

Routing Performance in the Presence of Unidirectional Links in Multihop Wireless Networks

Mahesh K. Marina
Department of ECECS
University of Cincinnati
Cincinnati, OH 45221-0030 USA
mmarina@ececs.uc.edu

Samir R. Das
Department of ECECS
University of Cincinnati
Cincinnati, OH 45221-0030 USA
sdas@ececs.uc.edu

ABSTRACT

We examine two aspects concerning the influence of unidirectional links on routing performance in multihop wireless networks. In the first part of the paper, we evaluate the benefit from utilizing unidirectional links for routing, as opposed to using only bidirectional links. Our evaluations are based on three transmit power assignment models that reflect some realistic network scenarios with unidirectional links. Our results indicate that the marginal benefit of using a high-overhead routing protocol to utilize unidirectional links is questionable.

Most common routing protocols, however, simply assume that all network links are bidirectional, and thus may need additional protocol actions to remove unidirectional links from route computations. In the second part of the paper, we investigate this issue using a well known on-demand routing protocol, Ad hoc On-demand Distance Vector (AODV), as a case study. We study the performance of three techniques for AODV for efficient operation in presence of unidirectional links, viz., BlackListing, Hello, and ReversePathSearch. While BlackListing and Hello techniques explicitly eliminate unidirectional links, the ReversePathSearch technique exploits the greater network connectivity offered by the existence of multiple paths between nodes. Performance results using ns-2 simulations, under varying number of unidirectional links and node speeds, show that all three techniques improve performance by avoiding unidirectional links, the ReversePathSearch technique being the most effective.

Categories and Subject Descriptors

C.2.2 [Computer-Communication Networks]: Network Protocols.

General Terms

Performance, Algorithms.

Keywords

Multihop wireless networks, Ad hoc networks, Unidirectional links, Asymmetric links, Routing, On-demand routing, Multipath routing.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MOBIHOC'02, June 9-11, 2002, EPFL Lausanne, Switzerland.
Copyright 2002 ACM 1-58113-501-7/02/0006 ...\$5.00.

1. INTRODUCTION

A unidirectional link arises between a pair of nodes in a network when only one of the two nodes can directly communicate with the other node. In multihop wireless networks (also known as ad hoc networks), unidirectional links originate because of several reasons. These include difference in radio transceiver capabilities of nodes, the use of transmission range control, difference in wireless channel interference experienced by different nodes. Depending on such conditions, unidirectional links can be quite common.

This paper addresses routing in the presence of unidirectional links. Most of the previous work on this problem concentrated on developing routing protocols [1, 10, 39, 31, 33, 2], or techniques such as tunneling [23] to allow the use of unidirectional links. But the resulting performance advantages and tradeoffs are not well understood. Our approach in this work is to empirically study the influence of unidirectional links on routing performance.

Utilizing unidirectional links along with bidirectional links for routing has two conceivable advantages over using only bidirectional links. First, they can improve the network connectivity. For example, removal of unidirectional links in Figure 1 partitions the network. Second, they can provide better, i.e., shorter, paths. In Figure 1, node B can communicate to node C directly in one hop by using the unidirectional link $B \rightarrow C$ where as the alternate bidirectional path $B - E - C$ requires two hops.

But routing using unidirectional links is complex and entails high overheads. Main difficulty comes from the asymmetric knowledge about a unidirectional link at its end nodes. A node downstream of a unidirectional link (for example, node F in Figure 1) immediately knows about the incoming unidirectional link (link $E \rightarrow F$) on hearing a transmission from the upstream node (node E); but the upstream node may not know about its outgoing unidirectional link until the downstream node explicitly informs it over a multihop reverse path (say, $F - A \rightarrow B - E$ path). Learning about the unidirectional link thus incurs higher overhead than when the link is bidirectional.

There is evidence in the literature that routing protocols finding unidirectional paths (paths with one or more unidirectional links) are subject to higher overheads than those finding only bidirectional paths. For distance-vector protocols, Gerla et al. [10] and Prakash [33] independently make this observation. In the realm of on-demand protocols for ad hoc networks also, similar observation can be made. DSR [15] requires two route discoveries to discover unidirectional paths — one from the source and the other from the destination, as opposed to a single route discovery to find bidirectional paths. Although pure link-state protocols such as OSPF [21] may be able to support unidirectional links with least additional overhead, they already have very high overheads compared

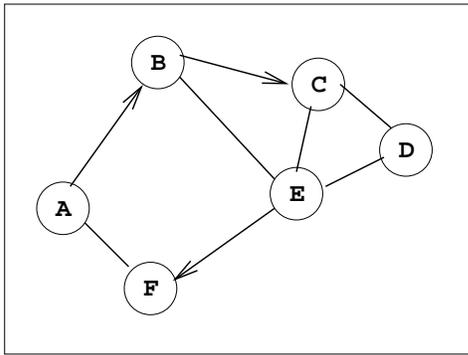


Figure 1: A network with unidirectional links.

to other competing protocols for ad hoc networks [6]. Exactly for this reason, efficient variants of link-state protocols (e.g., TBRPF [24], OLSR [5]) have been developed. But these protocols work only with bidirectional links.

Besides, the use of unidirectional links poses problems to existing link-layer protocols. Many common link-layer protocols for medium access and address resolution not only assume bidirectional links, but also very much depend on two-way handshakes and acknowledgments for their operation. For example, the well-known IEEE 802.11 DCF MAC protocol [14] depends on the exchange of RTS-CTS control packets between the sender and the receiver to prevent hidden terminal collisions, and also expects acknowledgment from the receiver to judge correct packet reception of unicast transmissions.

Thus it is important to evaluate the potential benefit of unidirectional links to know whether employing a seemingly high overhead unidirectional link routing protocol is justified. We evaluate this benefit in the first part of the paper. We compare two routing approaches: (i) using both unidirectional and bidirectional links; (ii) using only bidirectional links. We look at the idealized routing performance obtained from these two approaches — independent of specific routing protocols or associated overheads — to find out any performance advantages of utilizing unidirectional links. To accomplish this goal, we simulate a large number of random multihop network topologies with unidirectional links. We study the connectivity and path cost metrics of these topologies when unidirectional links are used, and when they are ignored. In order to create unidirectional links, we use three models that assign variable transmission ranges to nodes. These models reflect some realistic wireless network scenarios having unidirectional links. Our results show that the connectivity advantage using unidirectional links is almost non-existent, but shortest path costs show some improvement with unidirectional links. However, these improvements too go away when hop-by-hop acknowledgment costs are accounted.

Utilizing unidirectional links for routing purposes may not be efficient, and most routing protocols indeed work with only bidirectional links. But these protocols must still need additional mechanisms to “eliminate” unidirectional links from route computations when they are present. We investigate the importance of such mechanisms using a well-known on-demand routing protocol called Ad hoc On-demand Distance Vector (AODV) [29, 28]. The basic AODV protocol works only with bidirectional links. We propose a new technique called ReversePathSearch to handle unidirectional links in AODV. This technique takes advantage of multiple paths between nodes to overcome unidirectional links. We also consider two other techniques — BlackListing and Hello — that explicitly

eliminate unidirectional links in AODV. Using ns-2 simulations, we evaluate the performance of these three techniques relative to basic AODV.

The rest of the paper is organized as follows. Section 2 evaluates the idealized performance advantage of routing using unidirectional links. Section 3 investigates the issue of avoiding unidirectional links from route computations in the context of AODV. Here we present and evaluate three techniques to improve basic AODV performance in networks with unidirectional links. Section 4 reviews related work. Finally, we present our conclusions in Section 5.

2. BENEFIT OF UNIDIRECTIONAL LINKS FOR ROUTING

2.1 Unidirectional Link Scenarios

Generally speaking, there are two principal reasons behind the presence of unidirectional links in multihop wireless networks. First, difference in radio transmission power level (or receiver sensitivity) of the nodes give rise to unidirectional links. When two nodes (say A and B) have widely different radio transmission ranges¹ so that node A can transmit to node B but not the other way, a unidirectional link forms from node A to node B. Nodes may naturally have different transmission ranges in a heterogenous network where there is an inherent difference in radio capabilities. Alternatively, they may have different transmission ranges when power control algorithms are used for energy savings or topology control.

Second, difference in interference (or noise) at different nodes cause asymmetric links. Asymmetric links occur between a pair of nodes if the link quality is different in each direction. An extreme case of link asymmetry leads to unidirectional links. Nodes may experience different interference levels because of wireless channel imperfections such as multipath fading and shadowing. Hidden terminals can be another cause of wide variation in interference levels (as also mentioned in [33]).

To study the benefit of unidirectional links for routing, we only consider unidirectional links arising from difference in transmission ranges. Note that for unidirectional links to be effective for routing, they should exist long enough for the routing protocol to compute routes through them and to later use such routes to forward some data. Unidirectional links caused by variation in interference levels presumably happen on much smaller time-scales than would be needed for routing. So here we limit ourselves to unidirectional links from variable transmission ranges. However, later in the paper, we do take into account the issue of interference to some extent, specifically from hidden terminals, when studying the negative influence of unidirectional links on routing.

Among the power control algorithms, only a particular class is relevant here. Some power control algorithms prescribe a common power level for all nodes in the network (e.g., [22]). These algorithms do not create any unidirectional links. Other algorithms that do allow variable power levels either assign power levels on a per-transmission basis (e.g., [41, 20]), or assign power levels independent of any single transmission, but may be used for several successive transmissions (e.g., [13, 37, 34, 40]). Since the former set of algorithms may result in short-term unidirectional links, we limit our attention to the latter class only.

¹We will sometimes use transmission range instead of transmission power to simplify description. Note that it is usually straightforward to compute the “nominal” transmission range of a node given the transmission power, large-scale path loss model, and the radio parameters.

2.2 Models for Variable Transmit Power Assignment

Based on the above discussion, we use the following models in our evaluations to create networks with unidirectional links.

TwoPower In this model, the transmission range of a node can take one of two possible values with equal probability. A node can have either a short or a long range corresponding to low and high transmission power levels, respectively. The fraction of low power nodes is the variable parameter. A similar model was used in [38, 26]. This model represents a heterogeneous network with two widely different radio capabilities. For example, transmission power levels of vehicular and man-pack radios in a battlefield scenario can differ by as much as 10dB.

RandomPower With this model, each node is assigned a random transmission range that is uniformly distributed between minimum and maximum range values. This model is representative of two practical scenarios where unidirectional links might occur: (i) a generalization of TwoPower level model described above, i.e., a network of nodes with multiple different power levels; (ii) a snapshot of a network in which each node adjusts its transmit power based on the available energy supply to conserve its battery power.

Rodopl and Meng (R&M) This model is based on the distributed topology control algorithm proposed by Rodopl and Meng [37]. Topology control algorithms (e.g., [13, 37, 34, 40]) adjust node transmit powers in order to obtain a topology that optimizes a certain objective such as network capacity, network reliability or network lifetime. Almost all topology control algorithms in the literature try to guarantee some form of network connectivity while optimizing one or all of the above criteria. We have chosen the R&M algorithm because it is the *only* algorithm in our knowledge that considers unidirectional links and ensures strong connectivity possibly using some unidirectional links. This feature of the algorithm provides a favorable case for the use of unidirectional links for routing to potentially provide better network connectivity. All other algorithms guarantee connectivity using bidirectional links alone and thus are not good candidates for our evaluation.

Here we briefly review the R&M algorithm for the benefit of the readers. The algorithm aims at achieving energy efficiency through transmit power adjustment. Because of the nature of wireless communication, it is sometimes energy efficient for a node to use a lower transmit power and communicate with a farther node using intermediate relaying nodes than to use higher power and communicate directly. The algorithm uses this observation to advantage. Central to the algorithm is the notion of an *enclosure*. Enclosure of a node represents its immediate locality. As long as each node maintains links with nodes in its enclosure, strong connectivity is guaranteed. So each node reduces its transmit power from the maximum value to a level where it can reach only nodes in its enclosure. The algorithm assumes that each node knows its position. Every node computes its enclosure set by exchanging position information with all reachable nodes (using maximum power).

2.3 Evaluation Methodology

Our goal here is to assess the benefit attainable from using unidirectional links for routing in multihop wireless networks. To ac-

complish this goal, we evaluate the idealized routing performance of two approaches: (i) utilizing unidirectional as well as bidirectional links; (ii) utilizing only bidirectional links. Our evaluation process involves static simulation of large number (over a thousand for each data point) of random multihop network topologies containing unidirectional links, and comparing the average connectivity and path cost metrics with and without unidirectional links. The three models for transmit power assignment described in the previous subsection are used to generate networks with unidirectional links.

To measure connectivity, we compute the average number of strongly connected components and largest strongly connected component over all random graph samples. For comparing the path quality, we consider the average shortest path cost, per-hop acknowledgment cost and the total communication cost. Note that all path costs are in number of hops and are averaged over all node pairs having a bidirectional path between them, for each random graph sample — averaged over all such samples. The per-hop acknowledgment cost is computed as the average cost of traversing a shortest path between a pair of nodes hop-by-hop in the reverse direction. The total communication cost is simply the sum of the shortest path and acknowledgment costs.

We experiment with a wide variety of node densities and radio ranges. All our experiments are for 100 node networks. Each random network topology consists of nodes randomly placed in a square field. To vary node density (measured as nodes/sq. km), the dimensions of the field are varied. In all three range assignment models, a fixed maximum transmission range of 250 m is used. In the TwoPower model, the fraction of low power nodes is varied to get different range assignments. The long range is same as the maximum range, while the short range is always set to 125 m. Note that we experimented with different short range values, but the results are not very sensitive to these values. In the RandomPower model, the minimum range is changed for variation in ranges. In the R&M model, the radio range of a node is controlled by the algorithm, and cannot be artificially varied.

2.4 Simulation Results

2.4.1 Variation in Node Density

Here we study the effect of node density on connectivity and path cost metrics in all three range assignment models. The fraction of low power nodes in the TwoPower model is set to 0.5 as this value results in the most number of unidirectional links. In the RandomPower model, the minimum range value is kept constant at 125 m which is chosen somewhat arbitrarily. We did experiments with other values of minimum range, but the results did not vary much. Node density is varied so that all network configurations are covered starting from very sparse and disconnected networks to highly dense and connected networks. Note that number of unidirectional and bidirectional links (data not shown) increase with increase in density in all three models. However, the relative number of these links and their rate of increase with density is very much dependent on the specific model used. Also, we noticed that the mean radio range in the R&M model shrank as nodes become denser. This is expected, however, given the nature of the underlying algorithm.

The first set of plots (Figure 2) study the network connectivity properties with and without using unidirectional links. The number of strongly connected components and the size of the largest components are very similar regardless of whether or not unidirectional links are used. Note that unidirectional links do not improve connectivity in the R&M model even though they are explicitly taken into account by the algorithm. Furthermore, we found that connec-

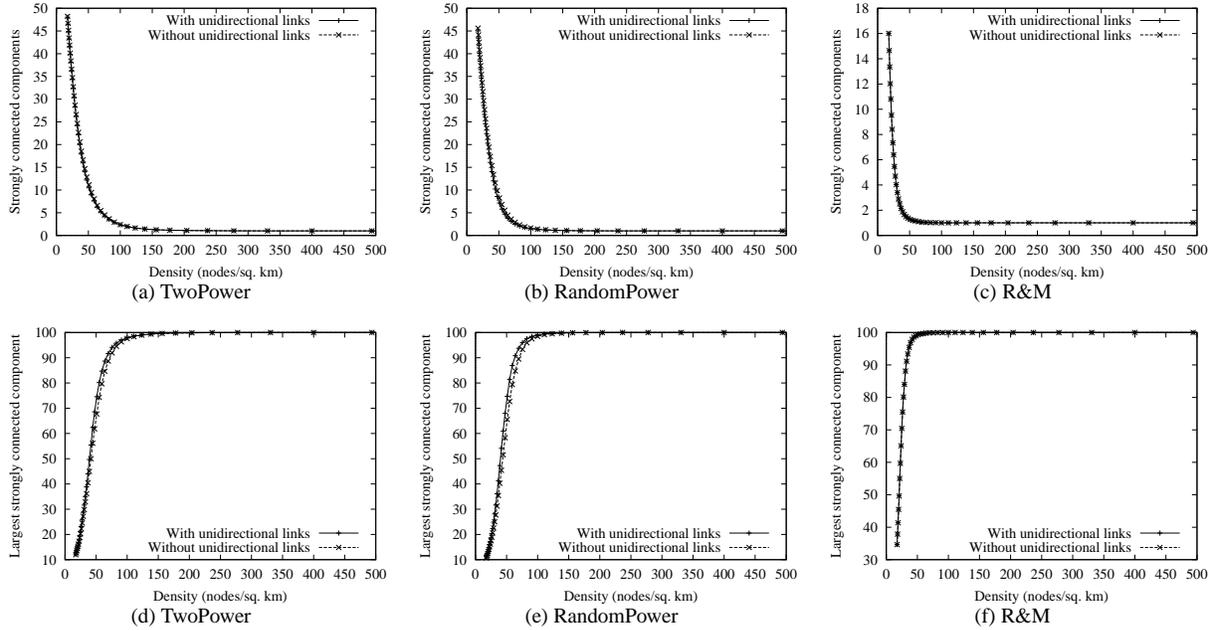


Figure 2: Connectivity metrics in all three models with varying density.

tivity metrics in this model are exactly identical to the case where all nodes use the maximum range. Both these observations suggest that it may be somewhat unlikely in random topologies for two sub-components to be connected by two unidirectional links (between different node pairs).

The second set of plots (Figure 3 (a, b, c)) shows the average cost of the shortest path. The initial hump in the plots is because of the sharp transition from disconnected to connected networks within a small range of densities. Observe that ignoring unidirectional links only marginally increases the shortest path cost in TwoPower and RandomPower models (Figure 3 (a, b)) except when the density is between 50–100 nodes/sq. km, where the increase is more significant. In the R&M model (Figure 3 (c)), the increase is marginal for lower density, but increases with increasing density.

However, note that the shortest path cost is only a part of the overall picture. Many ad hoc network protocols use some sort of per-hop acknowledgment either in the network or the link layer to guarantee reliable transmission and also to detect link breaks. Use of unidirectional links will cause such acknowledgments to traverse multiple hops – possibly in the network layer (see [23] for an idea based on tunneling). This will increase the overall communication cost. As expected, the hop-by-hop acknowledgement costs are more in all three models when unidirectional links are used (Figure 3 (d, e, f)). The overall communication cost in TwoPower and RandomPower models (Figure 3 (g, h)) is approximately the same with or without unidirectional links. In the R&M model (Figure 3 (i)), they are still similar for lower density, but the use of unidirectional links brings down the cost a bit (up to 10%) when the node density is very high.

2.4.2 Variation in Radio Range

Till now, in TwoPower and RandomPower models, the variability in range has been fixed and node density has been varied. Here we study the effect of variation in ranges for a fixed density. We set the node density to 100 nodes/sq. km which yields a connected net-

work when all nodes use the maximum range. Using this density value allows us to meaningfully evaluate the connectivity advantage from unidirectional links. A higher value of density will produce more number of bidirectional links and thus benefits the case without unidirectional links. On the other hand, a lower density value will not allow us to explore the whole range of connectivities, as network will not get connected for any range assignment.

Figure 4 shows all metrics with varying fraction of low power nodes in the TwoPower model. Note that connectivity (Figure 4 (a, b)) improves only slightly (less than a few percents) by using unidirectional links. On the other hand, the average total communication cost (Figure 4 (e)) improves (up to about 7%, but mostly lower) when a large fraction of nodes is low power; the costs are similar when a small fraction of nodes is low power.

The effect of variability in node ranges in the RandomPower model is shown in Figure 5 for different values of the minimum range. There is some noticeable improvement (about 15%) in the largest components (Figure 5 (b)) with unidirectional links when minimum range is very small. However, for higher values of minimum range, the improvements start to drop. This is somewhat expected because the variability in node ranges decreases with increase in the minimum range. Similar observation applies for communication cost (Figure 5 (e)) as well; there is up to 10% improvement with unidirectional links when the minimum range is small.

The general observation from the foregoing evaluations is that unidirectional links provide only incremental benefit. They do not improve connectivity in most cases. They do improve shortest path cost in general. But with per-hop acknowledgments, the overall benefit is small and is restricted to only certain densities and radio ranges.

3. ELIMINATION OF UNIDIRECTIONAL LINKS FROM ROUTE COMPUTATIONS

Majority of the protocols developed for multihop wireless networks assume bidirectional links (e.g., [17], DSDV [27], AODV [28],

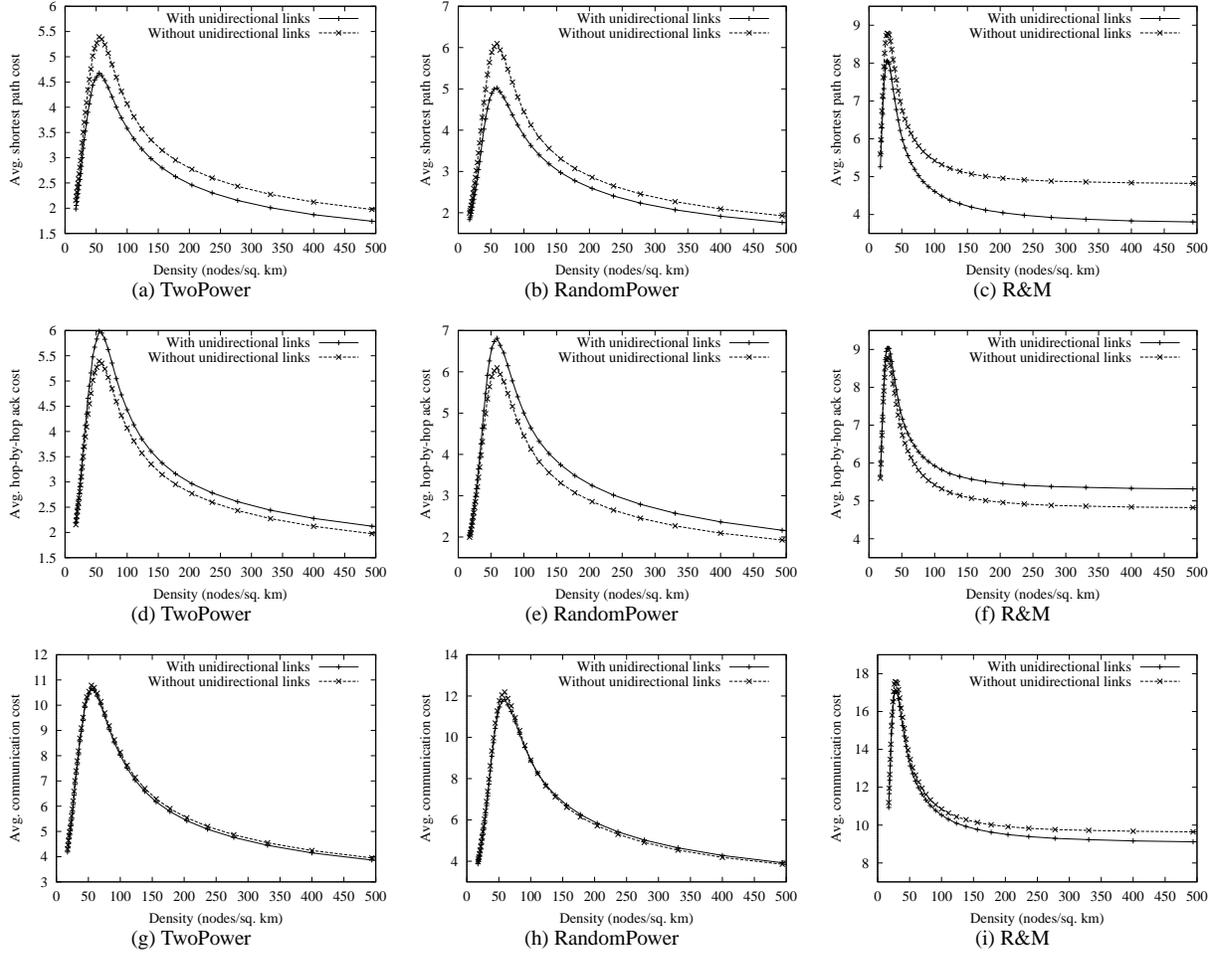


Figure 3: Path cost metrics in all three models with varying density.

TBRPF [24], OLSR [5]). But for correct operation in the presence of unidirectional links, they require additional mechanisms to eliminate unidirectional links from route computations. Our goal in this section is to understand the importance of such mechanisms and their effect on the overall performance of the routing protocol. We investigate this issue using AODV as a case study. While general comments are difficult to make, we do believe that other protocols will also benefit from the mechanisms we develop and our observations from the performance evaluation.

3.1 AODV Protocol

AODV is an on-demand routing protocol. It is loosely based on the distance-vector concept. In on-demand protocols, nodes obtain routes on an as needed basis via a route discovery procedure. Route discovery works as follows. Whenever a traffic source needs a route to a destination, it initiates a route discovery by flooding a route request (RREQ) for the destination in the network and then waits for a route reply (RREP). When an intermediate node receives the first copy of a RREQ packet, it sets up a reverse path to the source using the previous hop of the RREQ as the next hop on the reverse path. In addition, if there is a valid route available for the destination, it unicasts a RREP back to the source via the reverse path; otherwise, it re-broadcasts the RREQ packet. Duplicate copies of the RREQ are immediately discarded upon reception at every node.

The destination on receiving the first copy of a RREQ packet forms a reverse path in the same way as the intermediate nodes; it also unicasts a RREP back to the source along the reverse path. As the RREP proceeds towards the source, it establishes a forward path to the destination at each hop. AODV also includes mechanisms for erasing broken routes following a link failure, and for expiring old and unused routes. We do not discuss them, as they are not relevant here.

The above route discovery procedure requires bidirectional links for correct operation. Only then RREP can traverse back to the source along a reverse path and form a forward path to the destination at the source. Many common MAC protocols check link bidirectionality only for unicast transmissions. For example, IEEE 802.11 DCF MAC [14] protocol uses an RTS-CTS-Data-ACK exchange for unicast transmissions; receipt of CTS following an RTS or ACK following the data transmission on a link ensures that it is bidirectional. Broadcast transmissions, however, cannot detect the presence of unidirectional links. Since AODV RREQ packets typically use link-layer broadcast transmissions, some unidirectional links can go undetected and as a result reverse paths may contain unidirectional links (directed away from the source). RREP transmissions along such reverse paths will fail, as they are unicast.

Route discovery fails when none of the RREPs reach the source.

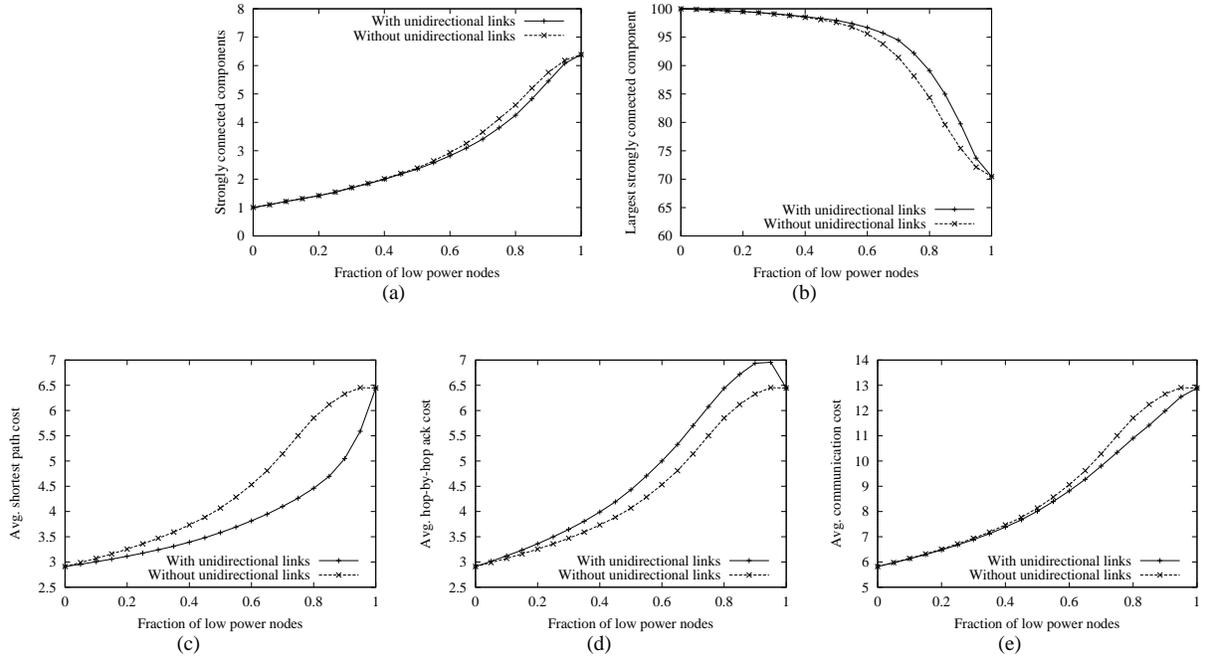


Figure 4: TwoPower model: connectivity and path cost metrics with varying fraction of low power nodes.

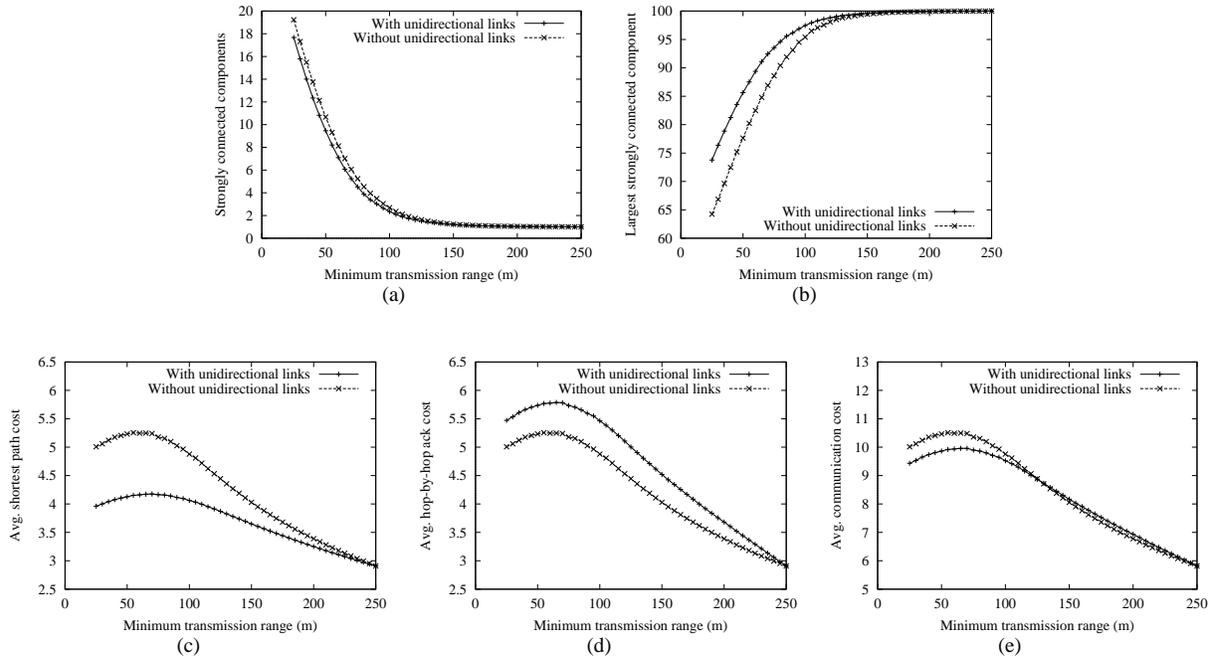


Figure 5: RandomPower model: connectivity and path cost metrics with varying values of minimum range.

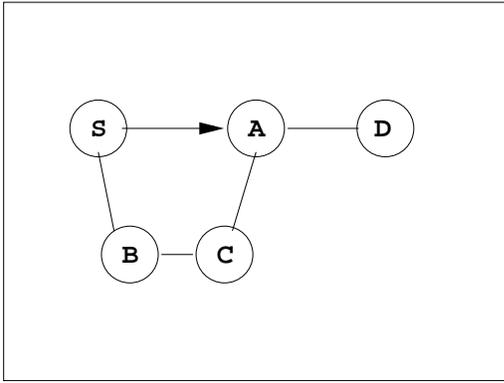


Figure 6: All links are bidirectional except the one with an arrow. *A* receives first copy of RREQ from *S* for *D* via *S – A* path, and forms a reverse path *A – S*; subsequent RREP transmission from *A* to *S* will fail. This scenario will repeat for later route discovery attempts from *S*. The alternate, longer path (*S – B – C – A – D*) will never be discovered.

It can fail even when there is a bidirectional path between the source and the destination. This is because only the first copy of a RREQ packet — which may arrive via a unidirectional path from the source — is considered by intermediate nodes and the destination to form reverse paths and send back RREPs; later copies are simply discarded even if they take bidirectional paths. See Figure 6 for an illustration.

In the worst case, this scenario will result in repeated route discovery failures. Thus additional mechanisms are needed to avoid the above problem in networks with unidirectional links.

3.2 Techniques for Handling Unidirectional Links in AODV

In the following we describe three techniques to alleviate this problem. The first two techniques — “BlackListing” and “Hello” — are known techniques. The third technique, “ReversePathSearch” is our contribution in this paper.

BlackListing This technique reactively eliminates unidirectional links. It is included in the latest AODV specification [29]. Here, whenever a node detects a RREP transmission failure, it inserts the next hop of the failed RREP into a “blacklist” set. The blacklist set at a node indicates the set of nodes from which it has unidirectional links. For example, in Figure 6 node *A* will blacklist node *S*. Later when a node receives a RREQ from one of the nodes in its blacklist set, it discards the RREQ to avoid forming a reverse path with unidirectional link. This gives a chance for RREQ from an alternate path (e.g., via *C* in Figure 6) to provide a different reverse path. BLACKLIST_TIMEOUT specifies the period for which a node remains in the blacklist set. By default, this period is set to the upper bound of the time it takes to perform the maximum allowed number of route discovery attempts by a source.

This technique is simple and has little overhead when there are few unidirectional links. However, when there are many unidirectional links, this approach is inefficient because these links are blacklisted iteratively one at a time. Several route discoveries may be needed before a bidirectional path, if exists, is found. Another difficulty with this technique is in setting an appropriate value for the BLACKLIST_TIMEOUT.

Setting it to a small value may reduce the effectiveness of the technique. On the other hand, setting it to a very large value affects connectivity when there are many short-term unidirectional links.

Hello In the contrast to the BlackListing technique, this technique proactively eliminates unidirectional links by using periodic one-hop Hello packets. A similar idea has also been used in OLSR [5] to record only bidirectional links. In each Hello packet, a node includes all nodes from which it can hear Hellos (i.e. its set of neighbors). If a node does not find itself in the Hello packet from another node, it marks the link from that node as unidirectional. Just as in the BlackListing technique, every node ignores RREQ packets that come via such unidirectional links. Note that this hello packets are identical to the AODV hello packets [29] except for the additional neighborhood information.

The advantage of this technique is that it automatically detects unidirectional links by exchanging Hello packets. But the periodic, large Hello packets can be a significant overhead. Although the size of the Hello packets may be reduced by using “differential” Hellos [24], the periodic packet overhead is still a concern. However, in situations when Hellos must be used for maintaining local neighborhood and to detect link failures (e.g., when the link layer cannot provide any feedback about link failures), incremental overhead for unidirectional link detection may not be very much.

ReversePathSearch Unlike the above two techniques, this technique does not explicitly remove unidirectional links. Instead, it takes a completely different approach. Each unidirectional link is viewed as a “fault” in the network and multiple paths between the nodes are discovered to perform fault-tolerant routing. The basic idea is as follows. During the RREQ flood, multiple loop-free reverse paths to the source are formed at intermediate nodes and the destination. Using a distributed search procedure, multiple RREPs explore this multipath routing structure in an attempt to find one or more bidirectional paths between the source and the destination. This search procedure is somewhat similar to the well-known depth first search algorithm. When RREP fails at a node, the corresponding reverse path is erased and the RREP is retried along an alternate reverse path, if one is available; when all reverse paths fail at a node during this process, the search backtracks to upstream nodes² of that node with respect to the source and they too follow the same procedure. This continues until either one or more bidirectional paths are found at the source, or all reverse paths are explored. This technique is described in more detail in the following subsection.

3.3 Reverse Path Search Technique

In a prior work [19], we have investigated a multipath extension to AODV, called AOMDV, where route update rules to maintain multiple loop-free paths are described. We use the same AOMDV route update rules here in the ReversePathSearch technique to maintain multiple loop-free paths to a destination at every node.

In the ReversePathSearch algorithm, *all* RREQ copies including duplicates are examined at intermediate nodes and the destination for possible alternate reverse paths to the source. However, reverse

²A node *X* is upstream of a node *Y* with respect to a node *D* if *Y* appears on a path from *X* to *D*. Conversely, *Y* is a downstream node of *X* for *D*.

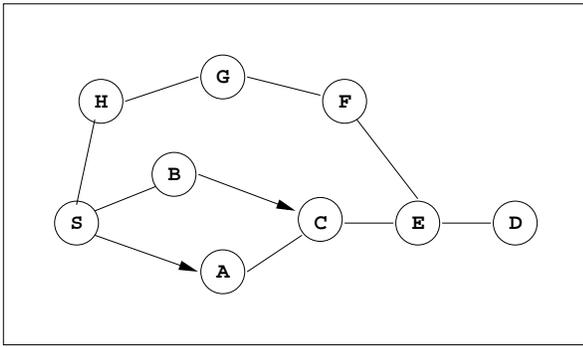


Figure 7: Demonstration of reverse path search. Multiple reverse paths to the source S formed during the RREQ flood are shown; some of the reverse paths contain unidirectional links such as $C - A - S$. Consider the RREP propagation via $C - A - S$. The transmission from A to S fails causing a BRREP transmission at A . C erases the reverse path via A and transmits a RREP to B in order to explore the reverse path $C - B - S$. This also fails, causing C to transmit BRREP. E then erases the reverse path via C and transmits RREP to F in order to explore path $E - F - G - H - S$. This will be successful.

paths are formed only from those copies that satisfy route update rules and provide loop-free paths [19]. Other copies are simply discarded. Note that this is different from basic AODV where only the first copy is looked at.

When a RREQ copy at an intermediate node creates or updates a reverse path, and the intermediate node has no valid path to the destination, the RREQ copy is re-broadcasted provided it is the first copy that yields a reverse path; this is somewhat similar to basic AODV where only first copies of RREQ are forwarded to prevent looping during the flood. On the contrary, if the intermediate node does have a valid path to the destination, it checks whether or not a RREP has already been sent for this route discovery. If not, it sends back a RREP along the newly formed reverse path and remembers the next hop used for this RREP; otherwise, the RREQ copy is dropped.

When the destination receives copies of RREQs, it also forms reverse paths in the same way as intermediate nodes. However, unlike intermediate nodes, the destination sends back a RREP along each new reverse path. Multiple replies from destination allow exploration of multiple reverse paths concurrently — thus speeding up the search for a bidirectional path. In contrast, allowing multiple destination replies in basic AODV has little benefit unless those replies take non-overlapping paths to the source. This is because intermediate nodes have at most one reverse path back to the source.

When an intermediate node receives a RREP, it follows route update rules in [19] to form a loop-free forward path to the destination, if possible; else, the RREP is dropped. Supposing that the intermediate node forms the forward path and has one or more valid reverse paths to the source, it checks if any of those reverse paths was previously used to send a RREP for this route discovery. If not, it chooses one of those reverse paths to forward the current RREP, and also remembers the next hop for that reverse path; otherwise, the RREP is simply dropped.

RREP transmission failure (as a result of transmission over a unidirectional link, for example) at an intermediate node results in that node erasing the corresponding reverse path, and retrying an alternate reverse path. If no such alternate path is available, the in-

termediate node transmits (broadcast transmission) a new message called the “backtrack route reply” (BRREP) to inform its upstream nodes (with respect to the source) to try other reverse paths at those nodes. A BRREP is also generated by an intermediate node, if that node does not have any reverse path upon a RREP reception. On receiving a BRREP, an intermediate node upstream of the BRREP source (meaning it has last sent a RREP to the BRREP source for this route discovery) takes a similar action as on a RREP failure; nodes that are not upstream of the BRREP source simply discard the packet on reception. When the destination encounters a RREP failure, or receives a BRREP, it only erases corresponding reverse paths. See Figure 7 for an illustration.

Note that the above procedure is guaranteed to terminate. To see this, observe that every RREP failure erases the corresponding reverse path. So reverse paths cannot be explored indefinitely since there are only finite number of them. On the other hand, alternate reverse paths are not explored at the intermediate nodes as long as RREPs successfully go through. Also note that in the above description, some details have been omitted for the sake of brevity. For instance, algorithms actions to cope with BRREP loss are not mentioned.

Multiple loop-free reverse paths used by the above algorithm are in general a subset of all possible reverse paths. Thus, sometimes it is possible that the multiple reverse paths explored by the algorithm do not include a bidirectional path between source and destination although such a path exists. But in dense networks that we consider, often there is more than one bidirectional path and the above possibility is rare.

Finally, multiple replies from the destination in the algorithm yield multiple forward paths at intermediate nodes and the source. This ability to compute multiple bidirectional paths in a single route discovery is highly beneficial in mobile networks for efficient recovery from route breaks.

3.4 Performance Evaluation

In this section we evaluate the performance of the three techniques described in the previous subsection relative to basic AODV under varying number of unidirectional links and node speeds. Two primary goals from this evaluation are: (i) to understand the impact of unidirectional links on basic AODV performance; and (ii) to evaluate the effectiveness of the three techniques in handling unidirectional links.

3.4.1 Simulation Environment

We use a detailed simulation model based on *ns-2* [8]. The Monarch research group in CMU developed support for simulating multi-hop wireless networks complete with physical, data link and MAC layer models [4] on *ns-2*. The distributed coordination function (DCF) of IEEE 802.11 [14] for wireless LANs is used as the MAC layer. The radio model uses characteristics similar to a commercial radio interface, Lucent’s WaveLAN. WaveLAN is a shared-media radio with a nominal bit-rate of 2 Mb/s and a nominal radio range of 250 meters. More details about the simulator can be found in [4, 8]. This simulator has been used for evaluating performance of earlier versions of the AODV protocol (e.g., [4, 30]).

The AODV model in our simulations is based on the latest protocol specification [29], except that the expanding ring search is disabled for all protocol variations. Note that the expanding ring search introduces an additional reason for route discovery failures, i.e., when a smaller ring size (TTL value) is used for the search. This makes analysis somewhat difficult. So in our AODV model, a source does network-wide route discoveries appropriately spaced

in time until either a route is obtained or a maximum retry limit (3) is reached; when the limit is exceeded, source reports to the application that the destination is unreachable and drops all buffered packets for the destination — we term this event as “route search failure” in our evaluations. Besides, link layer feedback is used to detect link failures in all protocol variations. The 802.11 MAC layer reports a link failure when it fails to receive CTS after several RTS attempts, or to receive ACK after several retransmissions of DATA. Note that in the Hello technique, link failures are detected using hello messages as well as the feedback, whichever detects the link breakage first.

TwoPower model is used to create unidirectional links, primarily because it does not perform power control. Use of power control makes the analysis here difficult as the number of unidirectional links then will become heavily dependent on the choice of the actual power control algorithm and the frequency at which the algorithm is invoked. Note that the frequency of invocation did not play any role in our earlier evaluations as we used only static topologies. We modified the ns-2 simulator to allow variable transmission ranges for nodes. In all our experiments, short and long ranges in the TwoPower model are set to 125 m and 250 m, respectively. We vary the fraction of low power nodes to vary the number of unidirectional links.

We consider 100 node networks with nodes initially placed at random in a rectangular field of dimensions 575 m x 575 m. The field size is chosen to guarantee a connected network across the parameter space. We use random waypoint model [4] to model node movements. Pause time is always set to zero and the maximum speed of the nodes is changed to change mobility.

Traffic pattern in our experiments consists of fixed number of CBR connections (20) between randomly chosen source-destination pairs and each connection starts at a random time at the beginning of the simulation and stays until the end. Each CBR source sends packets (each of size 512 bytes) at a fixed rate of 4 packets/s.

All our experiments use 500 second simulation times. In the case of static networks, each data point in the plots is an average of at least 50 runs with different randomly generated initial node positions and range assignments in each run. In the mobility experiments, we average over 25 randomly generated mobility scenarios and range assignments. Identical scenarios are used across all protocol variations.

3.4.2 Performance Metrics

We evaluate four key performance metrics: (i) *Packet delivery fraction* — ratio of the data packets delivered to the destination to those generated by the CBR sources; (ii) *Average end-to-end delay* of data packets — this includes all possible delays caused by buffering during route discovery, queuing delay at the interface, retransmission delays at the MAC, propagation and transfer times; (iii) *Route search failures* — total number of route search failure events (see previous subsection) at sources; (iv) *Normalized routing load* — the number of routing packets “transmitted” per data packet “delivered” at the destination. Each hop-wise transmission of a routing packet is counted as one transmission.

3.4.3 Simulation Results

We present two sets of experiments. In the first set, the network is static and the number of unidirectional links is varied. In the second set, node mobility is considered. Figure 8 shows the packet delivery fraction, average delay, and the route search failures as a function of the number of unidirectional links. We change the fraction of low power nodes from 0 to 0.5 to increase the number of unidirectional links. With increase in unidirectional links, basic

AODV drops the highest number of packets (as many as 20%) and also experiences most number of route search failures (Figure 8 (a, c)). This is because the basic AODV protocol does not take notice of the unidirectional links and repeatedly performs route discoveries without any benefit. Note that after every route search failure, all packets buffered for the destination at the source are dropped.

The drop in packet delivery is less drastic for BlackListing compared to basic AODV. But it still drops as many as 14% of the packets because of its slowness in eliminating unidirectional links one by one. It still has a large number of route search failures; but performs somewhat better than basic AODV. The delay performance of AODV and BlackListing (Figure 8 (b)) is similar as route discovery latency dominates the delay in both cases.

Both Hello and ReversePathSearch deliver almost all packets always (Figure 8 (a)), as both are able to successfully find routes always (Figure 8 (c)). However, their delay performance (Figure 8 (b)) is quite different. Delay for the Hello technique is similar to basic AODV and BlackListing, while ReversePathSearch has significantly lower delay than others. This shows that ReversePathSearch can effectively overcome unidirectional links by exploring multiple reverse paths. However, the delay performance of the Hello technique is counter-intuitive. One would normally expect the performance of Hello to be independent of unidirectional links because it proactively eliminates unidirectional links in the background without burdening AODV route discovery mechanism. Below we will analyze the reason behind this unexpected behavior.

By additional instrumentation, we found that the sharp increase in route discovery attempts with increase in unidirectional links (Figure 9 (a)) explains the delay performance of the Hello technique. What is more interesting is the reason behind the rise in route discoveries itself. This we found was because of an undesirable interaction of this protocol with the 802.11 MAC layer. We noticed that in general MAC collisions increased with increase in the number of unidirectional links. This is because of the hidden terminal interference via unidirectional links, and the insufficiency of the RTS-CTS handshake in 802.11 MAC to avoid such hidden terminals (See [32] for a similar observation). For the Hello technique, however, the increase in collisions is quite dramatic (Figure 9 (c)) because of large and periodic (every second) broadcast hello packets. Consequently, the efficiency of the MAC layer is negatively affected resulting in more number of unsuccessful transmissions and triggering of link failure signals to the routing protocol. Such link failure signals cause route breaks (Figure 9 (b)). Note that in the basic AODV protocol, every route break will result in a new route discovery attempt. Since the Hello technique is identical to basic AODV except for the additional hello exchanges, it does more number of route discovery attempts as unidirectional links grow in number. Note that basic AODV and BlackListing are not affected very much by the above phenomenon, as they spend most of their effort finding a route. Even though ReversePathSearch is subject to this problem to some extent, the availability of redundant forward paths prevents a big drop in its performance.

Figure 10 shows the routing load comparison for the four protocols. Overall, ReversePathSearch has the lowest overhead followed by BlackListing, basic AODV and Hello. The high overhead of the Hello technique is expected because of the periodic hello messages. Also, low routing overhead in ReversePathSearch indicates that the higher per route discovery costs in ReversePathSearch due to more (backtrack) route replies is very well offset by the significant reduction in route discoveries (Figure 9(a)). Relative performance in terms of the byte overhead (not shown) is same as above. In summary, the ReversePathSearch technique allows for a much more effective elimination of unidirectional links from route computations

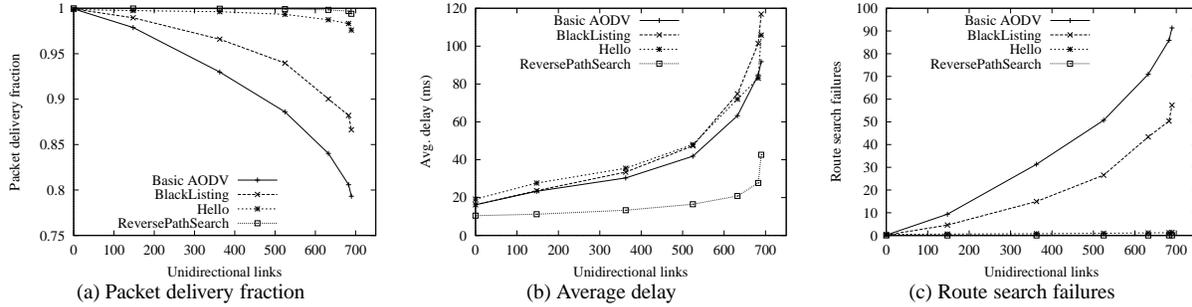


Figure 8: Performance with varying number of unidirectional links.

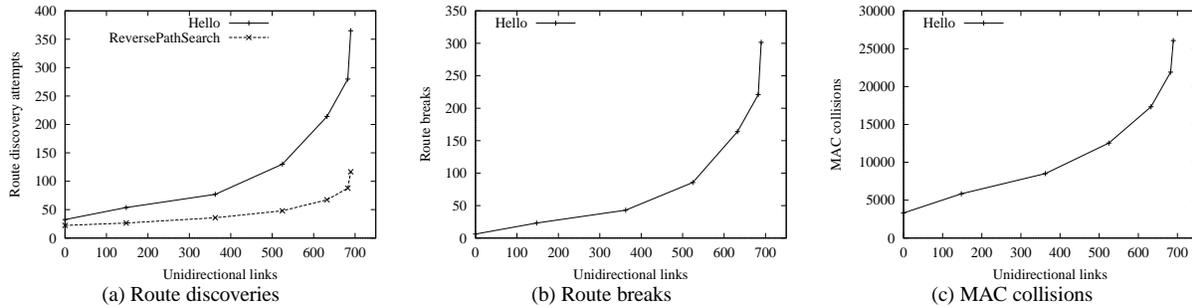


Figure 9: Route discoveries, route breaks and MAC collisions for the Hello technique with varying number of unidirectional links.

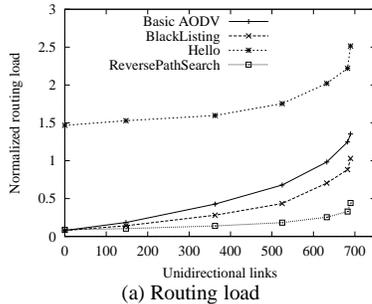


Figure 10: Routing load with varying number of unidirectional links.

compared to BlackListing, however with much lower overhead cost compared to Hello.

The effect of node mobility on all metrics is shown in Figure 11. Here we vary the maximum node speed between 0 and 20 m/s. To stress the protocols, we set the fraction of low power nodes to 0.5 which results in the most number of unidirectional links in static networks. Mobility also affects the number and duration of unidirectional links. But these are somewhat hard to quantify in mobile networks. As expected, performance in terms of packet delivery, delay and overhead degrades for all protocols with increase in mobility (Figure 11 (a, b, d)). However, observations about the relative performance of the four protocols pretty much remain same as in the static network case. Note that the difference in routing load between basic AODV, BlackListing, and Hello shrinks as mobility

starts to play a dominant role. ReversePathSearch, however, continues to perform significantly better than the rest. With the built-in redundancy in the route discovery process, it not only eliminates unidirectional links from the route computations by exploring alternate reverse paths, but also, in a similar vain, avoids broken reverse paths due to mobility (Figure 11 (c)). In addition, by computing multiple forward paths at the source and intermediate nodes, it obviates the need for frequent route discovery attempts in response to route breaks caused by node mobility.

4. RELATED WORK

Although many common routing protocols assume bidirectional links, there is still considerable amount of literature available on routing using unidirectional links (e.g., [1, 10, 39, 31, 2]). These protocols are mainly targeted towards two network environments, namely, mixed satellite and terrestrial networks, and multihop wireless networks, where unidirectional links commonly occur. Of the several unicast routing proposals for multihop wireless networks within the IETF MANET working group [18], only DSR [15] and FSR [9] can fully support unidirectional links, while ZRP [12, 11] and TORA [25] can partially support unidirectional links. There have also been attempts to extend existing protocols to support unidirectional links [33, 38, 16]. But none of the above efforts contain any simulation or experimental evaluation of the impact of unidirectional links on routing performance in realistic scenarios.

Support for unidirectional links below the network layer also received some attention. Link-layer tunneling approach has been explored in [7, 23]. The main motivation behind this approach is to hide the unidirectional nature of a link from higher layer protocols so that they can operate over unidirectional links without any modifications. This is basically achieved by forming a reverse

tunnel (possibly via a multihop path) for each unidirectional link using information gathered by the routing protocol. In [36], an alternative approach to tunneling, but with a similar goal is proposed. The idea here is to introduce a sub-layer beneath the network layer to find and maintain multihop reverse routes for each unidirectional link. There is also some work on using multihop acknowledgements to discover unidirectional links [26], and GPS-based approaches for enabling link-level acknowledgements [16] over unidirectional links.

Work on channel access protocols for multihop wireless networks with unidirectional links is starting to get attention [35, 3]. Ramanathan [35] makes an important observation that many unidirectional links hurt channel access protocol performance. This is somewhat related to our observation that utilizing unidirectional links does not provide any significant additional benefit.

5. CONCLUSIONS

Unidirectional links commonly occur in wireless ad hoc networks because of the differences in node transceiver capabilities or perceived interference levels. Unidirectional links can presumably benefit routing by providing improved network connectivity and shorter paths. But prior work indicates that routing over unidirectional links usually causes high overheads. With this in mind, we evaluated performance advantages, in terms of connectivity and path costs, of routing using unidirectional links under ideal conditions. Our evaluations were done with respect to three variable transmission range assignment models that reflect some realistic network scenarios with unidirectional links. Main conclusion from this study is that unidirectional links provide only incremental benefit. Thus, protocols that avoid unidirectional links demand a closer look.

Many common routing protocols, however, simply assume that links are bidirectional. They will require some additional protocol operations to detect and eliminate unidirectional links from route computations. To assess the difficulty of doing this, we have presented a case study with the AODV protocol where three such techniques are presented and evaluated. It is observed that the ReversePathSearch technique performs the best because of its ability to explore multiple paths. It exhibits a dual advantage, both in terms of immunity from unidirectional links and from mobility-induced link failures. While this case study has been performed only for AODV, we expect that other protocols that share certain characteristics with AODV, such as the on-demand nature or the distance vector framework, will also benefit from these ideas.

Besides, our performance study also revealed that 802.11 MAC performance degrades in the presence of unidirectional links. A similar observation was also made in [32]. These observations suggest the need for more efficient MAC protocols to handle unidirectional links, as such links may be inevitable in certain ad hoc network scenarios (for example, a network of nodes with heterogeneous powers).

Acknowledgments

This work is partially supported by NSF CAREER grant ACI-00961-86 and NSF networking research grant ANI-0096264. Mahesh Marina is supported by an OBR computing research award in the ECECS department, University of Cincinnati.

6. REFERENCES

[1] Y. Afek and E. Gafni. Distributed Algorithms for Unidirectional Networks. *SIAM Journal of Computing*, 23(6):1152–1178, 1994.

[2] L. Bao and J. J. Garcia-Luna-Aceves. Link-state Routing in Networks with Unidirectional Links. In *Proceedings of International Conference on Computer Communications and Networks (IC3N)*, pages 358–363, 1999.

[3] L. Bao and J. J. Garcia-Luna-Aceves. Channel Access Scheduling in Ad Hoc Networks with Unidirectional Links. In *Proceedings of Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIALM)*, pages 9–18, 2001.

[4] J. Broch, D. Maltz, D. Johnson, Y.-C. Hu, and J. Jetcheva. A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols. In *Proceedings of IEEE/ACM MOBICOM*, pages 85–97, 1998.

[5] T. Clausen et al. Optimized Link State Routing Protocol. <http://www.ietf.org/internet-drafts/draft-ietf-manet-olsr-06.txt>, Mar 2002. IETF Internet Draft (work in progress).

[6] S. R. Das, R. Castaneda, and J. Yan. Simulation-based Performance Evaluation of Routing Protocols for Mobile Ad hoc Networks. *ACM/Baltzer Mobile Networks and Applications (MONET)*, 5(3):179–189, 2000.

[7] E. Duros, W. Dabbous, H. Izumiyama, N. Fujii, and Y. Zhang. A Link-Layer Tunneling Mechanism for Unidirectional Links. RFC 3077, 2001.

[8] K. Fall and K. Varadhan(Eds.). *The ns Manual*. <http://www.isi.edu/nsnam/ns/ns-documentation.html>, 2002.

[9] M. Gerla, X. Hong, and G. Pei. Fisheye State Routing Protocol (FSR) for Ad Hoc Networks. <http://www.ietf.org/internet-drafts/draft-ietf-manet-fsr-02.txt>, Dec 2001. IETF Internet Draft (work in progress).

[10] M. Gerla, L. Kleinrock, and Y. Afek. A Distributed Routing Algorithm for Unidirectional Networks. In *Proceedings of IEEE GLOBECOM*, 1983.

[11] Z. J. Haas, M. R. Pearlman, and P. Samar. The Interzone Routing Protocol (IERP) for Ad Hoc Networks. <http://www.ietf.org/internet-drafts/draft-ietf-manet-zone-ierp-01.txt>, June 2001. IETF Internet Draft (work in progress).

[12] Z. J. Haas, M. R. Pearlman, and P. Samar. The Intrazone Routing Protocol (IARP) for Ad Hoc Networks. <http://www.ietf.org/internet-drafts/draft-ietf-manet-zone-iarp-01.txt>, June 2001. IETF Internet Draft (work in progress).

[13] L. Hu. Topology Control for Multihop Packet Radio Networks. *IEEE Transactions on Communications*, 41(10):1474–1481, 1993.

[14] IEEE Standards Department. Wireless LAN medium access control (MAC) and physical layer (PHY) specifications, IEEE standard 802.11–1997.

[15] D. B. Johnson, D. A. Maltz, Y. Hu, and J. G. Jetcheva. The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks (DSR). <http://www.ietf.org/internet-drafts/draft-ietf-manet-dsr-07.txt>, Feb 2002. IETF Internet Draft (work in progress).

[16] D. Kim, C. K. Toh, and Y. Choi. On supporting Link Asymmetry in Mobile Ad Hoc Networks. In *Proceedings of IEEE GLOBECOM*, pages 2798–2803, 2001.

[17] G. Lauer. Packet-Radio Routing. In M. Steenstrup, editor, *Routing in Communication Networks*, chapter 11. Prentice Hall, 1995.

[18] J. Macker and S. Corson. Mobile Ad hoc Networks (MANET). <http://www.ietf.org/html.charters/manet-charter.html>, 1997. IETF Working Group Charter.

[19] M. K. Marina and S. R. Das. On-demand Multipath Distance Vector Routing in Ad Hoc Networks. In *Proceedings of IEEE International Conference on Network Protocols (ICNP)*, pages 14–23, 2001.

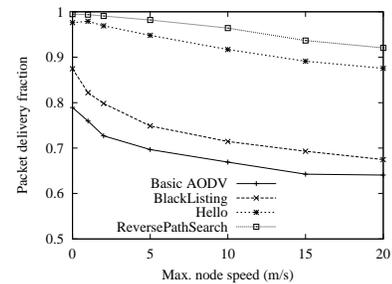
[20] J. P. Monks, V. Bharghavan, and W. W. Hwu. A Power Controlled Multiple Access Protocol for Wireless Packet Networks. In *Proceedings of IEEE INFOCOM*, pages 219–228, 2001.

[21] J. Moy. OSPF version 2. RFC 1247, 1991.

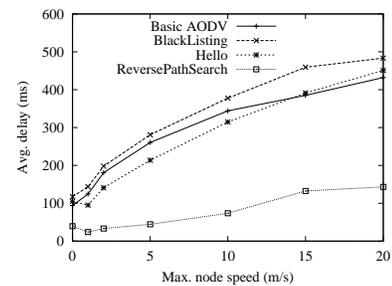
[22] S. Narayanaswamy, V. Kawadia, R. S. Sreenivas, and P. R. Kumar. Power Control in Ad-Hoc Networks: Theory, Architecture, Algorithm and Implementation of the COMPOW Protocol. In *Proceedings of European Wireless Conference*, pages 156–162, 2002.

[23] S. Nesargi and R. Prakash. A Tunneling Approach to Routing with Unidirectional Links in Mobile Ad-Hoc Networks. In *Proceedings of International Conference on Computer Communications and Networks (IC3N)*, pages 522–527, 2000.

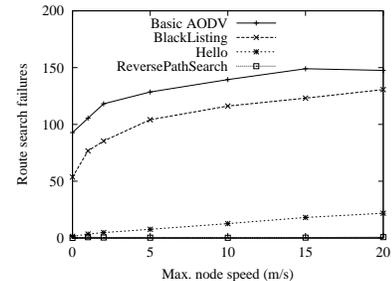
- [24] R. Ogier, F. L. Templin, B. Bellur, and M. G. Lewis. Topology Broadcast Based on Reverse-Path Forwarding (TBRPF). <http://www.ietf.org/internet-drafts/draft-ietf-manet-tbrpf-05.txt>, Mar 2002. IETF Internet Draft (work in progress).
- [25] V. Park and S. Corson. Temporally-ordered routing algorithm (TORA) version 1 functional specification. <http://www.ietf.org/internet-drafts/draft-ietf-manet-tora-spec-04.txt>, July 2001. IETF Internet Draft (work in progress).
- [26] M. R. Pearlman, Z. J. Haas, and B. P. Manvell. Using Multi-Hop Acknowledgements to Discover and Reliably Communicate over Unidirectional Links in Ad Hoc Networks. In *Proceedings of Wireless Communications and Networking Conference (WCNC)*, pages 532–537, 2000.
- [27] C. E. Perkins and P. Bhagwat. Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers. In *Proceedings of ACM SIGCOMM*, pages 234–244, 1994.
- [28] C. E. Perkins and E. M. Royer. Ad Hoc On-Demand Distance Vector Routing. In *Proceedings of IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, pages 90–100, 1999.
- [29] C. E. Perkins, E. M. Royer, and S. R. Das. Ad hoc On-Demand Distance Vector (AODV) Routing. <http://www.ietf.org/internet-drafts/draft-ietf-manet-aodv-10.txt>, Jan 2002. IETF Internet Draft (work in progress).
- [30] C. E. Perkins, E. M. Royer, S. R. Das, and M. K. Marina. Performance Comparison of Two On-demand Routing Protocols for Ad Hoc Networks. *IEEE Personal Communications*, 8(1):16–28, 2001.
- [31] C. Pomalaza-Raez. A Distributed Routing Algorithm for Multihop Packet Radio Networks with Uni- and Bi-Directional Links. *IEEE Transactions on Vehicular Technology*, 44(3):579–585, 1995.
- [32] N. Poojary, S. V. Krishnamurthy, and S. Dao. Medium Access Control in a Network of Ad Hoc Nodes with Heterogeneous Power Capabilities. In *Proceedings of IEEE ICC*, pages 872–877, 2001.
- [33] R. Prakash. A Routing Algorithm for Wireless Ad Hoc Networks with Unidirectional Links. *ACM/Kluwer Wireless Networks*, 7(6):617–625, 2001.
- [34] R. Ramanathan and R. Rosales-Hain. Topology Control of Multihop Wireless Networks using Transmit Power Adjustment. In *Proceedings of IEEE INFOCOM*, pages 404–413, 2000.
- [35] S. Ramanathan. A Unified Framework and Algorithm for Channel Assignment in Wireless Networks. *Wireless Networks*, 5(2):81–94, 1999.
- [36] V. Ramasubramanian, R. Chandra, and D. Mosse. Providing a Bidirectional Abstraction for Unidirectional Ad Hoc Networks. In *Proceedings of IEEE INFOCOM*, 2002. To appear.
- [37] V. Rodoplu and T. Meng. Minimum Energy Mobile Wireless Networks. *IEEE Journal on Selected Areas in Communications, Special Issue on Ad Hoc Networks*, 17(8):1333–1344, Aug 1999.
- [38] P. Sinha, S. V. Krishnamurthy, and S. Dao. Scalable Unidirectional Routing with Zone Routing Protocol (ZRP) Extensions for Mobile Ad-hoc Networks. In *Proceedings of Wireless Communications and Networking Conference (WCNC)*, pages 1329–1339, 2000.
- [39] G. Tel. Directed Network Protocols. In *Proceedings of International Workshop on Distributed Algorithms (WDAG)*, pages 13–29, 1987.
- [40] R. Wattenhofer, L. Li, P. Bahl, and Y. Wang. Distributed Topology Control for Power Efficient Operation in Multihop Wireless Ad Hoc Networks. In *Proceedings of IEEE INFOCOM*, pages 1388–1397, 2001.
- [41] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides. On the Construction of Energy-Efficient Broadcast and Multicast Trees in Wireless Networks. In *Proceedings of IEEE INFOCOM*, pages 585–594, 2000.



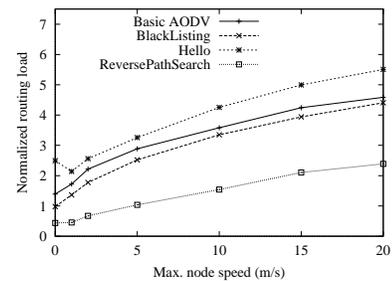
(a) Packet delivery fraction



(b) Average delay



(c) Route search failures



(d) Routing load

Figure 11: Performance with varying mobility.