

# Probabilistic Multi-path vs. Deterministic Single-path Protocols for Dynamic Ad-Hoc Network Scenarios

Christopher L. Barrett \*

Stephan J. Eidenbenz

Lukas Kroc

Madhav V. Marathe

James P. Smith

Basic and Applied Simulation Science (CCS-5)  
Los Alamos National Laboratory  
MS M997, P.O. Box 1663, Los Alamos, NM 87544, USA  
{eidenben, kroc, marathe, jpsmith}@lanl.gov

## ABSTRACT

We investigate the performance of different protocol stacks under various application scenarios. Our method of choice is a full-fledged simulation in QualNet, testing the complete protocol stack over fairly large-scale networks. We find that the relative ranking of protocols strongly depends on the network scenario, the session load, the mobility level, and the choice of protocol parameters. We show that the Parametric Probabilistic Protocols, which we generalize from their original definition, can outperform standard routing protocols, such as AODV or Gossiping or Shortest-Path, in a variety of realistic scenarios.

## Categories and Subject Descriptors

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

## Keywords

Wireless Networks, Mobility, Scenarios, Probabilistic Routing

## 1. INTRODUCTION

In the world of wireless communication, it would be wonderful to have a single, simple protocol that worked perfectly under any conditions, but protocols all behave differently in different circumstances. Every scenario has properties that can make one protocol inherently better than another. Examples of this include situations with static networks compared to cases with frequent or significant topology changes, or extreme cases of data flow, such as a network with a single uplink node compared to a network where all nodes are peers and communicate with each other. Dynamic topology is an especially important factor in wireless networks,

\*Current affiliation is Virginia Bioinformatics Institute, Virginia Tech, Blacksburg, VA 24061, USA. Email: cbarrett@vbi.vt.edu

and can be induced by node failures or temporary power loss, device mobility, or signal interference. Network load is another major discriminator: a protocol that produces the least load when there is very much network traffic may not do so when there is very little.

In this paper, we undertake a systematic empirical analysis of the most commonly referenced application scenarios for ad-hoc wireless network communication. We consider three different scenarios: (i) alarm and mapping tasks for sensor networks with a single destination node that needs to be reached, (ii) full any-to-any networking with all network nodes as possible sources and destinations, and (iii) a futuristic, but likely most realistic, multiple-basestation scenario with nodes connecting to one of several equivalent uplink destinations (i.e., base stations) in an ad hoc fashion.

We take a holistic view of the protocol stack (i.e., optimizing the whole stack rather than individual layers) from the data link to the network layer. As MAC protocols, we study CSMA and 802.11, and to a lesser extent, TDMA. On the routing layer we include some well-known protocols: AODV, Shortest-Path, Gossiping, and Flooding. We also consider a generalization of Parametric Probabilistic Routing (PPR), introduced in [3] as a family of multi-path routing protocols whose members are characterized by the function describing the probability of forwarding received messages. We design a generic form of the PPR retransmission probability that merges the two major protocol classes in [3] and Gossiping. We also provide a more intuitive explanation of the parameters in the retransmission function in one of the classes. Interaction between layers in a protocol stack have important performance implications. In [2], statistical methods to formally characterize the interaction between the MAC and the routing layers have been used. The objective of our experimental analysis is to study the combination of MAC/routing protocols that works best for a given scenario, session load and mobility, and to test the efficacy of our modified form of PPR in real-world simulations.

We use simulations to carry out our empirical analysis. All protocols are fully implemented in the network simulator QualNet [18], in which we simulate the entire protocol stack across a very large parameter space. A key finding is that the ranking of the routing protocols changes with increasing mobility for most scenarios, where a scenario is defined by a multitude of parameters. Intuitively, this prominent role of mobility is due to the fact that the protocols rely (to different extents) on the accuracy of global information (such as hop distance information) that is propagated through the net-

©2005 Association for Computing Machinery. ACM acknowledges that this contribution was authored or co-authored by a contractor or affiliate of the [U.S.] Government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

ACM/SAC March 2005, Santa Fe, New Mexico, USA  
Copyright 2005 ACM 1-58113-964-0/05/0003 ...\$5.00.

work by various means; this accuracy is crucially impacted by mobility. Our experiments are also grouped into three main scenarios in which we present a detailed comparison of the diverse protocols. We find that Shortest-Path wins in terms of best performance for low mobility in the single-destination scenario; AODV wins in any-to-any networks with no or low mobility, Gossiping wins in any-to-any networks for high mobility values; Directed-Transmission wins at all mobility levels in multiple-destination scenarios and in the high-mobility single-destination scenario.


**Related Work.** The work most closely related to ours is [3], where Parametric Probabilistic Sensor Network Routing was proposed and tested in a light-weight simulation environment, not taking into account issues of dissemination of global data or protocol layers other than routing. In addition to consideration of mobility and other classes of scenarios than simple sensor networks, the present simulation results are done using a full-fledged implementation including information distribution models and all network layers. Routing in wireless and ad hoc networks has been the subject of active research for quite some time; see [15] for the state of the art until 1995 and more recent survey papers [14, 5, 12]. Probabilistic routing and Gossiping have been applied in various settings, such as spreading distributed database updates [6] and multicast techniques (see [11] for a recent result and an overview). Non-constant probability Gossiping has been developed by Kempe et al. [9], in which the authors use a Gossip-based approach to spread alarm messages, where the probability of forwarding to a particular neighbor depends on the geographic distance to that neighbor. Multiple path routing has been applied in [7] in order to make a network resilient against node failures.

**Organization.** The remainder of this paper is organized as follows. We propose a generalization to the idea of Parametric Probabilistic Routing in Section 2. Section 3 presents the remaining protocols that we study in more detail. The different scenarios that we study are introduced in detail in Section 4. The setup of our simulations is explained in Section 5 and the results are given in detail in Section 6. Concluding remarks are in Section 7.

## 2. RECASTING PARAMETRIC PROBABILISTIC ROUTING

Parametric Probabilistic Routing (PPR) protocols are a family of light-weight adaptive routing protocols, where each member is fully described by a retransmission probability function and its parameters, indicating with what probability a node will forward a received message to all of its neighbors. The retransmission probability function can depend on various factors such as the hop-distance to the destination, the hop-distance of the source to the destination, the number of hops that the packet has already traveled, the number of times a node has already forwarded the same packet, the number of neighbors a node has, etc. PPR protocols thus perform directed, controlled flooding resulting in multiple packet copies, as opposed to classic single-path, single-packet-copy methods. Two family members, Directed-Transmission and Destination-Attractor were introduced for sensor networks in [3].

We adopt a slightly modified notation from [3]. Let  $S$  denote the source node (the origin of data),  $D$  the final destination node,  $R$  a node currently holding a packet and



**Figure 1: Distance labels on the network. Solid lines are local Shortest-Path estimates, dashed lines are paths actually traveled. In the bottom diagram, the solid curved line is the shortest SD path through R.**

deciding whether to retransmit,  $h$  the number of hops that the packet at node  $R$  has traveled so far, and  $c$  the number of copies of the same packet that node  $R$  has received so far. These distances are illustrated in Figure 1. The same node can hold a packet on different timesteps, and the same packet can be on multiple nodes at the same time. We denote the shortest-path hop-distance between two nodes, say  $S$  and  $D$ , by  $d_{SD}$ .

We now generalize the form of the retransmission probability for a PPR protocol that contains the Directed-Transmission and Destination-Attractor members as special cases. For node  $R$ ,

$$P = \exp[-k_1(d_{RD} - d_{SD} + k_2 \cdot h) - k_3 \cdot c] \quad (1)$$

Parameters  $k_1, k_2, k_3$  can be used to create different protocols. Parameter  $k_3$  controls the duplicate dropping behavior. If we want duplicates to be dropped unconditionally, we set<sup>1</sup>  $k_3 = \infty - 1$ . Parameter  $k_2$  determines the influence of the number of hops  $h$  that the packet has already traveled on the retransmission probability, essentially aging the packet, and parameter  $k_1$  influences the main distance terms. If we set  $k_2 = 0$  and  $k_3 = \infty - 1$ , the retransmission probability is exactly Destination-Attractor from [3]. Setting  $k_2 = 1$  and  $k_3 = \infty - 1$  yields Directed-Transmission.

This is essentially an information-enhanced gossiping protocol. In the Directed-Transmission version, instead of proceeding blindly through the network with a fixed retransmission probability, the probability is modified to keep the packets near the shortest path. There is a parameter to determine how far from the shortest path a packet may stray before being terminated,  $k_1$ . We now introduce a reformulation of this protocol in terms of more intuitive parameters. We define  $\Delta$  to be the deviation of the packet from the shortest path between the source and destination, and  $\delta$  as the deviation of the packet from the shortest path between the source to the present node  $R$

$$\Delta = d_{RD} + d_{SR} - d_{SD} \quad (2)$$

$$\delta = h - d_{SR} \quad (3)$$

This is also shown in Figure 1.  $\Delta$  is the minimum number of additional hops (over the shortest path hops) that is now

<sup>1</sup>We use the value of  $\infty - 1$  to express that  $k_3$  yields 0 when multiplied by 0, and  $\infty$  when multiplied by any positive number

possible from  $R$  to  $D$ , and  $\delta$  is the number of additional hops (over the shortest path hops) taken in getting to  $R$ . Substituting these in Equation 1 with  $k_2 = 1$  and  $k_3 = \infty - 1$ , Directed-Transmission becomes

$$P = \exp[-k_1(\Delta + \delta)] \quad (4)$$

It is now clearly seen that a packet that stays on the shortest path retains  $P = 1$  for its entire lifetime, regardless of the value of  $k_1$ . An infinite  $k_1$  is therefore equivalent to a pure shortest path algorithm. A packet that makes one extra hop, anywhere along the path, will have subsequent probabilities reduced by a fixed amount parameterized by  $k_1$ . This formulation also makes it possible to do analytic calculations of total delivery probability in the presence of well-defined mis-information, although this paper concentrates on experimental results.

### 3. OTHER PROTOCOLS

We aim to consider all relevant layers of the protocol stack. For the physical layer, we limit ourselves to the standard protocol set by QualNet. We are most interested in the MAC (data-link) and the network layer, while we feel that the transport layer can be modeled individually in the experiments by simply setting the number of packets that are to be delivered in a session without considering congestion or flow control issues. The various other protocols used in the scenarios considered are described here.

**MAC layer protocols:** We considered both CSMA and 802.11 paired with each routing protocol in our experiments. For use with the PPR family, including Gossiping, we found that CSMA consistently outperforms 802.11, which is due to the following: Since we only do broadcasts, CSMA is equivalent to 802.11 for our purposes, but it is more lightweight and provides us with the same reliability. We show results for both CSMA and 802.11 coupled with Shortest Path. For AODV, we show results for only 802.11 as this combination works consistently better than the one with CSMA.

**Gossiping:** Gossiping is a very simple routing protocol, where the retransmission probability function is a constant  $k$  that can be set by a network designer. Gossiping in this form was introduced by Haas et al. [8]. Gossiping can be regarded as a member of the PPR family, where the retransmission probability is simply a constant. Flooding is a special case of Gossiping (and thus also a member of PPR family) with retransmission probability set to one.

**Shortest Path:** In the presence of perfect, static global information, nothing will out-perform a Shortest-Path algorithm, which keeps only one copy of the packet alive at any given time and chooses the next node in the path uniformly among those that are closest to the destination. Distributing this information in the network is non-trivial. Shortest path routing is clearly an optimum protocol in ideal circumstances, but in reality the shortest paths to the destination are not always known and also change over time. With poor information, it is possible for a packet to “get lost”, visiting the same nodes more than once and never reaching the destination.

**AODV:** One of the best routing protocols in ad hoc networks is AODV, Ad-hoc On-demand Distance Vector routing. AODV is based on routing tables that indicate the next node to which a message should be forwarded in order to reach the destination quickly. Routing tables are updated only as necessary. If a link in the path is broken and a

packet can no longer be forwarded, that node initiates a local Flooding asking its neighbors if they know a route to the destination, which they then propagate. For a more detailed description of AODV, see [13] and [17].

**Propagation of Control Information:** Most routing protocols rely on information such as hop-distance to destination, next hop on shortest path, or numbers of hops already traveled. Pure Gossiping protocols and Flooding make no use of network information, and there is therefore no network information to propagate. The collection of information necessary to route a packet in AODV is specified as part of the protocol (route discovery), and the reader is referred to the references. The other protocols under consideration that require information are the PPR protocols and the Shortest-Path protocol.

Information relating to the network topology, like hop-distances for PPR and shortest path neighbors for Shortest-Path, requires additional messages to be sent around the network. Choosing the best method to propagate this information has a profound impact on the performance of a protocol; moreover, due to potential node mobility, no method can guarantee to compute the exact values for distances.

The method called *beacon flooding* works as follows: A destination periodically sends out a beacon signal, which is received by its neighbors that then know they are at distance one from the destination. The neighbors then propagate their distance information to all of their neighbors who know that they are at distance two from the destination, and so on (this can be done using small packets that carry just one number: distance to destination). Once the flood is finished, every node knows its hop-distance to the destination, moreover it also knows which of its neighbors is closest to the destination and thus is its shortest-path neighbor. The accuracy of the information is directly affected by the beacon frequency: more frequent beacons produce more load.

We assign an age limit  $T$  to each piece of information, after which the node will assume it to be invalid and reset its estimate to the minimum of the most recently obtained information: thus, the distance estimate of a node is one unit larger than the minimum of the estimates of its neighbors that the node received within the last  $T$  seconds. The problem this resolves is that otherwise, distance estimates could only monotonically decrease in time.

### 4. APPLICATION SCENARIOS

We consider three scenario classes in our experiments, designated by the number of nodes that can act as base stations or uplinks.

In the first scenario, only a single dedicated node can act as the destination node, we thus call the scenario the *single-destination scenario*. An example of this is an alarm task for a large sensor network: upon detection of an unusual value of the measured environmental variable, the corresponding sensor needs to inform the base station about this event. The destination node acts as a base station collecting alarm messages from the sensors and relaying it to the fixed infrastructure network. The limited transmission range of the sensors forces a multi-hop path from source to destination with intermediate sensors forwarding the message. Another example is nodes measuring an environmental variable and periodically send the measurements to the base station. This mapping task results in higher load of the network and typically concurrent sessions. In our experiments,

we model the alarm task by having 50 sessions, while having 1000 simultaneous sessions for the mapping task; both have the destination centered in the simulation area and stationary. For both sensor tasks we assume that only the values of variables are reported, thus only very few bytes that comfortably fit into a single packet need to be transmitted per session. For an overview of the state-of-the-art on sensor networks, see the survey [1].

In the second scenario class, a small number of designated nodes can act as destinations. We thus call it the *multiple-destination scenario*. Communication takes place between a source node and any destination node over a multi-hop path. Each destination is assumed to have a direct link to a fixed infrastructure network. Thus, all destinations or base stations are equivalent: one endpoint of a connection is a mobile node, the other can be any one of the destination nodes, no matter which one it is. It is even acceptable to send consecutive packets of the same session to different destinations. In our experiments we choose the number of destinations to be 9 so the average number of hops was neither too large nor too small; the destinations are distributed uniformly over the simulation area and do not move.

In the third scenario class, all nodes can act as destinations, which gives the name *any-to-any scenario*. All nodes are equivalent and communication can occur between any pair. Packets are larger and there is more than one packet in a session. This is the usual ad-hoc network scenario used in mobile communications research.

We believe that the multiple destination scenario is the most realistic scenario for the following reasons: the single destination scenario has an obvious single point of failure in the single base station and in order to achieve redundancy the multiple destination scenario seems like good alternative; the any-to-any scenario suffers from the well-known scalability issues encountered when trying to deploy ad hoc networks. A type of small-scale hybrid network similar to the multiple destination scenario was recently investigated, and found to have large performance improvement over either ad hoc or cellular networks on the same set of nodes [10].

## 5. EXPERIMENTAL DESIGN

In order to compare the protocols to each other, we ran simulations in version 3.5.2 of the QualNet network simulator [18]. An overview of the parameters defining our experimental setup is given in Table 1. In order to reduce the parameter combinations to a manageable and presentable level we ran several preliminary tests to determine reasonable values. We chose the random waypoint model as our mobility model, which has established itself as the standard model in the community despite known deficiencies (see [4]). All reported results are average values for 1000 runs of the same case with different random seeds.

Tables 2 and 3 give an overview over the parameter choices for the different protocols. All choices shown in these tables were influenced by a significant amount of preliminary testing, whose results can be obtained from the authors. As can be seen from these tables, we chose to use some standard parameters in draft-9 of AODV, limiting us to 72 instances of AODV protocols. We also explored all sensible regions of the parameter space of the PPR protocols, but show here only those listed in the tables.

### Measurements:

In order to present our performance comparison, we first

**Table 1: Overview of Experimental setup**

1. <b>Number of runs:</b> 1000 for <i>each</i> set of parameter values
2. <b>Simulated time:</b> 100 seconds
3. <b>Number of nodes:</b> 1000
4. <b>Grid size:</b> 6678 m x 6678 m
5. <b>Transmission range:</b> 376 m (resulting in an average connectivity of 10)
6. <b>Initial node distribution:</b> uniform random
7. <b>Mobility model:</b> Random waypoint
(a) <b>Speed intervals [meters/second]:</b> 0, 5–10, 10–20, 20–40, 40–80, 80–160, 160–320, 320–640
(b) <b>Waiting time:</b> zero seconds
8. <b>Number of sessions per run:</b> (spread randomly in the simulated interval)
(a) Single-destination scenario: 50, 1000
(b) Multiple-destination scenario: 1000
(c) Any-to-any scenario: 1000
9. <b>Number of packets per session:</b> 1 (single- and multi-destination scenario); 1, 3, 10, 25 (any-to-any)
10. <b>Packet size:</b> 16 B, 1 KB

define the performance measures we used. We define the *reliability* as the percentage of successfully delivered packets, with success defined as at least one packet copy reaching the destination. The *load* imposed on the network during a simulation is defined as the total amount of energy consumed by all transmitting nodes in milli-Watt hours (mWh). This is proportional to the total number of packets when packets are all the same size, but when a protocol has different size packets (e.g. small control and large data packets), it can be quite different. We feel these two measures effectively characterize some of the most important trade-offs a protocol can make, and can essentially quantify the energy used for every packet delivered. Preliminary investigations also looked at packet collisions, signal error, separating data and control energy, packet numbers, and other measures, but we think load and reliability are most useful.

When we plot load versus reliability, we fix a scenario and show one point for each combination of protocol and parameters. The points are connected into lines in subsets where for each protocol, all but one parameter are held fixed, and the remaining parameter creates a line. In Gossiping, that variable parameter is the retransmission probability, in Directed-Transmission it is the  $k_1$  parameter, and in Shortest-Path it is the beacon period. In most cases, all points for a given set of parameters for a protocol perform better than all points for another set. That is, lines do not cross. It is then easy to show only the one best curve, as is done in all cases for Gossiping and Directed-Transmission. There is no clear trend for AODV, in which case we show all experiment points (not connected by lines). For Shortest-Path no MAC protocol clearly dominates the other, thus we usually have to draw two curves for Shortest-Path. Thus each plot of load and reliability will have a set of points for AODV, one connected curve for Gossiping and Directed-Transmission, and two curves for Shortest-Path (one corresponding to CSMA and one to 802.11).

To better illustrate the shift in relative performance among

**Table 2: Protocol parameters for Directed-Transmission and Gossiping**

<p>1. <b>Directed-Transmission parameter:</b></p> <p>(a) <b>MAC protocol:</b> CSMA (tests with 802.11 showed that CSMA consistently outperforms 802.11)</p> <p>(b) <b>Tunable parameter (<math>k_1</math>):</b> 30.0, 10.0, 3.0, 1.0, 0.3, 0.1, 0.03, 0.01</p> <p>(c) <b>Duplicate handling (<math>k_3</math>):</b> <math>\infty - 1</math>, i.e., all duplicate packets dropped (tests with other values showed dropping is always a superior strategy)</p> <p>(d) <b>Information propagation:</b> beacon with update period 120, 22.5, 7, 3.5, 1.75, 0.875 seconds.</p> <p>(e) <b>Control information validity timeout:</b> set to same value as beacon update frequency</p> <p>2. <b>Gossiping parameter:</b></p> <p>(a) <b>MAC protocol:</b> CSMA (tests with 802.11 showed that CSMA consistently outperforms 802.11)</p> <p>(b) <b>Tunable parameter (forwarding probability):</b> 0.00.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.91.0 (value of 1.0 corresponds to Flooding)</p> <p>(c) <b>Duplicate handling:</b> all duplicate packets dropped (tests showed that other values lead to inferior performance)</p> <p>(d) <b>Information propagation method:</b> not needed</p>
--

protocols for the different scenarios, we define the utility of a method by combining both load and reliability into one variable called *goodness*:

$$goodness = (reliability) - (relative\ load) \quad (5)$$

*Relative load* is the energy used during a session scaled by energy used when a full flood is performed (may be greater than 1 in case of a packet long wandering around). Therefore, the higher the goodness, the “better” the method (where both a unit of relative load and a unit of reliability have the same weight). Naturally, information will be lost in going to a scalar, but we are primarily interested in the best and worst, where both load and reliability are extremal, not the trade-off region where the relative weighting of load and reliability would matter.

## 6. RESULTS

Due to space limits, we only present a sample of the most interesting results; additional plots are available from the authors. We evaluate the effect of mobility in a separate subsection before discussing the results for the three different application scenarios in detail. The effect of the number of packets per session is shown only in the context of the any-to-any scenario, and the effect of the number of sessions is shown only in the context of the single destination. Except for Figure 4, all figures show experiments of 1000 packets with other parameters shown on the axes or as separate curves.

### Node Mobility:

For only one destination and no mobility at all, Shortest-Path is clearly the best protocol: the one beacon at the

**Table 3: Protocol parameters for Shortest-Path and AODV**

<p>1. <b>Shortest-Path parameter:</b></p> <p>(a) <b>MAC protocol:</b> CSMA, 802.11</p> <p>(b) <b>Duplicate handling:</b> duplicate packets never dropped (preliminary tests showed that other values lead to inferior performance)</p> <p>(c) <b>Information propagation:</b> beacon update period 120, 22.5, 7, 3.5, 1.75, 0.875 seconds</p> <p>(d) <b>Control information validity timeout:</b> set to same value as beacon update frequency</p> <p>2. <b>AODV parameter:</b> set according to values recommended in draft 9 of the AODV protocol</p> <p>(a) <b>MAC protocol:</b> 802.11 (CSMA was consistently outperformed by 802.11)</p> <p>(b) <b>Active Route Timeout:</b> 1, 3</p> <p>(c) <b>Route Request Retries:</b> 2, 4, 16</p> <p>(d) <b>Route Deletion Constant:</b> 1, 5, 25</p> <p>(e) <b>Local Repair:</b> yes, no</p> <p>(f) <b>Search Better Route:</b> yes, no</p>
---

very beginning disseminates precise information about the distances, and this information will not become inaccurate. Therefore, all subsequent data traffic will be routed along shortest paths. The cost of the first beacon (flooding) needs to be recouped in the cost of data traffic. Therefore, as the number of sessions simulated increases, the advantage of Shortest-Path becomes more apparent. This is illustrated in the left panel of Figure 2 for the single-destination scenario with 1000 sessions and no mobility (and in Figure 4 for nonzero mobility and varying numbers of sessions).

There is another feature of the left panel of Figure 2 that requires explanation, associated with the Directed-Transmission protocol. As the parameter  $k_1$  is increased from zero (a full flood in the upper right of the plot) to large values, the load decreases monotonically. However, the delivery fraction is bitonic with respect to the parameter value (i.e., informally speaking, there is a bump in the curve). In these experiments, all duplicate packets are dropped unconditionally. Because of the realistic MAC layer, it sometimes happens that the packet that arrives first at an intermediate node does not follow the shortest path to that node from source. Directed-Transmission penalizes packets that do not travel on the shortest path ( $\Delta + \delta > 0$ ), so this packet has a lower chance to survive all future retransmission decisions. When the shortest-path packet arrives at the node in consideration (later than the non-shortest-path packet), it is dropped because it is a duplicate. When the  $k_1$  parameter is low, there are many packets and the retransmission probabilities are always very close to one; then this effect is unimportant. When the  $k_1$  parameter is very high, there are very few simultaneous packet copies, so the chance of a “not-shortest-path” packet arriving at a node before a “shortest-path-packet” is very low. Finally, when the value of  $k_1$  is in between, the above feature becomes prominent. In short, this is a MAC layer issue for this routing protocol. When we changed it to TDMA with enough time slots for each node to have its own frequency (therefore no collisions) this effect disappeared.

As the mobility increases, the methods relying heavily on the accuracy of the distance information (Shortest-Path and AODV) quickly become worse. The cost of keeping the information sufficiently accurate becomes too high. The use of back-up copies of the packets on the network becomes advantageous. Directed-Transmission is the best performer for a wide range of speeds, beating both methods that rely more on the network information (Shortest-Path and AODV) and methods that use no network information at all (Gossiping and Flooding). See the right panel of Figure 2.

As the mobility increases even more, all information is inaccurate and misleading, and the cost of keeping it usable, even for Directed-Transmission, is too high. In this region, Gossiping (or Flooding) is the method of choice. (Recall that Directed-Transmission with  $k_1 = 0$  is Flooding; see below for more detailed explanations.)

### Single-destination Scenario:

The single-destination scenario results have already been used to illustrate the general effects of mobility in Figures 2. Those are plots of load versus reliability, from which the “goodness” of a protocol can be determined for a particular scenario (Equation 5). Because the goodness is a linear combination of the measurements on the x- and y-axes of the load-reliability plots, the goodness is essentially the L1 distance from the lower right corner. Each experiment (point in the load-reliability plots) produces a goodness value, and the value plotted in the goodness plot (Figure 3) will be the highest goodness among all parameter values for a protocol that was explored under the given conditions. Thus the plotted goodness represents the best a protocol can do in those circumstances. Left panel of Figure 3 shows the resulting goodness with respect to mobility for the various protocols in the single-destination scenario.

As seen in left panel of Figure 3, in the case of no mobility, Shortest-Path is the method of choice, but as the mobility increases, Directed-Transmission soon becomes the best, finally overthrown by Gossiping for very high values of speed. The slope of the curves in this plot represents the sensitivity to misinformation. The lower the slope, the more rapidly the method deteriorates with increasing mobility (and more imprecise information), and therefore the more sensitive to misinformation. (AODV’s Goodness is negative for speeds above 50 m/s in this scenario, and not shown in this figure.)

The cost of keeping information accurate depends, in part, on the number of sessions: the more active sessions during a certain time period, the more the information is used for a given beacon frequency, and the lower the cost of the beacon. This is illustrated in Figure 4. In the Shortest-Path case, when there is mobility, there are two separate curves that do not connect. The left part (lower fraction delivered) is from a CSMA MAC layer, and the right (higher fraction delivered) is 802.11. Both are shown because it is difficult to determine which is better. These results show that with an increased number of simultaneous sessions, AODV decreases relative load per session, but also decreases the overall reliability. Gossiping increases relative load while reducing reliability. Directed-Transmission remains nearly the same. Shortest-Path with either CSMA or 802.11 significantly reduces its relative load, while maintaining its delivery fraction. Because the effect of the number of sessions is similar in all scenarios, and because the few sessions case is not necessarily appropriate for modeling more general traffic, it is only shown for the single-destination scenario.

### Multiple-destination Scenario:

As far as PPR protocols are concerned, the multiple-destination scenario removes one of the biggest obstacles to performance in a single destination/uplink, namely the congestion near the destination. By having multiple equivalent destinations spread throughout the network, each uplinking to the “rest of network”, there is much larger subset of nodes that must be used by every session (those near the destination(s)). Instead, they split their load over all equivalent uplinks.

Middle panel of Figure 3 shows the goodness-mobility for the multiple-destination scenario. The load-reliability results are similar to those in Figure 2 for the single-destination scenario, and are not shown. The trends that are apparent from the plot are all similar to those for the single-destination scenario: Directed-Transmission, independent of mobility, performs better than the other protocols. AODV performs better than Gossiping for low values of mobility, but is worse for medium and high mobility. The transition point where Gossiping performance surpasses AODV is at a higher mobility in multiple-destination networks than in the simple single-destination scenario, because the shorter paths to the destinations make it easier for a single path method to succeed.

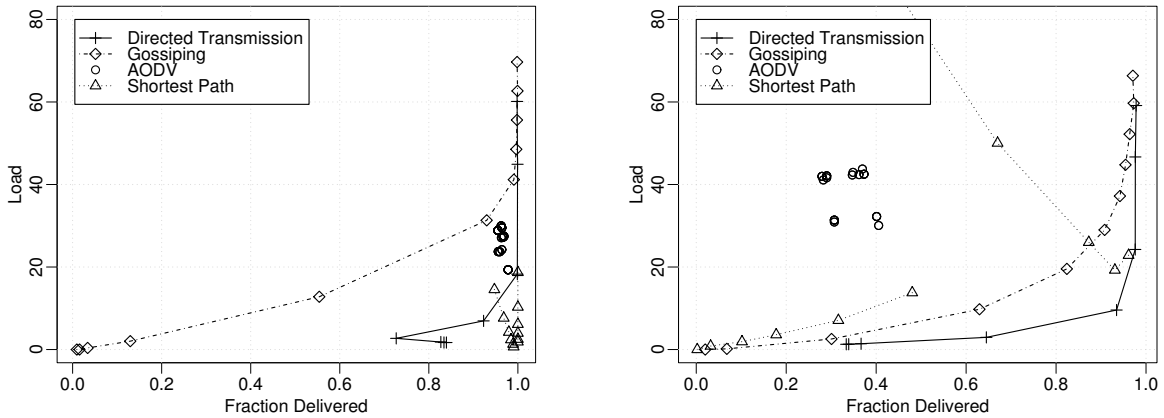
The Shortest-Path method performs poorly in this scenario. The Goodness levels are below zero and are therefore not shown. The reason is that the cost of distributing the distance information (using beacons) is too high for multiple destinations. The same is true for the any-to-any scenario.

### Any-to-any Scenario:


AODV wins most any-to-any scenarios. This is because of its compiling routes on demand, which requires no network information to be stored and updated. It does not win unconditionally, however. We show an example of a class of scenarios where it is only the best protocol for a limited range of low mobility. In this subset of parameter space, the data/control ratio necessary for AODV to maintain its paths surpasses that produced by simple Flooding because the discovered paths are breaking too frequently.

Unlike the single-destination scenario, the any-to-any scenario has two additional parameters: the number of packets in a session and the size of the packets. We did experiments on both of those parameters and found that for most reasonable combinations, AODV performed significantly better than the other protocols. Some results are shown in Figure 5. In these experiments there were 1000 sessions and no mobility. It can be seen that for sessions with only one small packet, the overhead induced by AODV in route discovery is such that a Gossip flood can have both lower load and higher delivery fraction. When this is increased to three packets per session, both protocols retain about the same reliability, but the Gossiping load nearly triples. Above ten packets per session, Gossiping has reached a congestion limit, where there are too many packets competing for limited resources. It is not possible to flood every packet throughout the network: the load has peaked at the maximum possible, and reliability continues to drop with increased session sizes.

Right panel of Figure 3 shows the results of our experiments on 100 10-packet sessions, which have the same number of unique packets on the network as previous experiments, where the congestion limits shown in Figure 5 were



**Figure 2: Load-Reliability for two levels of mobility in the single-destination scenario: the left panel has no mobility, and the right averages 60 m/s. Note that both Gossiping and Directed-Transmission improve with increased mobility, while AODV and Shortest-Path get worse.**



**Figure 3: Goodness vs. Mobility plots in the single-destination, multiple-destination and any-to-any scenarios (from left to right, respectively). Only the best achieved Goodness is shown for any level of mobility**

not reached. The transition point where AODV's control overhead makes it less efficient than simple Gossiping occurs just over 50 m/s. (Note that this speed is not necessarily meaningful; the result also depends strongly on the density of nodes, network load and average connectivity.)

The Directed-Transmission results are not very favorable in this scenario; it performs significantly worse than Gossiping for all mobilities. This is essentially because of the information dissemination. Because a periodic beacon flood from every node is too big, information is obtained only by piggybacking distances onto packets already being transmitted (method available from the authors). With only 10 packets in a session, and no beacon, very little information can be distributed throughout the network by piggybacking alone, and Directed-Transmission essentially behaves as Gossiping. In cases where there are many more packets in a session, however, piggybacking can work, and the Directed-Transmission performance improves a great deal. This is especially true with small packets.

## 7. CONCLUSIONS

In addition to formalizing the Parametric Probabilistic Routing protocols introduced by [3], through a comprehensive set of realistic simulations, we have demonstrated the importance of considering the dependence of performance on topology changes through mobility, and on the scenario being simulated (Figures 2 and 3). We have also investigated the effect of the total session load (Figure 4) and number of

packets in a session (Figure 5).

## 8. REFERENCES

- [1] Ian F. Akyildiz, Weilian Su, Yogesh Sankarasubramaniam, Erdal Cayirci; *A Survey on Sensor Networks*; IEEE Communications Magazine, August 2002.
- [2] C. Barrett, M. Drozda, A. Marathe and M. Marathe, Characterizing the Interaction Between Routing and MAC Protocols in Ad-hoc Networks. *Proc. 3rd ACM international Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, Lausanne, Switzerland, June 2002.
- [3] C.L. Barrett, S. Eidenbenz, L. Kroc, M.V. Marathe, and J.P. Smith. *Parametric Probabilistic Sensor Network Routing*; Proceedings of The Second ACM International Workshop on Wireless Sensor Networks and Applications (WSNA 2003), 2003.
- [4] C. Bettstetter, H. Hartenstein, X. Pérez-Costa. Stochastic Properties of the Random Waypoint Mobility Model. accepted for *ACM/Kluwer Wireless Networks Special Issue on Modeling & Analysis of Mobile Networks*, to appear in 2004
- [5] J. Broch, D. Maltz, D. Johnson, Y. Hu, and J. Jetcheva. Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols. *Proc. 4th Annual ACM/IEEE International Conference on*

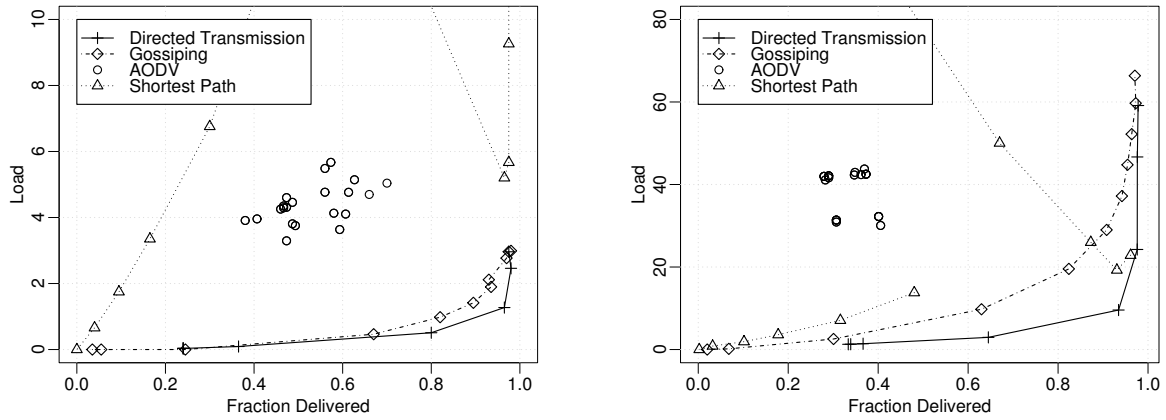


Figure 4: Load with respect to Reliability for mobility with average speed of 60 in the single-destination scenario. The left panel has 50 and the right 1000 sessions.

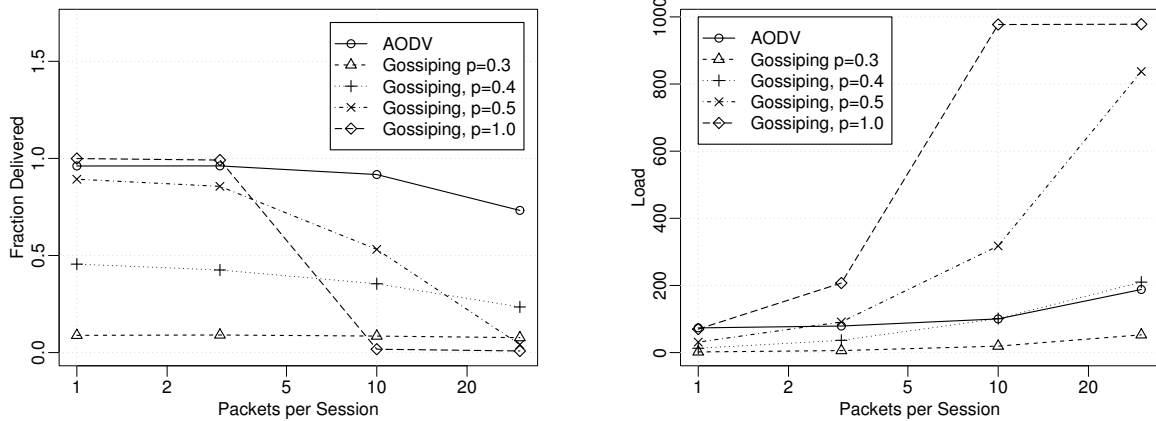


Figure 5: Reliability and Load with respect to the number of packets in each session for AODV and Gossiping protocols in the any-to-any scenario.

*Mobile Computing and Networking*, ACM, Dallas, TX, October 1998.

- [6] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, D. Terry; *Epidemic algorithms for replicated database maintenance*; Proceedings ACM Symposium on Principles of Distributed Computing (PODC), pp. 1-12, 1987.
- [7] Deepak Ganesan, Ramesh Govindan, Scott Shenker, Deborah Estrin; *Highly Resilient, Energy Efficient Multipath Routing in Wireless Sensor Networks*; Mobile Computing and Communications Review, Vol. 1, No. 2.,
- [8] Zygmunt Haas, Joseph Y. Halpern, Li Li; *Gossip-Based Ad Hoc Routing*; Proceedings of INFOCOM 2002, 2002.
- [9] David Kempe, Jon Kleinberg, Alan Demers; *Spatial Gossip and Resource Location Protocols*; Proceedings of 33rd ACM Symposium on the Theory of Computing (STOC), 2001.
- [10] Haiyun Luo, Ramachandran Ramjee, Prasun Sinha, Li Li, Songwu Lu, UCAN: A Unified Cellular and Ad-Hoc Network Architecture, ACM MOBICOM'03, p.353, 2003.
- [11] Jun Luo, Patrick Th. Eugster, Jean-Pierre Hubaux; *Route Driven Gossip: Probabilistic Reliable Multicast in Ad Hoc Networks*; Proceedings of INFOCOM 2003, 2003.
- [12] M. Mauve, J. Widmer and H. Hartenstein, A Survey on Position Based Routing in Mobile Ad-hoc Networks, *IEEE Network*, Nov. 2001.
- [13] Charles E. Perkins and Elizabeth M. Royer; *Ad hoc On-Demand Distance Vector Routing*; Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications, New Orleans, 1999.
- [14] S. Ramanathan and M. Steenstrup, A survey of Routing Techniques for Mobile Communication Networks. *Mobile Networks and Applications*, Vol. 1-2, 1996, pp. 89-104.
- [15] *Routing in Communication Networks*, M. Steenstrup, Editor, Prentice Hall, New Jersey, 1995
- [16] A. Vahdat and D. Becker. Epidemic Routing for Partially Connected Ad Hoc Networks. Technical Report CS-200006, Duke University, April 2000.
- [17] draft-08 and draft-09 of AODV are available from <http://moment.cs.ucsb.edu/AODV/aodv.html>
- [18] QualNet, <http://www.scalable-networks.com/>.