. 5. (a)Left: A random scenario. (b)Right: Throughput variation with beamwidth.

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 5(a) shows one such random topology and outlines the routes taken by DDSR and DSR for one of the three flows. Figure 5(b) shows the average throughput of DSR and DDSR, simulated in random topologies (recall that DSR beamwidth equals 360 degrees). Clearly, DDSR performs better than DSR. Closer examination reveals that higher transmission ranges of directional antennas may be effective in bridging “voids” in the topology. Consider Figure 5(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 a and b , bridging, what may be called a “void region” for omnidirectional transmissions.

The advantage of smaller hop-count, combined with spatial reuse, can lead to higher aggregate throughput with DDSR as compared to DSR. This is evident in Figure 5(b). The graph shows the average of aggregate throughput over 25 random topologies for increasing beamwidth, hence decreasing transmission ranges. While aggregate throughput decreases with decrease in transmission range, at very high transmission ranges the trend is reversed. This is because at higher transmission ranges (i.e., at narrower beamwidths), the number of beams required for a single sweep increases proportionally. This not only increases *sweeping delay* but exacerbates *deafness* as well. Nodes that are capable of forming the shortest route often do not receive the RREQ. Sub-optimal routes are discovered, and performance can degrade with extremely narrow beamwidths.

Routing Overhead

Sweeping requires N sequential transmission of the same packet in N different

directions, increasing control overhead. We compare control overhead (of DSR and DDSR) using the following overhead metric α .

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

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 interference footprint of that transmission. The total control overhead is thus equivalent to




Fig. 6. (a) Left: Propagation of RREQ when using *selective forwarding* optimization. (b) Right: Comparing routing overhead α for DSR and DDSR

Number of Control Packets \times *Area Blocked by each packet*. This overhead is in exchange of the total *Number of Useful Data Packets* transmitted in the network. We use the ratio of these two quantities as the metric α .

The *sweeping* overhead of directional antennas is clearly larger than omnidirectional antennas. To address this in DDSR, we propose the following optimization.

Selective Forwarding Optimization: In Figure 6(a), when node Y receives a RREQ from node X, it may not be necessary to forward the RREQ 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 [10], since control packets are unnecessarily broadcast in the channel. Using directional antennas, we consider an optimization whereby a node forwards a control packet

with only n ($n \leq N$) beams. A node initiating the control packet, however, forwards it using all beams. We define a configuration parameter $\gamma = \lceil \frac{n}{N} \rceil$, $\gamma = \frac{1}{2}$ 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. Figure 6(a) shows how a RREQ may propagate when $\gamma = \frac{1}{2}$ and $N = 6$. Observe that node X receives the RREQ using beam 5 and forwards it using beams 1, 2 and 3.

Figure 6(b) plots routing overhead α against beamwidth of directional antennas. α was calculated from 25 simulations of random, static topologies. This optimization reduces the routing overhead of DDSR, although it remains higher than DSR.




Fig. 7. Aggregate throughput and Average end to end delay for mobile scenarios.

Mobility and Route Error

Our observations until this point are under the absence of node mobility. In the presence of node mobility, link failures may happen when nodes move too far away, or when a node moves out of the radiation pattern of one beam to that of the other – called *handoff*. Look-up tables thus get stale. In contrast, due to higher transmission range of directional antennas, links may persist for a longer duration when the end nodes of the link are moving away from each other. This reduces the frequency of route errors.

We have handled the problem of handoff at the MAC layer and DDSR remains unaware of it. If DiMAC does not receive a reply from a neighboring node after a small number of attempts, it assumes handoff, and attempts to find the appro-

priate beam that could re-establish communication. The node performs a partial *scan*, whereby the transmitter node sends “hello” packets using K beams, adjacent to the beam that was in use previously (detailed in [8]). Figure 7(a) shows a graph that measures aggregate throughput for several mobile scenarios and different beamwidths. A *random waypoint* mobility model [6],[3] is assumed. 10 source-destination pairs are chosen, and the time of simulation is proportionally reduced when mobility is increased. From Figure 7(a), aggregate throughput of DDSR degrades with mobility. This happens because the time for route discovery is almost a constant. 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, and data transmission is disrupted for 5 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, DDSR achieves higher aggregate throughput, even at higher mobility. This is because, due to the higher transmission range of DDSR, links break less frequently, which in turn reduces number of route rediscoveries. Average end-to-end delay of DDSR is also observed to be lower than DSR, although between the variations of DDSR the differences are small.

5 Conclusion

We identify several problems that emerge from executing DSR over directional antennas, and propose suitable modifications - leading to Directional DSR (DDSR). Our strategies may be applicable to other routing protocols that use directional antennas. 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 handoff does not pose a serious problem. Our evaluation and optimizations have produced results that are encouraging, suggesting the need to utilize directional antennas more efficiently in ad hoc networking.

References

1. S. Bandyopadhyay, D. Saha, S. Roy and T. Ueda “*A Network Aware MAC and Routing Protocol for Effective Load Balancing in Ad Hoc Networks with Directional Antennas*”, Proc. of ACM/SIGMOBILE MobiHoc 2003.
2. IEEE, “*Wireless LAN Medium Access Control(MAC) and Physical Layer(PHY) Specifications*”, IEEE Standard 802.11, June 1999.
3. D. Johnson, D. Maltz and J. Broch, “*Dynamic Source Routing for Mobile Ad Hoc Networks*”, Mobile Ad Hoc Network Working Group, IETF, March 98.
4. 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.
5. A. Nasipuri et al. “*On Demand Routing Using Directional Antennas in Mobile Ad Hoc Networks*”, Proc. of the IEEE WCNC 2000.
6. “*Qualnet 2.6.1*”, Scalable Network Technologies, www.scalable-networks.com.
7. Ram Ramanathan, “*On the Performance of Beamforming Antennas in Ad Hoc Networks*”, Proc. of the ACM/SIGMOBILE MobiHoc 2001, October 2001.
8. R. Roy Choudhury and N. Vaidya, “*Ad Hoc Routing Using Directional Antennas*”, Technical Report, <http://www.crhc.uiuc.edu/nhv/publications.html>
9. R. Roy Choudhury, X. Yang, R. Ramanathan and N. Vaidya, “*Using Directional Antennas for Medium Access Control for Ad Hoc Networks*”, Proc. of the ACM MOBICOM 2002, September 2002.
10. Sze-Yao Ni, Yu-Chee Tseng, Yuh-Shyan Chen and Jang-Ping Sheu, “*The Broadcast Storm Problem in a Mobile Ad Hoc Network*”, Proc. of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM 99), WA, USA.
11. Mineo Takai, Jay Martin, Rajiv Bagrodia and Aifeng Ren, “*Directional Virtual Carrier Sensing for Directional Antennas in Mobile Ad Hoc Networks*”, Proc. of the ACM/SIGMOBILE MobiHoc June 2002.