

Using Directionality in Mobile Routing

Bow-Nan Cheng
ECSE Department
Rensselaer Polytechnic Institute
chengb@rpi.edu

Murat Yuksel
CSE Department
University of Nevada - Reno
yukse@cse.unr.edu

Shivkumar Kalyanaraman
ECSE Department
Rensselaer Polytechnic Institute
shivkuma@ecse.rpi.edu

Abstract—The increased usage of directional methods of communications has prompted research into leveraging directionality in every layer of the network stack. In this paper, we explore the use of directionality in layer 3 to facilitate routing in *highly mobile environments*. We introduce Mobile Orthogonal Rendezvous Routing Protocol (MORRP) for mobile ad-hoc networks (MANETs). MORRP is a lightweight, but scalable routing protocol utilizing *directional communications* (such as directional antennas or free-space-optical transceivers) to relax information requirements such as coordinate space embedding, node localization, and mobility. This relaxation is done by introducing a novel concept called the directional routing table (DRT) which maps a *set-of-IDs* to each interface *direction* to provide *probabilistic routing information based on interface direction*. We show that MORRP achieves connectivity with high probability even in highly mobile environments while maintaining only probabilistic information about destinations. Additionally, we compare MORRP with various proactive, reactive, and position-based routing protocols using single omni-directional interfaces and multiple directional interfaces and show that MORRP gains over 10-14X additional goodput vs. traditional protocols and 15-20% additional goodput vs. traditional protocols using multiple interfaces. MORRP scales well without imposing DHT-like graph structures (eg: trees, rings, torus etc). We also show that high connectivity can be achieved *without the need to frequently disseminate node position resulting increased scalability even in highly mobile environments*.

I. INTRODUCTION

A recent trend in wireless communications has been the desire to leverage directional forms of communications (e.g. directional smart antennas [11], Free-Space-Optical transceivers [13], and sector antennas) for more efficient medium reuse, increased scalability, enhanced security and potential for higher achievable bandwidth. Previous work in directional antennas focused heavily on measuring network capacity and medium reuse [11] [12]. In these works, it was shown that with proper tuning, capacity *improvements* using directional over omnidirectional antennas are dramatic - even just 8 directional interfaces results in a theoretical capacity gain of 50X.

Additionally, there has been a large push in the free space optical (FSO) community to use FSO to compliment traditional RF methods [7]. FSO has several attractive characteristics like (i) *dense spatial reuse*, (ii) *low power usage*, (iii) *license-free band of operation*, and (iv) *relatively high bandwidth compared to RF* but suffers from (i) *the need for line of sight (LOS) alignment* and (ii) *reduced transmission quality in adverse weather conditions*. Yuksel et al. [13] proposed several ways to mitigate these issues by tessellating low cost

FSO transceivers in a spherical fashion and replacing long-haul point-to-point links with short, multi-hop transmissions.

Given the seemingly large increases in medium reuse and potential for higher bandwidth in directional forms of communications, it becomes interesting to investigate how *directionality* can be used to complement and even enhance wireless networks in all layers of the stack. There are several challenges associated with using directionality in mobile networks. Unlike omnidirectional antennas where neighbor reach depends almost exclusively on range, nodes using directional antennas need also take into account the neighbor's direction and map it to a specific interface in that direction. The problem is complicated even further as nodes closer to a source seemingly incur more dynamism (even small movements can affect perceived direction dramatically) while nodes farther away incur less change. In this paper, we address these issues and propose utilizing directionality for a novel purpose: to facilitate layer 3 routing in *highly mobile environments* without the need for flooding either in the route dissemination or discovery phase. Most prior work on leveraging directional antennas in the routing layer focus on adapting routing protocols to simply *utilize* directional communications [4][5]. Our work is novel in that we utilize *local* directionality as a property *to* route packets itself. To the authors' best knowledge, this is the first paper that attempts to use directionality to address issues with *high mobility*.

Our protocol, Mobile Orthogonal Rendezvous Routing Protocol (MORRP) is based on two fundamental primitives: a) local directionality is sufficient to maintain forwarding of a packet on a straight line, and b) two sets of orthogonal lines in a plane intersect with high probability even in sparse, bounded networks. Cheng et al. [3] showed that in *static* wireless mesh networks, by forwarding packets to nodes intersected by a pair of orthogonal lines originating from a source and destination, one can successfully route packets to a high degree of connectivity (98%) without the need for coordinate space. Furthermore, it was shown that forwarding using this method state-scales to $O(N^{3/2})$ with the states spread evenly throughout the network, while incurring a path stretch vs. shortest path of only 1.2. Unfortunately, the proposed protocol fails under even slight mobility because straight-line paths and rigid "destination - next-hop" routing tables are hard to maintain in the presence of node dynamism.

MORRP facilitates high mobility by abstracting the concept of rendezvous *points* to rendezvous *regions* and forwards pack-

ets *probabilistically* based on which direction a destination or rendezvous node is most likely found. These directions shift accordingly to a node's *local* velocity. For example, if a source node is moving north, a node originally *east* of the source will *seem* to be moving *south*.

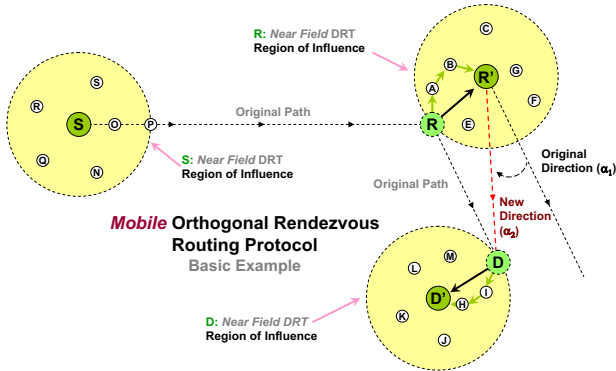


Fig. 1. Basic MORRP Example: Source sends to rendezvous node R found in region illustrated which in turn sends to destination D found in the DRT region given.

Figure 1 illustrates a basic example. Suppose source S wants to send packets to destination D and through announcement and route request (RREQ) packets, the path “Original Path” is established between S and D with node R as the rendezvous node. After some time, node R has moved to R’ and node D has moved to D’. With *infrequent updates* in a mobile environment, node R wishes to maintain a general direction to node D based solely on *local* information (its own mobility pattern) and adjusts its direction of sending to D from angle α_1 to α_2 . All nodes maintain a “field of influence” where each node knows the relative direction to all nodes in its region. The data packets S sends to D will traverse the original path, “gravitating” toward R’ once it hits R’s field of influence. Then, it will be sent in the modified direction of D until it hits D’s field of influence and “gravitates” toward the destination.

MORRP routes packets using directionality in highly mobile environments by 1) shifting destination node directions based on a node’s *local* velocity and 2) increasing probability of finding nodes by introducing “fields of influence”. All of this is done through a novel replacement to routing tables we formulate called the *directional routing table* (DRT).




Fig. 2. Directional Routing Tables (DRTs) map a direction to a set-of-IDs stored in bloom filters

The concept behind DRTs is simple: instead of maintaining *destination IDs* to *next-hop IDs*, we map a *probabilistic set-of-IDs* to each *interface direction* as shown in figure 2. The set-of-IDs are stored in bloom filters that are aggregated and sent to neighbors who merge them with the set-of-IDs associated with the interface of receipt. The information in

the filter becomes less useful as we progress in time and space and thus we decay (remove bits) from each bloom filter before sending it to its neighbors to capture this effect. Closer nodes have “more” information because the rate at which they are being updated by the source node is higher. Because DRTs only maintain information on each interface rather than on specific routes in a network, it adds more robustness to mobility as it provides several alternative paths for reaching a destination. In short, any next-hop in a particular direction can take the packet forward. Naturally, the closer a packet gets to a destination node, the information intermediate nodes have about the location of the destination increases. For destinations too far for the source to have any information about the location, MORRP relies on route request (RREQ) packets sent in orthogonal directions to *rendezvous* with state information maintained by each node along an announcement path also disseminated in orthogonal directions. This lightweight method of information dissemination ensures low control overhead from being flooded network-wide.

Key contributions of MORRP include:

- **Using directionality to solve the issues caused by high mobility in MANETs** - Using only *local* information, any node is able to more efficiently “guess” the direction of a destination and forward probabilistically.
- **The Directional Routing Table** - A replacement for traditional routing tables based on purely *probabilistic routing*. DRTs map a *set-of-IDs* to a specific *direction* which eliminates the need to maintain exact routing information about nodes in a network while lessening the frequency of route dissemination.
- **Routing Based on Probabilistic Hints** - Traditional routing protocols have a hard limit on route expiration. With probabilistic routing, routing information is decayed with time and becomes less and less accurate. Below a certain threshold, the information becomes insignificant.

In comparing with several proactive, reactive, and position-based routing protocols, MORRP shows high data delivery (93%+), low packet overhead, and over 10-14X goodput gains vs. traditional routing protocols and 15-20% goodput gains vs. traditional routing protocols modified with multiple directional interfaces in highly mobile (30m/s) environments. These gains come from many *key design factors*:

- **Weak state information and probabilistic routing** - MORRP does not maintain complete paths and is thus more flexible to forward packets in mobile environments.
- **Local update of weak state information** - Adjusting the “general direction” of a destination node based on one’s *local* velocity “takes a packet forward” even with infrequent location updates.
- **Field of influence** - Enlarging the intersection area results in a greater probability of finding a path in highly mobile environments even with infrequent updates.
- **Leveraging local direction information** - Limited flooding is curtailed by using *local* directionality to forward in straight lines and rely on intersections of announcement

and route request packets to “find” potential paths. This results in “freeing up” the medium for data. This is especially important because MORRP is a hybrid proactive/reactive protocol.

The rest of the paper is organized as follows: Section II and III outline the concept of MORRP including a detailed explanation of DRTs and several decaying strategies as well as route information dissemination and maintenance. Section IV gives some simulation performance evaluations and section V concludes with some thoughts on future work.

II. THE DIRECTIONAL ROUTING TABLE

One of the underlying mechanisms behind MORRP’s *probabilistic* forwarding strategy is the directional routing table (DRT), a simplified method of storing route information by leveraging directional communications methods. Unlike traditional routing tables which map *destination-IDs* to *next hop IDs*, DRTs map a *set of IDs* to a specific interface direction. In other words, all the nodes covered by the transmission sector of a specific antenna are included in the entry for that interface in the DRT. The number of entries in the DRT remains constant based on the number of interfaces and does not grow even as the number of nodes in the network grows. This is done through bloom filters.

The concept of using *bloom filters* in probabilistic routing schemes is not new. Acer et al. [6] and Kumar et al. [17] have suggested novel, decentralized, and scalable ways on how information can be disseminated in various types of networks using bloom filters. Bloom filters are space efficient probabilistic data structures that are used to test whether an element is a member of a set. Given an array of bits A (the bloom filter) initialized to all 0 and a fixed number (k) of hash functions ($h_1(\cdot), \dots, h_k(\cdot)$), elements (x) are inserted into the bloom filter by evaluating the element in each hash function and mapping the resultant locations in the array to one ($h_i(x) = 1, i = 1, 2, \dots, k$). Lookups are done in the same way in that if the positions in the bit array corresponding to the hashes of an element all equal 1, then the element is a member of the set.

Kumar et al. [17] introduced exponential decay bloom filters (EDBF), a data structure based on the traditional bloom filter concept. Instead of testing whether an element is part of a set or not (absolute information), EDBFs *count* the number of 1’s in the bit array corresponding to the element hash in lookup ($\theta_x = |\{i | A[h_i(x)] = 1, i = 1, 2, \dots, k\}|$). The fraction of bits set to 1 over the number of hash functions can be used to interpret the certainty of an element being in the set. Bits are “dropped” (decayed) using various strategies. In this paper, we apply the concept of EDBFs to store a probabilistic set-of-IDs corresponding to neighbor nodes a sector antenna covers in a MANET. We generalize the term to decaying bloom filter (DBF) as there are many ways to decay bloom filters.

Figure 2 outlines the structure for the DRT. In short, a *set-of-IDs* stored in a decaying bloom filter is mapped to each specific *interface direction*. To find the certainty of reaching a node by sending out a specific interface, the DBF associated with the

interface is selected and the destination node ID is sent through each hash function. By counting the number of bits set to “1” in the locations where the hashes land, the level of certainty of reaching a destination node by sending out that interface is obtained. As time goes on and without frequent updates, the level of certainty decreases. To facilitate this idea, we decrease the level of certainty by “decaying” bits in the bloom filter (i.e. changing bits in the DBF from 1 to 0). Decaying methods can be broken up into two main thrusts: *intra-node decay* which handles how bits are removed to simulate that as time goes on, there is *less certainty* about information, and *inter-node decay* which dictate how bits are removed as information is passed from node to node, simulating that nodes farther away know *less* about a node than nodes closer to the node. In the following subsections, we detail each method.

A. Intra-Node Decay

1) *Time Decay*: Current routing strategies employ hard timeouts for routing entries, updating routing entries periodically through route dissemination or route discovery. While effective for low mobility situations, high mobility situations can cause routes to become stale quickly if the interval between route updates is not decreased. As a result, maintaining accurate routing entries network-wide poses a huge problem as it incurs a much higher overhead. MORRP attempts to mitigate this issue by decaying the likelihood a neighbor or destination is in the direction covered by a specific interface as time moves on. In stationary environments, the probability of a neighbor being in a specific region decays at a constant rate (bits from the bloom filter are removed randomly at a constant rate).




Fig. 3. Each interface has a different relative notion of how fast a specific node is traveling.

In mobile environments, we employ a different strategy to decay neighbor location probabilities. Figure 3 illustrates the basis for our formulation of a simple *time decay* heuristic in mobile scenarios. Assuming all things constant, as a node moves away from its original position, the probability of neighbors in the direction of movement should decay slower than the nodes directly opposite of the direction of movement. In short, the velocity with which each interface perceives itself to be moving at is dependent on the angle the transceiver is from the direction of movement. As we wish to split the *intra-node decay* between *time decay* and *spread decay*, we will only use half the bits in each bloom filter in our calculations.

We formulate our *time decay* heuristic as follows:

Step 1: Suppose v_x is the speed a node is moving in the “x” direction and ϕ is the angle a specific interface is from

the direction of movement. We define the velocity as seen by a specific transceiver v_ϕ as:

$$v(\phi) = v_x \cos(\phi) \quad (1)$$

Step 2: If we let R be the range of a transceiver, a node traveling directly *away* from a specific direction at velocity v_x would be out of the transmission region in $\frac{R}{v_x}$ seconds. As a result, we specify that *all* bits of the bloom filter in a specific interface direction must be decayed in $\frac{R}{v(\phi)}$ seconds.

Step 3: Assuming there are k bits of ones in the bloom filter for a specific interface and half of those k bits ($\frac{k}{2}$) are reserved for *time decay*, we linearly decay the number of bits in each bloom filter for each interface with respect to time and velocity. The number of bits to remove per time interval (δ_t) is:

$$\begin{aligned} \delta_t &= \frac{ktD_{tc}}{2} - \frac{ktv(\phi)}{2R} \\ \delta_t &= \frac{kt}{2} \left(D_{tc} - \frac{v_x \cos(\phi)}{R} \right) \end{aligned} \quad (2)$$

Where $k/2$ is the number of bits reserved for *time decay* ($\frac{1}{2}$ the total bits set to 1 in the bloom filter), t is the time, D_{tc} is the time decay factor in the stationary case (D_{tc} fraction of bits removed per second in the stationary case), R is the transceiver range, v_x is the velocity in a specific direction, and ϕ is the angle from the current interface to the direction of movement. These bits are removed and discarded.




Fig. 4. Each interface/transceiver has a specific coverage region. As a node moves in one direction, the spread overflows to regions covered by neighboring interfaces.

2) *Spread Decay:* In a mobile environment with directional communications, the probability a neighbor will be in a certain transmission region/sector is *stretched* over time. As time progresses, the area a neighbor is possibly located, increases. Figure 4a illustrates this concept. Suppose a neighbor announces its position to be within region 2. Without knowing what direction and velocity the neighbor is traveling at, as time progresses, there is a greater possibility that the neighbor will be in region 1 and region 3 and a lessened probability that the neighbor will be in region 2. We say that as time goes on, the “spread” for the area the neighbor is in, is increased.

In much the same way, a mobile node traversing in a certain direction will need a greater spread to cover the same area in the direction it is traveling in. Figure 4b illustrates this. As a node trying to cover range θ_1 moves in the “+x” direction, it will need a greater spread, θ_2 to cover the same transmission region *in* the direction it is traveling while at the same time, a smaller spread, θ_3 to cover the same region in the direction *away* from the direction it is traveling. Each

direction other than the direction the node is traveling in and the direction directly opposite has varied stretch in between these two extremes based on the angle from the direction the node is traveling.

Unlike in our *time decay* heuristic formulation, bits removed from the bloom filter are not discarded but instead, *relocated* to the surrounding directions. The inherent nature of bloom filters allows us to move bits in the DBF associated with a specific interface, to surrounding DBFs, keeping the bits set to 1 *in the same hash locations*. Due to space constraints, we do not go into details regarding spread strategies. For our simulations, we assume a simple heuristic that whatever bits were affected by the $\frac{v_x \cos(\phi)}{R}$ term in equation 2 are affected in the opposite way for spread decay (ie: if the bits were removed, they are spread). When there is no mobility, there is no spread decay. It is important to note the *duality* of *time* and *spread* decay: A neighbor in the direction of travel will incur *less* time decay but at the same time, *more* spread decay.

B. Inter-Node Decay




Fig. 5. Neighbor information is decayed going farther from the source.

The general idea behind decaying the information transferred *between* nodes is that nodes “closer” to a specific source will most likely have more accurate information about the location of the source than nodes “farther” away. Nodes that are much farther away from the source will have so little information on the source that it will be indistinguishable from “noise”. Figure 5 illustrates this principle: Node A is a 1-hop neighbor of Node B. Node B aggregates its information about all its neighbors and decays this information before sending it to node C. Node C does the same thing with all its neighbors and what results is less and less accurate information about any node in a network depending on the distance that node is from the source.

1) *Exponential Distance Decay:* Updates are easily created by aggregating the DBFs associated with each interface in the DRT. We follow much of the same aggregation techniques presented in [17] in decaying bits exponentially with number of hops. Exploration of various distance decay methods are beyond the scope of this paper.

To perform distance decay, the local node ID is first hashed into a DBF U . Then, each of the DBFs associated with each interface is bitwise decayed according the decaying function $f_b(r)$ and bitwise OR-ed with U . U is then broadcasted out all interfaces to all neighbors.

Upon receipt of the aggregated and decayed DBF from a neighbor, a node will take the max number of bits counted for each entry and bit-wise OR the received DBF with the

DBF associated with the interface it receive the packet for the max number of bits for that entry. The reason we cannot simply bit-wise OR the entry with the received DBF is because with increased number of hash functions, probabilities will be biased toward directions with more neighbors since there is higher probability that even if neighbors have same amount of information about a specific node, the bits associated with that information will be more spread out.

Dissemination, which occurs periodically, only takes place between 1-hop neighbor nodes and requires no route/path maintenance. A common assumption in wireless routing protocols is neighbor discovery (each node knows its 1 hop neighbors) and this is usually achieved through periodically broadcasting *hello* packets to all nodes within transmission range. By piggy-backing dissemination information on these *hello* packets, we can therefore disseminate DRT information to our 1 hop neighbors without additional overhead.

C. Design Variables and Considerations

There are several factors to consider in designing routing algorithms based on DRTs. Table I lists several parameters that affect successful packet delivery using DRTs. Exploration of all the variables is beyond the scope of this paper, however, in section IV, we examine how varying some of the constraints affect routing in MORRP.

TABLE I

PARAMETERS AFFECTING SUCCESSFUL PACKET DELIVERY	
<i>Network Density</i>	Average number of neighbors
<i>Num of Interfaces (ϕ)</i>	The number of interfaces per node
<i>Time Decay Factor (D_t)</i>	Fraction of bits in bloom filter dropped per second per time interval D_i
<i>Time Decay Interval (D_i)</i>	The time interval to do decaying
<i>Dist. Decay Factor (D_d)</i>	Fraction of bits in bloom filter to drop per hop
<i>Near/Far-Field Threshold ($thresh \mid ff_thresh$)</i>	The number of bits found for it to be considered a positive result in searching in NF/FF DRT
<i>Spread Ratio (s_ratio)</i>	The ratio between bits used for spread decay and bits used for time decay
<i>Bloom Filter Size (m)</i>	The number of bits in each bloom filter
<i># of Hash Funcs (k)</i>	The number of hash functions

III. MOBILE ORTHOGONAL RENDEZVOUS ROUTING PROTOCOL

MORRP relies heavily on DRTs to provide probabilistic routes from source to destination. Because information about nodes farther away tend not to need to be refreshed as often as nodes closer to a source [10], MORRP is broken into two major arenas of operation, each with a separate DRT updated at different intervals: *near field* and *far field*. The near field handles direction changes and information about 2-3 hop “neighbors” while the far field handles everything beyond the near field’s “region of influence”. Near field operation including information dissemination is fairly straight forward and follows what is described in section II-A. In this section,

we will focus mainly on reaching nodes that are not in the immediate vicinity of the source (i.e. nodes in the far-field).

A. Assumptions

MORRP relaxes many of the assumptions made by position-based routing protocols (no need for location discovery and coordinate space embedding) while still providing connectivity even in *highly mobile* environments. To do so, MORRP assumes 3 givens: 1) Any given node will know (i) its 1-hop neighbors and (ii) the given direction/interface to send packets to reach this neighbor, 2) Each node must have its own local perception of direction with each interface aligned accordingly, and 3) RREQ and RREP send/receive time is negligible.

B. Far Field Operation

To facilitate routing in the *far-field*, MORRP sends RREQ packets in orthogonal directions (randomly choosing a neighbor in each orthogonal direction) and when one of these RREQ packets intersects the path of the destination’s announcement packets (also sent in orthogonal directions at periodic intervals), a RREP packet is sent back to the source. MORRP stores only *weak-state*[6] at each hop and because of infrequent updates, the far-field DRT is decayed at a slower rate than the near-field DRT. The specifics of this hybrid proactive-reactive model are described in the following subsections.




Fig. 6. 1: MORRP Announcements used to generate rendezvous node-to-destination paths 2-3: MORRP RREQ and RREP Packets to generate source-to-rendezvous node paths 4: Data path after route generation

1) *Proactive Element*: In order for a source and destination to agree upon a rendezvous node, pre-established “routes” from the rendezvous node to the destination must be in place. Because each node has merely a local sense of direction, making no assumption on position and orientation of other nodes in the network, it can only make forwarding decisions based on its own neighbor list. As mobility is increased however, routes become stale more quickly. Upon a set interval, each node sends MORRP announcement packets to its neighbors in orthogonal directions as shown in figure 6. When those neighbors receive these MORRP announcement packets, it hashes the ID of the source of the packet into the far-field DRT entry corresponding to the interface/direction it received the announcement packet and stores/updates the shortest number of hops associated with this announcement sequence number to the announcement source in a “hop count” table. This happens only if the sequence number of the packet is greater or the hop count is less than that recorded in the

table (better or newer path). Note that this “hop count” table is not maintained in any traditional sense and only updated once we have routes. The packet is then forwarded out the interface exactly opposite in direction from the interface it received the packet. If no neighbor is found in the opposite interface to send the MORRP announcement, a multiplier angle method (MAM) [3] is employed to attempt to maintain straight paths or forward along the perimeter as much as possible.

The entries in the far-field DRT are decayed in the same way as the near-field DRT with *intra-node* decay methods described in section II used. In this way, even moving nodes maintain a general sense of direction for any source they receive an announcement packet from. Time decaying methods ensure that positioning of nodes become less and less accurate with time and eventually, the information a specific node has about another node becomes negligible if not updated. Unlike the near-field DRT, however, far-field DRT is *not* shared with neighbors so *inter-node* decay is not used. This is to minimize indirection confusion.

2) *Reactive Element*: In order to build the path from source to rendezvous node, an on-demand, reactive element to MORRP is necessary. When a node wishes to send packets to a destination that is not within its immediate neighbor table or near-field DRT, it creates an entry in a simple *destination-rendezvous node* table and sends out a route request packet (RREQ) in all four of its orthogonal directions. Due to the fact that far-field DRTs only track nodes that send MORRP announcements or RREQ packets along the line, the *destination-rendezvous* table keeps track of which rendezvous nodes to forward to for a specific destination. Until a RREP is received, this entry is considered unusable.

When a neighbor node receives this RREQ packet, it hashes the node ID of the source into its far-field DRT and forwards the packet in the opposite direction utilizing MAM. Because one of the assumptions we made is that RREQ and RREP send and receive times are negligible compared to node movement, we need to add a short-timeout reverse path to the source so RREP packets can be sent back quickly. A simple *destination-next-hop* routing table with fast entry expiry times is used for this reverse-route back to the source.

In a 2-D Euclidian plane, by sending a RREQ packet in all 4 of its orthogonal directions, it is highly likely to encounter a node that has a path to the destination. When a node with a path to the destination (destination is either in neighbor table or destination ID is above threshold in near or far-field DRTs) receives the RREQ, it sends a RREP packet back to the source. Because we assume that nodes have not moved much in the process of the RREQ being sent, it is able to forward the RREP back efficiently. Finally, when the source receives the RREP, it hashes the rendezvous node’s ID into its far-field DRT and updates the *destination-rendezvous* table with the rendezvous node for a specific destination and “activates” that entry.

3) *Data Delivery*: For data delivery, if the packet is at the source, first the neighbor list and near-field DRT is queried for the destination. If destination is not found in these two tables, then the far-field DRT is checked to see if the number of bits

TABLE II

DEFAULT SIMULATION PARAMETERS

Parameter	Values
Trans. Radius / # Interfaces	250m / 8 Directional Interfaces
Topology Boundaries	1300m x 1300m, 2000m x 2000m
# of Nodes / Simulation Time	100, 200, 300 / 70s
Annnc. Interval / Mobility (m/s)	4.0s / RWP 0m/s - 30m/s
Distance Decay Factor (D_d)	0.7 (fraction of bits dropped per hop)
Time Decay Factor (D_t)	0.3 (fraction of bits dropped per sec)
Time Decay Interval (D_i)	0.5s
# of BF Hash Funcs / BF Size	30 / 16000 bits
NF Threshold / FF Threshold	6 bits / 6 bits
Spread Decay Ratio (s_{ratio})	0.5

associated with the destination hash is above the threshold. If destination is still not found in the far-field DRT, then the destination-rendezvous table is queried to see if there is a rendezvous node we need to send to. If it is found, then the far-field DRT is queried for the rendezvous node ID. If after all these steps the destination is unreachable, then a RREQ is sent out in orthogonal directions.

For forwarding packets, a similar approach is taken in that first the neighbor list and near-field DRT is checked for the rendezvous node if its present in the packet header and if not, the destination node. If it is not found in either, the far-field DRT is checked. If it is not found in any of the tables, the packet is simply forwarded to the opposite direction of receipt (the antenna exactly 180° from the receiving antenna).

IV. PERFORMANCE EVALUATION

In this section, we provide performance evaluations of MORRP under various parameters and against several proactive, reactive, and position-based routing protocols with one omni-directional interface and several directional interfaces. The simulations were performed using NS2 [19], with nodes using the standard IEEE 802.11 MAC and a 250m antenna range (NS2 default). Random waypoint mobility model with a node pause time of 5 seconds in $1300 \times 1300m^2$ and $2000 \times 2000m^2$ topologies used.

The performance metrics we evaluated are *packet delivery ratio*, *control packet overhead*, *average path length*, *aggregate network goodput*, *end to end latency*, and *far-field vs. near-field DRT usage*. We examine these metrics under conditions of varying *node mobility speeds*, *decay factors*, *transmission rates*, and *network densities*. All simulations were averaged over 3 runs of 5 different random topologies (total 15 trials). Table II outlines our default simulation parameters.

MORRP and ORRP were configured using n interfaces (divisible by 4) with each interface having a beam-width of $360/n$ degrees and announcement and RREQ packet TTL set to 10 hops. Announcement packets were sent every 4 seconds. We choose 30 hash functions and a bloom filter size of 16000 bits for simulations with MORRP to ensure minimum overlap of bits with 100 or so nodes and employ no bloom filter compression. The exploration of optimal hash function sizes to ensure minimal bit collisions are beyond the scope of the paper and more information can be found in [18].

For *reactive* routing protocols like AODV which require no periodic updates and position-based protocols like GPSR

with GLS as the location service, the standard NS2 defaults were used. For OLSR, we set the TC update interval to 4 seconds to match ORRP and MORRP announcement intervals. For all simulations, a hello interval of 2 seconds was used and MAC layer feedback employed for all the routing protocols. A potential future extension is MORRP with routing metrics. Traffic patterns varied for each test and are described in each subsection. Implementations and defaults for GPSR/GLS and OLSR can be found at [14] and [15] respectively.

In order to explore whether MORRP gains were merely from capacity gains with multiple directional antennas or *actual* design improvements, we modified AODV and OLSR implementations to support multiple directional interfaces in the same way as MORRP and ORRP. Since AODV and OLSR rely on omni-directional broadcast to disseminate information, sending out all interfaces simulates the behavior of AODV and OLSR broadcasts. Transmitting data packets, however, require only one interface to be active at a time freeing the medium and other interfaces for other nodes to use.

A. Evaluation of MORRP Parameters in Mobile Environments

1) *Effect of Time and Distance Decay Factors:* Sections II-A and II-B infer that knowing how many bits of the bloom filter to “decay” per time interval and per hop will affect reachability. In this section we evaluate how *time* and *distance* decay factors (what fraction of bits are dropped) affect reachability and far field DRT usage. The smaller the decay factor, the *less* fraction of the bits are decayed per time interval for *time decay* and per hop for *distance decay*. In our simulations, we fixed the default values given in Table II and vary the decay factors from 0.1 to 1.0. Figures 7 and 8 show our results for 100 nodes with an average of 10 1-hop neighbors.




Fig. 7. As time decay factor increases, reach probability drops due to confusing paths as old information is not decayed adequately. As distance decay factor increases (more bits dropped per hop), a plateau in reach occurs.

As expected, for various degrees of mobility, decreasing the time decay factor (dropping *less* bits per time interval) results in lower reach probability due to misinformation and bit accumulation. On the opposite spectrum, having too high of a time decay factor, thereby dropping a high number of bits per time interval also leads to less reach probability. This can be explained by the *far-field* DRT usage graph. As the time decay factor is low, the majority of data packets will be utilizing the far-field DRT to find a path because the near-field DRT coverage region will decay rather quickly. Because far-field




Fig. 8. Increased dependence on Far-Field DRT happens when more bits are dropped every time interval and distance.

TABLE III
COMPARISON OF ADHOC ROUTING PROTOCOLS

Type	Requires Localization	No Localization
Proactive	–	OLSR, ORRP, MORRP
Reactive	GPSR/GLS	DSR, AODV, ORRP, MORRP

DRT information is updated *less* frequently, too much reliance on it can yield inaccurate results. An optimal decay factor must be selected, therefore, to ensure high delivery success and a fair usage of both far-field and near-field DRTs.

Our results in figure 7 show that there is a gradual increase in reach probability when the distance decay factor goes from 0.1 to 0.6 and then plateau’s out. The low reach probability when the distance decay is lower results from saturation of bits to multiple interfaces resulting in confusing paths chosen. While the reach probability plateau’s at a distance decay factor of 0.6, the far-field dependence graph in figure 8 shows that there is still a gradual shift from using near-field DRT to route information at the source to far-field DRT dependence.

2) *Effect of Threshold:* Another interesting knob to adjust is the near-field and far-field bit threshold. As the threshold increases, the amount of information a node has about neighbors farther away decreases. It is expected that as the threshold for the near-field increases to greater than or equal to the number of hash functions, each node will only have information about itself. From our results, it was shown that when the threshold is low, there is a lot more confusion about path selection because of bit collisions and slow time decay. This results in poor path choices, low reachability, and greater dependence on the far-field DRT. As threshold increases, it approaches a point where each node has information about itself and its 1 hop neighbors (since 1 hop neighbors do not decay their node own ID hash when sending their DRT) resulting in high reachability. Furthermore, it was seen that reach probability peaks at roughly 6 bits which is roughly 20% of the number of hash functions used.

B. Comparison of MORRP against AODV, GPSR/GLS, OLSR and ORRP

In this subsection, we evaluate MORRP against reactive protocols like AODV [9], proactive protocols like OLSR [2], and position-based protocols like GPSR/GLS [1], [16] in terms of reach probability, average path length, control overhead, delivery success, aggregate network goodput, and end-to-

end delay under conditions of varying mobility speeds, data rates, and network densities. Table III shows a comparison of the classification of ad-hoc routing protocols measured, highlighting the need for node localization and key differences. Default parameters for MORRP given in table II were used in all scenarios unless otherwise stated. We focus heavily on *reachability/delivery success* in all these scenarios because in mobile adhoc networks, reachability comes primary over throughput, latency, etc. The reason is because our results show that for high mobility, even limited-flooding protocols like AODV and OLSR simply cannot deliver the majority of the packets (low reachability).

1) *Effect of Increased Velocity*: In this subsection, we evaluate the effect of increasing velocity on traditional routing protocols like AODV, GPSR/GLS, and OLSR and compare it to MORRP, ORRP, and multi-interfaced versions of AODV and OLSR. Our initial simulations involve relatively light load (1000 random 5 second connections). While protocols like GPSR/GLS provide high reach under light load, as the load increases to 10,000 connections, we see a significant drop in reachability. Figure 9 shows our results in comparing MORRP to traditional routing protocols with one omni-directional antenna under varying node and topology sizes.

It is clear that in conditions of high mobility with few connections, MORRP with atleast 8 interfaces provides high reach probability (93% for $1300 \times 1300m^2$ networks and 87% for $2000 \times 2000m^2$ networks) even under conditions of infrequent announcements sent (4 second intervals). As maximum velocity increases, AODV and OLSR fail because of stale routes. With high mobility, it becomes increasingly hard to maintain end-to-end routes without increasing state dissemination rate or route requests. Both options lead to network congestion.

GPSR with GLS performs rather well in all cases, providing high reach even with increased mobility. Simply using GPSR with GLS (or a version modified to support directional antennas), however, runs into several issues: 1) End-to-end packet latency with GPSR with GLS is fairly bad, reaching close to 3-4 seconds per packet (See figure 10) and 2) GPSR/GLS requires node localization which we currently assume is a given. Position information is often obtained through node localization protocols which incur additional overheads and function poorly in sparse network environments or additional hardware like GPS which require “sky” access.

To see where the multiple interfaces shine, we increase the load from 1,000 connections to 10,000 connections, each for 5 seconds. As figure 10 shows, under increased load, protocols that utilize omnidirectional antennas saturate the medium and fail to successfully deliver packets. It becomes important, therefore, to understand how much of the gains from figure 10 results from more efficient medium reuse due to directional interfaces vs. the gains coming from MORRP protocol design itself. We compare modified versions of AODV and OLSR to support multiple directional antennas with ORRP and MORRP with the results shown in figure 10. The modified versions of AODV and OLSR still broadcast (ie: send out all

interfaces) when performing route requests or dissemination due to the protocol design and as such, we expect to see better performance with MORRP.

AODV with 8 directional interfaces shows significant improvement in delivery success vs. the traditional AODV due to the directional interfaces causing less interference in data delivery. The primary gains come from utilizing multiple interfaces as AODV is a reactive protocol and sends out route request packets on-demand. OLSR, on the other hand, saw only gains vs. the single interface OLSR only under low mobility. When the mobility increases to 30m/s, these gains from capacity almost fade. This is due to OLSR’s proactive nature. The periodic rate of link-state exchange simply cannot keep up with mobility speed and as a result, most packets are not successfully transmitted due to stale information rather than medium saturation. ORRP delivery success drops with increased mobility because it cannot maintain straight line next hop paths without constant updates. MORRP performs consistently well, delivering over 93% of the packets even in highly mobile environments.

2) *Effect of Increased Network Density*: In this subsection, we evaluate how increasing network density affects *reachability* and *amount of control packet bytes* networkwide. The reason why we focus on control packet bytes rather than control packets is simply because MORRP sends bloom filters to its immediate neighbors. Although bloom filters are relatively small in size, the incurred overhead is larger than traditional packets. Figures 11 and 12 shows our results varying number of nodes from 50 to 300 with each node having a maximum velocity of 30m/s. 2500 random source and destination pairs are chosen and 512KB CBR packets sent for 20 seconds at a rate of 2Kbps. For fair comparison, we only evaluate MORRP against ORRP and the modified versions of AODV and OLSR to support multiple directional interfaces.

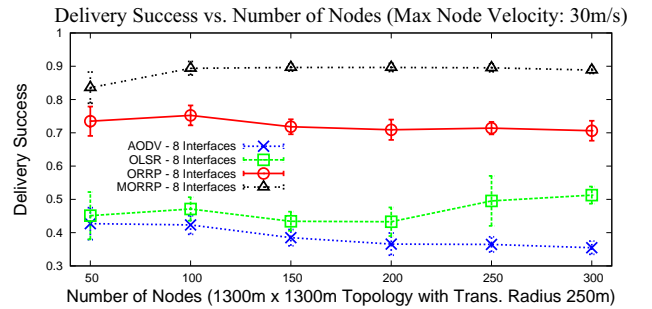


Fig. 11. MORRP scales well under increasing node density, providing over 90% delivery success in most cases.

It can be seen that as the number of nodes increases, AODV with multiple interfaces start dropping in reach due to its broadcast nature. OLSR fails because of stale routes due to high mobility. As the density increases, however, OLSR performs seemingly better because closer nodes are more within better reach. ORRP fails to deliver packets because in highly mobile environments, straight line paths are hard to maintain. MORRP delivers roughly 90% of the packets successfully.




Fig. 9. MORRP yields above 93% reachability (data delivery success) even in highly mobile environments for medium-sized networks (1300m x 1300m) and about 87% reachability for large-sized networks (2000m x 2000m) with moderate density.




Fig. 10. GPSR w/ GLS has much higher end-to-end latency than all the other routing protocols including MORRP. Under load, all other routing protocols with omnidirectional antennas fail. MORRP performs much better than traditional routing protocols modified with directional antennas.




Fig. 12. MORRP incurs low packet overhead even in high mobility.

It is interesting to note that MORRP seems to send out *less* control packets than ORRP despite it needing to periodically send DRT update messages to all neighbors. The reason for this is simple: In ORRP, RREQ packets travel in a line and a RREP is generated only when this packet intersects with a path generated by an ORRP announcement packet. With MORRP, however, RREQ packets stop being forwarded once it intersects with a destination's "field". Because these "fields" are two or three hops large, MORRP RREQ packets traverse less hops than ORRP RREQ packets. OLSR grows rapidly with network size because more nodes are periodically sending out link-state information. AODV grows despite the constant number of connections due to more nodes in the network forwarding RREQ and RREP packets.

3) *Effect of Increased Data Rate:* Although in mobile environments, high reachability naturally leads to high aggregate network goodput, it is important to quantify these gains. In this subsection, we evaluate the effect of increased data rate on network goodput. To do so, we make all-to-all connections simultaneously network-wide and send packets at a set data rate for 20 seconds. By slowly increasing the rate, we can measure the amount of data that actually gets sent. We expect the capacity constraints will be mostly dependent on medium usage. All nodes are moving at a uniformly distributed velocity




Fig. 13. MORRP achieves 10-14X more aggregate goodput compared to routing protocols with omnidirectional antennas due to efficient medium reuse with a max of 30m/s.

We first compare MORRP to AODV, OLSR, and GPSR/GLS to highlight the gains from simply moving from omnidirectional antennas to directional antennas. Figure 13 shows our results. As expected, MORRP with 8 interfaces achieves much higher goodput than all the other protocols (roughly 10-14X more than OLSR the closest competitor).




Fig. 14. MORRP achieves 15-20% more aggregate goodput over protocols with 8 directional interfaces because of use of directionality to limit flooding.

Comparing to AODV and OLSR with 8 directional interfaces and ORRP, MORRP still performs almost 15-20% better than OLSR and ORRP. Again, ORRP fails because it was

never designed for mobility and maintenance of straight-line paths becomes difficult in highly mobile environments. The gains from MORRP come from protocol design. Much like the majority of previous work in using directional interfaces in layer 3 routing [5][11], the modified versions of OLSR and AODV simply *adapt* the protocol to support directionality rather than leveraging the *inherent properties* of directionality to route. Whereas OLSR and AODV even with multiple directional interfaces simply “broadcast” out all intervals for topology control dissemination or route discovery, MORRP utilizes local directionality to disseminate packets along lines to limit flooding. Therefore, it is understandable to see large gains with MORRP over OLSR and AODV with multiple interfaces.

4) *Summary of Performance Evaluation:* Below we summarize our findings in evaluating MORRP:

- MORRP yields above 93% reachability even in highly mobile environments for medium-sized networks and 89% reach for large-sized networks with medium density.
- Routing using MORRP accounts for an almost 10-14x higher aggregate goodput compared to AODV, OLSR and GPSR/GLS. These gains come primarily through more efficient reuse of the medium under heavy load.
- MORRP yields 15-20% higher aggregate goodput compared to modified versions of AODV and OLSR for 8 directional interfaces and also ORRP. These gains come by using directionality constructively and scalably to overcome problems inherent with directionality.
- End to end packet latency is very low under MORRP compared to AODV, OLSR, and GPSR/GLS because of more efficient medium reuse.
- As node density increases, AODV, OLSR and GPSR/GLS data delivery success drops significantly due to network saturation but does not affect MORRP much.
- MORRP sends *less* control packets than ORRP and much less than AODV, and OLSR in highly mobile situations.

V. CONCLUSION

In this paper, we presented Mobile Orthogonal Rendezvous Routing Protocol (MORRP), an *unstructured, probabilistic, and highly mobility tolerant* forwarding paradigm based on directional communication methods and rendezvous abstractions. By utilizing directional routing tables (DRTs), a novel replacement for traditional routing tables, information about nodes in a specific region and nodes along a straight line path is maintained probabilistically. DRTs map interface *directions* to a probabilistic *set-of-IDs* which are decayed and spread locally within a node based on time and *local* node velocity and decayed by number of hops from the source. DRTs provide *regions* where a node can be found in the near-field case and *directions* to send in the far-field case.

When a destination is outside of the near-field region, MORRP relies on taking intersections of orthogonal lines originating from source and destination and forwarding packets from the source to rendezvous nodes which in turn hand them over to the destination providing simplified routing. We have

outlined several “knobs” associated with MORRP and evaluated *distance decay factor, time decay factor, and near-field and far-field threshold* under conditions of varying mobility. It can be seen that spread decay affects networks that are sufficiently dense and has very little effect on sparse networks. Additionally, we compared MORRP against AODV, OLSR, GPSR/GLS, and ORRP under varying conditions of mobility and node densities and found that MORRP provides higher reach probability, average path selection, and has much lower control packet overhead. In short, MORRP provides high connectivity even in highly mobile, dense, and unstructured environments.

While we have only considered the base case of MORRP in square topologies with random waypoint mobility, there are several directions for future work. First, it would be interesting to see how MORRP fits into hybrid routing environments with networks having a mixture of nodes with omnidirectional and directional communications. Additionally, it would be interesting to see how to incorporate routing metrics into MORRP and DRTs to provide for even better path selection and obstacle avoidance. Another area of consideration is a more detailed evaluation of MORRP under various topologies and traffic patterns.

REFERENCES

- [1] B. Karp and H.T. Kung, “GPSR: Greedy Perimeter Stateless Routing for Wireless Networks,” Proceedings of MOBICOM 2000.
- [2] OLSR RFC. <http://www.ietf.org/rfc/rfc3626.txt>
- [3] B. Cheng, M. Yuksel, S. Kalyanaraman, “Orthogonal Rendezvous Routing Protocol for Wireless Mesh Networks,” Proceedings of ICNP 2006.
- [4] A. Nasipuri, J. Mandava, H. Manchala, R. Hiromoto, “On Demand Routing Using Directional Antennas in Mobile Ad Hoc Networks”, Proceedings of IEEE WCNC 2000.
- [5] R. R. Choudhury, N. Vaidya “Performance of Ad Hoc Routing using Directional Antennas” Journal of Ad Hoc Networks, Nov. 2004.
- [6] U. Acer, S. Kalyanaraman, A. Abouzeid, “Weak State Routing for Large Scale Dynamic Networks”, Proceedings of MOBICOM 2007, Sept 2007.
- [7] D. Britz, R. Miller, “Mesh Free Space Optical System: A method to Improve Broadband Neighborhood Area Network backhaul”, Proceedings of IEEE LANMAN, Princeton, NJ, June 2007.
- [8] D. Johnson, D. A. Maltz, J. Broch. DSR: The Dynamic Source Routing Protocol for Multi-Hop Wireless Ad Hoc Networks. *Ad Hoc Networking*, Chapter 5, pp. 139-172, Addison-Wesley, 2001.
- [9] C. Perkins and E. Royer. “Ad hoc On-Demand Distance Vector Routing.” Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications, New Orleans, LA, February 1999, pp. 90-100.
- [10] BBN: Hazy Sighted Link State (HSLS) Routing Protocol <http://www.cuwireless.net/downloads/HSLS.pdf>.
- [11] R. Ramanathan. “On the performance of ad hoc networks using beam-forming antennas”, Proceedings of ACM MOBIHOC, October 2001.
- [12] S. Yi, Y. Pei, and S. Kalyanaraman, “On the Capacity Improvement of Ad Hoc Wireless Networks Using Directional Antennas,” Proceedings of ACM MOBIHOC, Pages 108-116, Annapolis, MD, June 2003.
- [13] M. Yuksel, J. Akella, S. Kalyanaraman, and P. Dutta, “Free-Space-Optical Mobile Ad-Hoc Networks: Auto-Configurable Building Blocks”, To appear in ACM/Springer Wireless Networks, 2007.
- [14] HLS patch for ns-2.29. <http://cn.uni-duesseldorf.de/staff/kiess/software/hls-ns2-patch>.
- [15] OLSR for ns-2.29. <http://masimum.dif.um.es/?Software:UM-OLSR>
- [16] J. Li, J. Jannotti, D. De Couto, D. Karger, R. Morris, “A Scalable Location Service for Geographic Ad Hoc Routing,” ACM Mobicom 2000.
- [17] A. Kumar, J. Xu, and E. Zegura, “Efficient and Scalable Query Routing for Unstructured Peer-to-Peer Networks,” in *IEEE Infocom*, 2005.
- [18] M. Mitzenmacher, “Compressed bloom filters,” *IEEE/ACM Transactions on Networking*, vol. 10, no 5, pp. 604-612, 2002.
- [19] The Network Simulator. ns-2. <http://www.isi.edu/nsnam/ns>.