



Parametric Probabilistic Routing in Sensor Networks

CHRISTOPHER L. BARRETT, STEPHAN J. EIDENBENZ, LUKAS KROC, MADHAV MARATHE
and JAMES P. SMITH

CCS-5, Los Alamos National Laboratory, MS M997, P.O. Box 1663, Los Alamos, NM 87545

Abstract. Motivated by realistic sensor network scenarios that have mis-informed nodes and variable network topologies, we propose an approach to routing that combines the best features of limited-flooding and information-sensitive path-finding protocols into a reliable, low-power method that can make delivery guarantees independent of parameter values or information noise levels. We introduce *Parametric Probabilistic Sensor Network Routing Protocols*, a family of light-weight and robust multi-path routing protocols for sensor networks in which an intermediate sensor decides to forward a message with a probability that depends on various parameters, such as the distance of the sensor to the destination, the distance of the source sensor to the destination, or the number of hops a packet has already traveled. We propose two protocol variants of this family and compare the new methods to other probabilistic and deterministic protocols, namely constant-probability gossiping, uncontrolled flooding, random wandering, shortest path routing (and a variation), and a load-spreading shortest-path protocol inspired by (Servetto and Barrenechea, 2002). We consider sensor networks where a sensor's knowledge of the local or global information is uncertain (parametrically noised) due to sensor mobility, and investigate the trade-off between robustness of the protocol as measured by quality of service (in particular, successful delivery rate and delivery lag) and use of resources (total network load). Our results for networks with randomly placed nodes and realistic urban networks with varying density show that the multi-path protocols are less sensitive to misinformation, and suggest that in the presence of noisy data, a limited flooding strategy will actually perform better and use fewer resources than an attempted single-path routing strategy, with the *Parametric Probabilistic Sensor Network Routing Protocols* outperforming other protocols. Our results also suggest that protocols using network information perform better than protocols that do not, even in the presence of strong noise.

Keywords: sensor networks, probabilistic routing, simulation

1. Introduction

1.1. Motivation

As the deployment of large-scale sensor networks is imminent, the need for efficient communication protocols that govern the interaction of the sensors with each other continues to be of utmost importance. A sensor network consists of a large number of mobile devices that measure an environmental variable such as temperature, pressure, radioactivity, or geographic position. The sensors are equipped with transceivers that enable them to transmit and receive data to and from other sensors. In our model, most sensors only have a very small transmission range, while a single sensor is equipped with a powerful transmitter device (along with a more powerful battery) that enables it to forward data to a fixed base station. The need for communication between a regular sensor, called source, and the powerful sensor, called destination can arise at any time possibly triggered by unexpected changes in the environment variable or by a timer; we assume that a communication session consists of sending a single packet of data to the powerful sensor who then forwards it to the base station; this assumption is reasonable as the data of an environmental variable only consist of a few bytes. If a source sends out its data packet, chances are that the packet will not reach the destination in a single hop due to the typically very limited transmission range of the source; in order to obtain a functioning system, other sensors that receive the packet transmitted by the source need to forward the data

packet themselves until it reaches the destination. This model of communication results in a multi-hop path from source to destination with intermediate sensors forwarding the packet.

The problem of how to route a data packet has been the subject of active research in wire-line and wireless settings. The unique requirements of a sensor network setting (e.g., limited battery power, high network node density, and frequent network topology changes) have also led to proposed routing protocols specifically tuned for sensor networks such as SMECN, Flooding, Gossiping, SPIN, SAR, and LEACH, see [1] for an overview. Traditionally, if a source and a destination is given, most routing protocols—particularly in ad hoc networks—first compute a routing path and then route data packets along the computed path. In our setting, this approach seems to be somewhat overkill as we only need to send a single packet to the destination. In fact, a straight-forward flooding approach (i.e., every sensor who receives the packet retransmits it to all its neighbors) would solve our routing problem, but at the cost of involving all sensors in every transmission and thus unnecessarily depleting their batteries. A wandering approach alleviates this problem: in the WANDERER protocol, a sensor that receives a packet forwards the protocol to only one of its neighbors, which it chooses according to some given probability distribution. Such an approach, however, may severely increase path lengths. A third approach, PURE GOSSIP tries to combine the advantages of FLOODING and WANDERER by deciding at each node whether to forward the packet to all

neighboring nodes with a given probability. A well-known result from percolation theory (see [6] for details) says such an approach will result in either almost all nodes receiving the packet or almost no nodes receiving the packet.

1.2. Our contribution

In this paper, we propose a family of routing protocols called Parametric Probabilistic Sensor Network Routing protocols, which substantially improve the performance of controlled flooding methods like gossiping by making the probability of retransmission a function of several parameters. Parametric Probabilistic Sensor Network Routing protocols are completely described by the retransmission probability function: each node in the network—upon receiving a packet—retransmits the packet to its neighbors with a certain probability according to the probability function. While this basic principle is straight-forward, we must be careful not to overload the probability function: the retransmission probability could depend on parameters as diverse as the numbers of copies of the packet already received by a certain node or the distance to the destination. In order to keep the routing protocol as light-weight as possible, we try to keep the retransmission probability as simple as possible with respect to the complexity it takes to compute or estimate the parameters that the function depends on. We propose two variations of Parametric Probabilistic Sensor Network Routing: The retransmission probability function of our first variation DESTINATION ATTRACTOR depends on the distance counted in hops from the source node to the destination node and the hop-distance from the node currently holding a packet to the destination; in our second variation DIRECTED TRANSMISSION, the probability depends on the same two factors and additionally on the number of hops that the packet has already traveled.

We measure the performance of DESTINATION ATTRACTOR and DIRECTED TRANSMISSION against each other as well as against a number of other proposed sensor network routing protocols through simulation in different settings. These four settings are:

- Random distribution of nodes in a square with uniform transmission ranges resulting in an average node degree of 6.7 (i.e., a sensor has 6.7 neighbors on average) and 13.7 respectively,
- Node distribution that reflects possible deployment of sensors in an urban environment with uniform transmission ranges resulting in average node degrees of 6.7 and 13.7 respectively.

The other protocols we consider are FLOODING, PURE GOSSIP, and WANDERER as described above. We also compare against a SHORTEST PATH protocol that always reaches the destination in a minimum number of hops and a variation called SHORT PATH. The SHORTEST PATH protocol serves as a theoretical optimum that no real-life protocol can ever match: the large number of sensors and their level of mobil-

ity make the cost of updating routing tables prohibitive. It is included to investigate the degree to which our Parametric Probabilistic Sensor Network Routing protocols achieve this optimum performance. Finally, we compare against a protocol we call SHORTEST PATH COUNTING, in which a packet is sent to a neighbor who is closer to the destination and sits on a shortest path to the destination; the probability of choosing a specific neighbor depends on the number of shortest paths on which the neighbor sits. SHORTEST PATH COUNTING is a generalization of an earlier protocol [14] that works on general graphs (rather than only on grids as in [14]). Most of the methods we use¹ rely on global information that changes as the sensors change positions. As exact and continuously updated computation of this information would incur prohibitively high overhead in these protocols, estimator methods are required. Rather than implementing a particular estimator method, we model sensor mobility through inaccurate data: the parameter values that the sensors use are “noised”.

In our simulations, we investigate the trade-off between quality of service measures (in particular the fraction of times that at least one packet reaches the destination and the average lag before a packet reaches the destination), load incurred in the network (measured as the number of times that a sensor emits a packet summed over all time steps and nodes) for different levels of noise. The key finding of our simulations is that between the two protocols from the Parametric Probabilistic Sensor Network Routing family, DIRECTED TRANSMISSION outperforms DESTINATION ATTRACTOR on all levels of noise achieving better quality of service results while incurring less load in the network. Both Parametric Probabilistic Sensor Network Routing protocols in turn clearly outperform PURE GOSSIP on all noise levels despite the fact that PURE GOSSIP is not sensitive to noise at all as it does not rely on global information. These three multi-path protocols incur more load in the network than the single-path protocols, but only for low levels of noise; for high noise levels, single path algorithms such as SHORT PATH can break leading to excessive lag times and even load, while delivering only a small fraction. Among the single-path protocols, SHORT PATH outperforms SHORTEST PATH COUNTING, while WANDERER and SHORTEST PATH come in last. Our results thus suggest that multi-path algorithms are more robust to mis-information, and can actually have a higher delivery rate, lower lag and lower load under the same conditions. Surprisingly, these results are largely independent of the chosen setting (dense vs. sparse network, realistic distribution vs. random distribution).

1.3. Related work

Routing in wireless and ad hoc networks has been the subject of active research over the last two decades, see [15] for the state of the art until 1995 and more recent survey papers

¹Namely, DESTINATION ATTRACTOR, DIRECTED TRANSMISSION, SHORTEST PATH, SHORT PATH and SHORTEST PATH COUNTING.

[3, 12, 13]. The tutorial by N. Vaidya at MobiHoc 2002 also provides an excellent upto date survey of recent protocols. As argued in [10], one possible reason for researchers proposing a large number of routing protocols is that efficient routing aims to optimize a variety of measures, including, path lengths, load balancing, robustness, etc. Here we concern ourselves with the following design constraints that are particular to sensor networks: (i) low energy requirements, (ii) simple and local routing methods, (iii) robustness under node/edge failures. Probabilistic routing is a simple method that designed correctly satisfies all the above constraints.

Probabilistic routing and gossiping in particular have been applied in various setting, such as spreading distributed database updates [4] and multicast techniques (see [11] for a recent result and an overview). Previous work on gossip usually assumes a complete graph and random set of unicast or multicast messages. Here, we make no assumptions on the structure of the underlying network. Here, we are concerned with wireless communication and thus assume omnidirectional transceivers that either forward a message to all of their neighbors within range or to none of them. This model has been introduced by Haas et al. [6], where the authors propose to use gossiping as a means to reduce flooding in the route discovery phase of ad hoc routing protocols such as AODV. Most previous work on gossiping assumes a constant probability for forwarding messages and tries to exploit the bimodal behavior of gossiping. Non-constant probability gossiping has been developed by Kempe et al. [7], 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.² Multiple path routing has been applied in [5] in order to make a network resilient against node failures. While Parametric Probabilistic Sensor Network Routing protocols proposed in this paper also make a network resilient against node failures, we do not explicitly compute alternative paths (as is done in [5]) but rely on statistical properties of the retransmission probability function; also, Parametric Probabilistic Sensor Network Routing protocols incur very little overhead.

The SHORTEST PATH COUNTING protocol is inspired by the recent work of Servetto, Barrenechea [14]. In that paper, the authors considered the problem of routing in sensor networks, when the sensors are placed on a uniform 2D-grid and sensors can fail with certain probability. This model is closely related to the classical work on percolation theory, with the important distinction that we are now interested in the length of paths between a given source and destination. The authors leave open the design of probabilistic local methods for unstructured networks. Our work is a step in this direction. In our work, we model network irregularities such as node failures or mobility through noise.

²The model consists of n sensors positioned at lattice points of a $\sqrt{n} \times \sqrt{n}$ region of the plane. An underlying mechanism that supports a point-to-point communication is assumed.

1.4. Organization

The remainder of this paper is organized as follows. Section 2 presents the Parametric Probabilistic Sensor Network Routing protocols. We briefly present the other protocols that we consider in Section 3, and present a method for estimating global information in Section 4. Section 5 presents the experimental design; Section 6 contains a discussion of our simulation results. Concluding remarks and directions for future research can be found in Section 7.

2. Parametric Probabilistic Sensor Network Routing protocols

Parametric Probabilistic Sensor Network Routing protocols apply a limited-flooding strategy. The key idea underlying our protocols is that the retransmission probability for a packet at a sensor node is a function of various parameters rather than a constant as used in previous gossiping approaches such as in [6]. We will show that this allows us to further limit flooding (and thus further reduce the load of the network) as compared to constant-probability gossip solutions.

Let S denote the source sensor (the origin of data), D denote the final destination sensor (the one we want to send the data to); in time step i , all sensors who received a packet in the previous time step $i - 1$ decide whether to retransmit the packet and, if affirmative, do so; let R_i denote a sensor holding a packet at time step i , and let R_{i-1} denote the sensor that sent the packet to R_i (if R_i received the packet from more than one sensor, R_{i-1} is defined to be any one of them).

The same sensor can hold a packet on different time steps, and the same packet can be on multiple sensors at the same time. Source S emits its packet in time step 0, thus $S = R_0$. We denote the hop-distance between two sensors, say S and D , by $d(S, D)$.

The retransmission probability of a protocol from the Parametric Probabilistic Sensor Network Routing family could depend on various parameters, such as $d(S, D)$, $d(S, R_i)$, $d(S, R_{i-1})$, $d(R_i, D)$, $d(R_{i-1}, D)$, the number of neighbors of R_i or R_{i-1} , or the number of received copies of the same packet. Using subsets of these parameters, we choose two forms that yield qualitatively different protocols that we call DESTINATION ATTRACTOR and DIRECTED TRANSMISSION.

2.1. Destination attractor

The DESTINATION ATTRACTOR evolves from a very simple and sensible concept: If the packet is getting closer to the destination, then its retransmission probability is increased; if it is getting further away from the destination, then the retransmission probability is reduced. For a packet at the i th node R_i in its path from the source, this probability P_{R_i} can




Figure 1. Retransmission Probability contours for DESTINATION ATTRACTOR and DIRECTED TRANSMISSION, where the source is located at coordinates (0,0) and the destination at (1,0).

be represented in a recursive manner as

$$P_{R_i} = \begin{cases} (1+k) P_{R_{i-1}} & \text{closer to destination} \\ (1-k) P_{R_{i-1}} & \text{further from destination} \\ P_{R_{i-1}} & \text{same or indeterminate} \end{cases} \quad (1)$$

where k is the relative amount by which the probability should be increased or decreased. The three cases of equation (1) can be combined if we use the change in the distance to destination $d(R_{i-1}, D) - d(R_i, D)$ in the last time step as an additional factor. In the case of accurate information, $d(R_{i-1}, D) - d(R_i, D)$ always has a value of ± 1 or zero, as in Equation (1). Thus,

$$P_{R_i} = \{1 + k[d(R_{i-1}, D) - d(R_i, D)]\} P_{R_{i-1}} \\ \approx \exp\{k[d(R_{i-1}, D) - d(R_i, D)]\} P_{R_{i-1}} \quad (2)$$

The second line of Equation (2) is derived from the similarity to the low order terms of a Taylor series expansion of the exponential function ($\exp(x) \sim 1 + x$). We arrive at a non-recursive form that is more amenable to analysis by evaluating the recursive terms back to the source $S = R_0$ and by using the fact that the sum of a contracting sequence is the difference of its end-points:

$$P_{R_i} = \exp \left\{ k \sum_{j=1}^i [d(R_{j-1}, D) - d(R_j, D)] \right\} \\ = \exp\{k[d(S, D) - d(R_i, D)]\} \quad (3)$$

Obviously, $P_{R_i} > 1$ is possible (in which case the probability is one) if the packet moves closer to the destination over several time steps; this is intended to allow a packet to build up “positive energy”, which can only be lost over a series of time steps moving farther away from the destination.

Thus, the retransmission probability for DESTINATION ATTRACTOR depends on the hop-distance of the source to the destination as well as the distance of the node currently holding a packet to the destination. The local values of the retransmission probability function for a graph obtained by placing infinitely many points uniformly at random in a 1×1 region of the plane³ are illustrated in figure 1 (left panel). In the fig-

ure the destination is located in the center of the butte and the source node is located at any place on the rim of the butte. This explains the name DESTINATION ATTRACTOR. Parameter k models the exponential decay constant whose increase will cause the contour to more closely resemble a pedestal with vertical sides; decreasing k will spread the tails of the contour. In both cases, the probability within the circle of distance $d(S, D)$ around the destination will always be unity. However, this is all done in the presence of perfect information. In the presence of noisy data, the contour in figure 1 (left panel) would acquire mesas extending out from the main pedestal in the areas of missing information. It would also naturally change in the presence of mis-information. Parameter k can be used to help alleviate problems caused by mis-information: decreasing k makes DESTINATION ATTRACTOR behave more like flooding. Thus, for a particular sensor network setting (for example sensor with excessive mobility in a tornado) we can tune parameter k so as to still guarantee a required QoS-measure.

2.2. Directed transmission

The retransmission probability function that we use for the DIRECTED TRANSMISSION protocol is based on the idea that the nodes that lie on a shortest path from the source to the destination should forward packets with a very high probability and the farther away a node is from the shortest path, the smaller its retransmission probability should be. Let R'_i denote any node that holds a packet at time i . The node minimizing $\min_{R'_i} d(R'_i, D)$ is the node closest to the destination. Thus, the retransmission probability P_{R_i} of node R_i holding a packet at time i is defined as:

$$P_{R_i} = \exp \left\{ k \left[\min_{R'_i} d(R'_i, D) - d(R_i, D) \right] \right\}$$

We have chosen an exponential form in order to remain amenable to analysis. Since we assume that the network topology does not change dramatically during a transmission (and if it does, we model it through noise), there is at least one packet that travels along a shortest path from source to destination; moreover, we know that such a packet will take

³We will call this an infinitely dense geometric graph.

$d(S, D)$ time steps to reach destination as it advances one hop in each time step. Thus:

$$\min_{R'_i} d(R'_i, D) = d(S, D) - i$$

This allows us to write the retransmission probability in a simpler form:

$$P_{R_i} = \exp\{k[d(S, D) - d(R_i, D) - i]\}$$

Even though their motivation is very different, DIRECTED TRANSMISSION and DESTINATION ATTRACTOR use the same two parameters (i.e., $d(S, D)$ and $d(R_i, D)$); in addition, the DIRECTED TRANSMISSION protocol uses the number of time steps i that the packet has traveled so far.

With an infinitely dense geometric graph, the local values of the retransmission probability function for DIRECTED TRANSMISSION is illustrated in figure 1 (right panel), where the destination is located on one end of the ridge and the source node is located at the other end of the ridge. Since the protocol aims to forward packets in the direction of the destination, we call it DIRECTED TRANSMISSION. Decreasing k will cause the shape to have less steep walls. Parameter k can be tuned by a system designer to achieve a certain QoS-level.

3. Other protocols and overview

This section describes the other routing protocols used in our simulations.

Wanderer

The simplest routing protocol that we use is called WANDERER, which is simply a random walk on the network of possible radio connections starting from the source. When transmitting a packet, the sender chooses only one of its neighbors to be the receiver. All neighbors receive the packet (since we assumed omni-directional radios), but only the neighbor specified as the receiver will retransmit. WANDERER does not use global information, but each node must maintain a list of its neighbors in order to choose one randomly. As a node may receive the same packet multiple times, the node keeps a record of how many times it has forwarded a particular packet. It may make sense to only forward the same packet a certain number of times; this number, called duplicate-forwarding parameter, can be set by a system designer.

Short and 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. Since this method breaks down in the presence of noise, we introduce a method we call SHORT PATH: instead of choosing the next node exclusively from those that are closest to the destination, a sensor node R_i chooses the next sensor node uniformly at random among

all the sensor nodes who claim to be closer to the destination than the present node R_i . For these methods, a system designer can set the duplicate-forwarding parameter, which makes sense for SHORTEST PATH only if we have noise.⁴

Shortest path counting

The SHORTEST PATH COUNTING protocol is an extension of SHORT PATH: sensor R_i currently holding the packet chooses to forward to one of its neighbors, where the probability of forwarding to neighbor R_{i+1} (with liberal use of notation) is a function of the number of shortest paths from the source or to the destination on which R_{i+1} lies. This was inspired by the work of Servetto and Barranechea [14]. Our modification is in the edge weighting, which is generalized to arbitrary graphs. The probability $P(R_i, R_{i+1})$ that sensor R_i chooses to forward to sensor R_{i+1} is

$$P(R_i, R_{i+1}) = \begin{cases} \frac{x_{i+1}}{\sum x_j} & \text{if } d(R_i, D) \leq d(R_i, S) \\ \frac{1/y_{i+1}}{\sum 1/y_j} & \text{if } d(R_i, D) > d(R_i, S) \end{cases} \quad (4)$$

where x_i is the number of shortest paths to the destination from sensor R_{i+1} and y_i is the number of shortest paths from the source to sensor R_{i+1} , and the sum is over all neighbors of sensor R_i . The design goals from [14] are different from ours: [14] designs a protocol that spreads the load across all equivalent paths (over many packet transmissions) on a grid topology; we want a protocol to be robust to arbitrary network topology and misinformation. In the absence of noise, this protocol performs as well as the shortest-path with the added feature of load-spreading-on-average, but having only one packet on the network at any time makes it inherently sensitive to misinformation. SHORTEST PATH COUNTING has again the duplicate-forwarding parameter, which can be set by a system designer.

Flooding

FLOODING is a protocol that requires no network knowledge and has no parameters: every node that receives a packet retransmits it once to all of its neighbors. Subsequent copies will be dropped, and the result is that every node on the network in the same component as the source will transmit the packet one time.

Gossiping

A simple parametric extension of FLOODING is PURE Gossip, which assigns a retransmission probability to a packet at the source, and every sensor to receive the packet will choose to retransmit with that probability. It tends to have a percolation behavior, in that for a given retransmission probability it is most likely that either very few devices on the network will

⁴Setting the duplicate-forwarding parameter prevents a packet from looping in the network infinitely long, which may happen in the presence of noise; if a node does no longer forwards a packet because of the parameter, the destination is not reached but on the positive side the network load incurred is limited.

Table 1
Protocol overview.

| | Uses network information | Multipath |
|-----------------------|--------------------------|-----------|
| WANDERER | ○ | ○ |
| SHORTEST PATH | ● | ○ |
| SHORT PATH | ● | ○ |
| FLOODING | ○ | ● |
| PURE GOSSIP | ○ | ● |
| DESTINATION ATTRACTOR | ● | ● |
| DIRECTED TRANSMISSION | ● | ● |

receive the packet, or most devices will receive the packet. The retransmission probability p is a parameter that can be set by a system designer.

All protocols that we consider can be classified as to whether they use network information and whether they are single-path or multi-path protocols. Table 1 gives an overview.

4. Estimating global information

The retransmission probability functions for DESTINATION ATTRACTOR and DIRECTED TRANSMISSION depend on knowing the hop-distance $d(R_i, D)$ from the node R_i to the destination D , distance $d(S, D)$ from source to destination, and time step i . The time step information is very simple to compute: we add a field into the header of the packet that contains the number of hops that it has already taken; this value is initialized to zero by the source and is incremented by every node that retransmits the packet. Computing hop-distances is harder: as the nodes are mobile, the distance values may change over time, but we assume that these changes will not be very substantial during the delivery of a single packet to the destination. Nevertheless, we cannot assume perfect knowledge as computing this distance takes several time steps such that the computed information might be outdated once it reaches the nodes. No matter what method we choose to estimate these distances, we may have wrong values. In our simulations, we model this issue by introducing noise such that a node will be likely to either underestimate or overestimate its current distance to the destination. Having global information and thus the tedious task to find estimates adds algorithmic complexity to our protocol, which is the price we pay for obtaining improved performance. Several methods could be developed to estimate $d(R_i, D)$. We outline only one of them that we believe to be one of the most light-weight methods possible. Other estimators may yield better results; in fact, finding improved estimator methods is a promising direction for future research.

A light-weight method of estimating the distance to the destination requires each sensor to include its current estimate of its distance to the destination in each packet that it retransmits; the destination “acknowledges” each packet it receives by sending out a packet containing the value zero in the distance-to-destination field (this acknowledgement packet is

only received by the destination’s direct neighbors; they do not forward the acknowledgement). Whenever a sensor R receives a packet from one of its neighbors R' , sensor R checks whether its current distance estimate for $d(R, D)$ is more than one unit larger than the distance estimate $d(R', D)$ of its neighbor R' ; if this is the case, sensor R updates its estimate by setting $d(R, D) = d(R', D) + 1$. In order to allow a sensor to increase its distance estimate, we assign each information a time-to-live T after which the sensor will assume it to be invalid and reset its estimate to the minimum of the most recently obtained: thus, the distance estimate of a sensor is one unit larger than the minimum of the estimates of its neighbors that the sensor received in most recent T time steps. Using this method, correct distance information propagates from the destination to the other nodes, where the wave front advances by one hop for each packet that the destination acknowledges: a sensor at (correct) distance d from the destination will be able to update its estimate to the correct value after retransmitting (or sending) d packets, if the network topology does not change during this time. Changes in the network topology get propagated equally slowly, however, if network traffic is not too infrequent, the update speed is still sufficient. Initially, the nodes set their distance estimate to some constant.

Rather than implementing this particular estimator method, we have opted to model the mobility and the inevitably resulting imperfectness of the distance information with noise in our simulations. The SHORTEST PATH COUNTING protocol uses number of shortest paths as parameter, which would require a different estimator but can be modeled through noise as well.

5. Experimental design

In order to assess the performance of the routing protocols described above, we analyze the results of simulations that we conducted under various levels of mis-information (noise) and through the entire reasonable parameter range. The value of the tunable parameter k are given as follows:

- for DESTINATION ATTRACTOR and DIRECTED TRANSMISSION $k \in \{.001, .01, .022, .046, .1, .22, .46, 1.0, 10.0, 100.0\}$,
- for PURE GOSSIP probabilities of retransmission are: 20%, 25%, 30%, 35%, 40%, 42.5%, 45%, 47.5%, 50%, 55%, 60%, 70%, 80%, 90% and
- for the duplicate-forwarding parameter for the single-path methods we used values of 0 (no retransmission of the same packet), 3, 10, 30, 100, 300, and infinity (no duplicate-dropping).

5.1. Setup

We use the following setup in our experiments for the different settings: In the random-distribution settings, we spread




Figure 3. Area of Portland where sensors were placed. (a) Portland road network with the area of interest marked by a square. (b) A zoomed-in view of the area of interest.

5000 sensors uniformly at random in a square field, where all sensors have equal range of transmission.

In order to generate spatial distributions in an urban setting, we used a section of downtown Portland, OR, measuring $2900 \text{ m} \times 2950 \text{ m}$. The particular area of Portland chosen for our study is shown in figure 3(a); an enlarged view of the area is shown in figure 3(b). Spatial distributions of sensors are obtained by placing them along a roadway system. The specific distance between consecutive sensors was simply obtained by running a traffic simulation program (TRANSIMS) developed at the Los Alamos National Laboratory and measuring the average distance between cars at different time instants [2]. (More uniform distributions can also be used and are currently being investigated.) A total of 1750 sensors were used to construct the sensor network.

We set the transmission range such that a communication graph with average node degree of approximately 6.7 or approximately 13.7 is induced. The lower average connectivity (i.e., 6.7) is chosen so that there is a huge connected component in the random-distribution settings, while there is still an interesting non-uniform structure in the network [8]. The longer transmission range (and thus higher average connectivity) makes the realistic downtown network almost fully connected, and results in a network that is about twice as dense.

Figure 4 illustrates the connectivity graph of an instance of a realistic (non-uniform) urban sensor network where each vehicle with a driver and all pedestrians on the streets carry a sensor device [2]. The figure depicts the communication graph that results from choosing transmission radii such that an average degree of 6.7 (left panel) respectively 13.7 (right panel) results. Choosing these two average degrees allows us to keep the experiments comparable to the random placement setting. Figure 5 shows the degree distribution for both average degrees for networks resulting from randomly placed nodes and networks resulting from the realistic setting. Clearly, very different degree distributions result.

For each simulation run, we place the destination in the middle of the sensor field, and we randomly choose a source sensor among all remaining sensors connected to the destination. We generate one packet at the source that needs to get to the destination (using the tested protocol) with no competing traffic on the network. In the random-distribution settings, each network and source-destination pair are used for only one simulation run, and 1000 runs are made for each scenario to obtain statistically sound results. In the realistic settings, the network stays the same for all runs, and 400 runs are made for each scenario using different source-destination pairs. A scenario consists of a protocol, a parameter value, a noise level, and a transmission range. We let each simulation run as long as there are packets on the network, or up to 5000 time steps (resp. 1750 for the realistic settings)—this only happens with single-path algorithms. It turns out to be necessary to impose this limit because the packet may travel in loops in the network when the information is inaccurate. In the simulations, we ignore the specifics of all layers in the protocol stack other than the network layer, where the routing takes place.

All multiple-packet protocols (i.e., FLOODING, PURE GOSSIP, DESTINATION ATTRACTOR, DIRECTED TRANSMISSION) use a duplicate-dropping functionality in our simulations: if a packet is received a second time, it is discarded unconditionally.⁵ For the single-path methods (i.e., SHORT PATH, SHORTEST PATH, SHORTEST PATH COUNTING, and WANDERER), the value of the duplicate-forwarding parameter determines this behavior.

5.2. Inaccuracy of information

An important aspect of our study is the dependence of the performance of a protocol on the accuracy of information it

⁵This requires that the sensors keep a list of packets already forwarded; in reality, such a list can be of limited constant length.




Figure 4. Realistic urban sensor network: communication graph of 1750 mobile devices in 8.5 km² of downtown Portland with an average degree of 6.7 (left panel) and 13.7 (right panel)




Figure 5. Degree distributions for random node placement and realistic placement for average degree 6.7 (left panel) and average degree 13.7 (right panel).

has available (for algorithms that use information about the network). We therefore test the protocols for several levels of “noise values”. The network information (distance to destination, etc.) is computed off-line in our simulations. The model used to add noise is such that each value representing a piece of information is replaced by an (integer) number randomly chosen from a uniform distribution with mean equal to the original (accurate) value and width of the interval equal to certain percentage of the original value. We use the width of the interval as parameter to model different noise-levels. The noise-levels we choose for each protocol are 0%, 3%, 10%, 30%, 100%, and 300% (values are bounded at 0, so that 300% noise is actually +300%/–100%). The noise is persistent in the sense that the information available to a sensor does not change with time. This is equivalent to the sensors not updating their information about the network in the course of one source-to-destination transmission. Higher noise-levels could also infer higher mobility levels: the faster nodes move, the faster and the more severely distance estimators become outdated, for all sensible (that is, reasonably light-weight) estimator methods.

5.3. Performance measures

We assess the performance of a protocol with a given parameter and a specified noise level by looking at the following measures:

- Load:** number of times any sensor transmits a packet (sum across all time steps and all sensors)
- Lag:** number of time steps it takes to deliver the (first) packet to the destination
- Fraction delivered:** fraction of runs where the destination receives the data (at least one packet copy reaches the destination)

5.4. Plot explanation

We present most of our results in plots of the type given in the illustrative figure 2. The x -axis depicts the measure fraction delivered and the y -axis stands for the average load, where averages are taken over all runs. A data point in this plot represents the average load and fraction delivered for a specific protocol with parameters set to specific values. The parameter

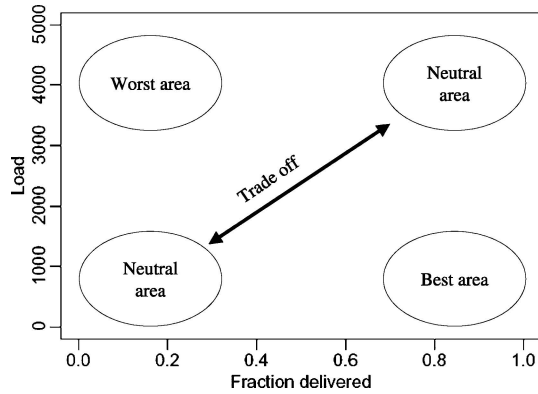


Figure 2. Trade-offs between performance measures.

settings result in a specific value of fraction delivered. For protocols with a tunable parameter, we usually obtain a curve in this plot, where we connect data points corresponding to adjacent parameter values. The optimal point in figure 2 is the lower right corner with very low load and 100% delivery; the worst point is the upper left corner with very high load and zero delivery; the points on the diagonal from lower right corner to upper right corner represent trade-offs between the two measures of load and fraction delivered. The closer a protocol lies to the lower right corner, the better it performs.

6. Comparative simulation results

We present our simulation results for the different settings in this section. We choose to present the results for random node

placement with average degree 6.7 in most detail. We restrict ourselves to pointing out the differences in behavior for the other settings.

6.1. Random node placement with average degree 6.7

In order to illustrate how the multi-path protocols DIRECTED TRANSMISSION, DESTINATION ATTRACTOR, and PURE GOSSIP work, we show snapshots of a run of these protocols on the same network in figures 6–8, respectively. The network is picked from the ones we used for the simulations and no noise is present in the information available to the sensors. The source is marked by S and destination by D . Time progresses from left to right, time steps 10, 20 and 30 are shown. Black nodes are sensors that have a packet at that time step, grey nodes are sensors without packets. In the chosen run, DIRECTED TRANSMISSION imposes less load to the network than DESTINATION ATTRACTOR, which in turn is lighter than PURE GOSSIP. As we shall see below, this ranking is actually representative for the three multi-path protocols. For animations showing all time steps of the simulation, see [9].

Protocol comparison

We now compare the algorithms' performance with respect to lag, fraction delivered, and load. For this comparison, we fix the value of noise to 30% (we will discuss the influence of noise on the protocols later). Figure 9 shows how the algorithms compare, where the x -axis depicts the measure fraction delivered and the y -axis stands for the average load




Figure 6. Sample snapshot of DIRECTED TRANSMISSION: This method does not leave any packets on the network by time step 20. It issues a cloud of packets and moves it towards the destination, keeping its spread roughly constant.




Figure 7. Sample snapshot of DESTINATION ATTRACTOR: Limited flooding is performed within a circle of radius roughly equal to the original distance of S and D (no flooding is done in the opposite direction to D). The wave front passes through D and slowly disappears as it gets further from D .



Figure 8. Sample snapshot of PURE GOSSIP: Its wave front propagates in all directions, eventually reaching all the sensors. It uses no information about the network to control the spread, although the extent of spreading is controlled by a parameter.

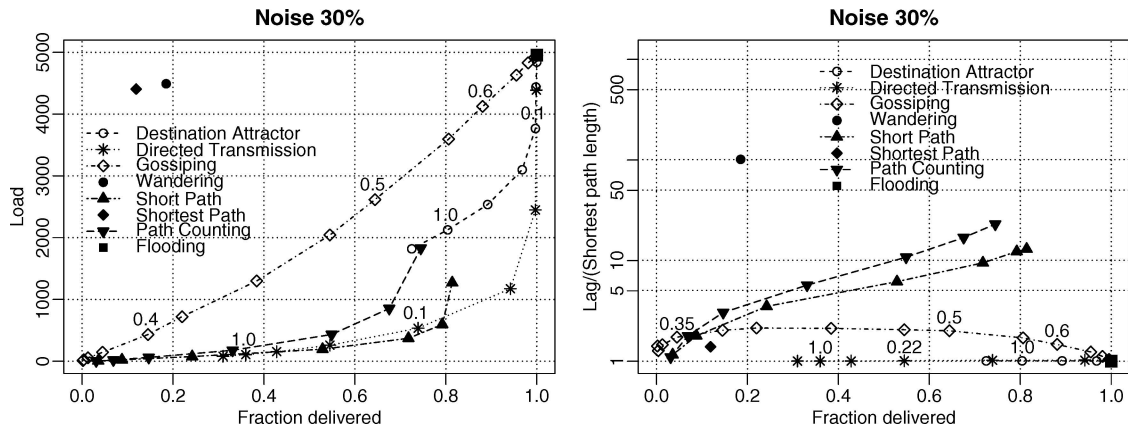


Figure 9. Load and Lag vs. Fraction Delivered for each algorithm (and various values of its parameter, if applicable) for noisiness of 30% on random networks with average degree 6.7. Note the logarithmic scale on the lag axis. The lag is divided by length of the shortest path from source to destination for each run to make the runs comparable.

(left panel) and average lag divided by the length of the shortest path (right panel), where averages are taken over the 1000 runs.

For all tunable methods, the tunable parameter determines the measure fraction delivered and these protocols are thus shown as a line going through points which correspond to the values of the parameter used in the simulations. The parameter values increase from left to right for PURE GOSSIP and decrease from left to right for DIRECTED TRANSMISSION and DESTINATION ATTRACTOR. For the single-path methods, which all have the duplicate forwarding parameter, the values increase from 0 (for no duplicate-forwarding) to infinity (for forwarding the same packet as often as it is received) from left to right. In order to keep figure 9 readable, we only show the value for infinity for WANDERER and SHORTEST PATH.

Recall that the optimal point in figure 9 (left panel) is the lower right corner, with very low load and 100% delivery and that the upper left corner represents the worst performance in both delivery and load. For the single-path methods, not surprisingly, WANDERER is in the poor performance region (it usually does not reach the destination, and wanders for the 5000 time steps until the simulation is stopped⁶) and FLOOD-

ING is in the upper right corner (delivery is guaranteed because all sensors receive the packet, but that imposes maximum load). The poor performance of SHORTEST PATH is somewhat surprising but indicates that this method is highly sensitive to noise in information (with perfect information, it is at the optimal point). For small values of the duplicate-forwarding parameter, SHORTEST PATH drops down to the origin very quickly (not shown in the figure).

SHORTEST PATH COUNTING performs adequately for this noise level and actually beats PURE GOSSIP; interestingly, the curve for SHORTEST PATH COUNTING seems to end where the curve for DESTINATION ATTRACTOR starts. SHORT PATH performs very well on this noise level; however, the highest fraction delivered achievable is at roughly 80%.

The fact that all lines lead from the upper right corner to the bottom left corner represents the trade-off between average delivery and average load as the value of the parameter is changed. If we restrict ourselves to the high-delivery part of the graph (75% delivery and higher), we can conclude that with 30% noise, DIRECTED TRANSMISSION is better than DESTINATION ATTRACTOR which is in turn better than PURE GOSSIP in the sense that for any given parameter to a method, there exists a value of a parameter for the other method such that it outperforms the first in *both* load and delivery. In the same sense, it can also be concluded that for

⁶If the duplicate-forwarding parameter is set to a small constant rather than infinity, the load decreases (not shown in the figure).

this level of noise, all single-path methods perform worse than DIRECTED TRANSMISSION. For the low-delivery part of the graph, DIRECTED TRANSMISSION and SHORT PATH perform almost equally well.

The right panel of figure 9 allows us to extend the same observations about the relative performance ordering of DIRECTED TRANSMISSION, DESTINATION ATTRACTOR and PURE GOSSIP to the lag measurements. Again, DIRECTED TRANSMISSION and DESTINATION ATTRACTOR outperform PURE GOSSIP. The single-path protocols are clearly inferior even if the duplicate-forwarding parameter is set to infinity as shown in the figure.

Dependence on noise level

No algorithm in our comparison in figure 9 reaches the optimal point, which is due to the noise in information available to the sensors. As we will see in the next plots, perfect information allows some methods to reach this point, but as the noise level increases, it is not possible to remain there. Methods that do not use any information about the network (WANDERER, FLOODING and PURE GOSSIP) are totally insensitive to noise, and their positions in the plot do not change with noise. Thus, when studying how the algorithms perform for different levels of noise levels, we can restrict ourselves to methods that actually make use of information: DIRECTED TRANSMISSION, DESTINATION ATTRACTOR, SHORTEST PATH COUNTING, SHORTEST PATH and SHORT PATH. Figure 10 illustrates this dependence: the x -axis again stands for fraction delivered and the y -axis stands for the average load, but the results are shown for different noise levels for DIRECTED TRANSMISSION (left panel), DESTINATION ATTRACTOR (middle panel), and SHORT PATH (right panel). The single-path methods SHORTEST PATH COUNTING, SHORTEST PATH and SHORT PATH all perform similarly with respect to changing noise: they move from the lower right to the upper left corner as noise increases. The speed of this transition (i.e., the sensitivity to noise) is different: SHORTEST PATH is most sensitive to noise, followed by SHORTEST PATH COUNTING, followed by SHORT

PATH, which is the only method shown in the right panel of figure 10.

Figure 10 shows that as the noise level increases, all methods shift in the direction from the optimal point towards the worst region. Both Parametric Probabilistic Sensor Network Routing methods move closer to the graph of PURE GOSSIP (which remains stationary), but retain their relative performance ordering. The speed at which particular protocols move away from the optimal point depends critically on how much and how efficiently they use the information available. The more efficiently information is used, the better the algorithm performs when the information is accurate, but the more rapidly it deteriorates when the information becomes inaccurate. This effect is most pronounced when comparing single-path to multiple-path methods. Single-path methods only rely on a single packet traveling in the network, which can easily get lost or take a long time to reach the destination if the routing decisions are based on unreliable information. In this way, they can impose a larger average load than multipath methods, which illustrates the trade-off between sensitivity to noise and the other measures, fraction delivered and load.

Tunable parameter dependence

As a final study, we investigate in more detail how the tunable parameter determines the QoS measure fraction delivered. Figure 11 shows this dependence for DIRECTED TRANSMISSION, DESTINATION ATTRACTOR, and PURE GOSSIP, where the y -axis is the concrete parameter value. The dependence is *monotonic* in all three cases, making it easy for a system designer to search for the parameter value resulting in the desired QoS level. The curves for all three methods have a flat region around a certain parameter value. This shows that close to that value, the probability of delivery is quite sensitive to the value of the parameter, which corresponds to the behavior of PURE GOSSIP as reported in [6]: for these critical values, either almost all of the sensors receive the packet or hardly any of the sensors receive the packet. As can be seen from the figure, this effect is less pronounced for DIRECTED




Figure 10. Load vs. Fraction delivered for noise levels of 0%, 3%, 10%, 30%, 100% and 300% for three different algorithms on random networks with average degree 6.7.




Figure 11. Parameter value vs. Fraction delivered for noise levels of 0%, 3%, 10%, 30%, 100% and 300% for three different algorithms on random networks with average degree 6.7. Note the logarithmic scale on the parameter axis for DIRECTED TRANSMISSION and DESTINATION ATTRACTOR.

TRANSMISSION and DESTINATION ATTRACTOR than it is for PURE GOSSIP.

Figure 11 also shows that for any given level of noise, DIRECTED TRANSMISSION and DESTINATION ATTRACTOR guarantee non-zero fraction delivered *no matter what the parameter value is*. This is very important, especially in situations when it is difficult to set the correct value of the parameter. Such a claim cannot be made for the case of PURE GOSSIP. We believe that the reason the curve for 300% crosses the other curves in the case of DIRECTED TRANSMISSION is that this level of noise causes the average distance to be overestimated (as explained in Section 5) and thus the method generates more packets (resulting in better fraction of packets delivered, but also higher load). A similar tendency can be observed for DESTINATION ATTRACTOR, but not as strong.

In summary, our results show that DIRECTED TRANSMISSION and DESTINATION ATTRACTOR outperform all other protocols for a given quality of service level and network load incurred if a certain noise level is exceeded (30%). In fact, if we look at three characteristics of a protocol, namely load, fraction delivered, and sensitivity to noise, we see that to optimize any two is simple: FLOODING optimizes fraction delivered and is insensitive to noise, SHORTEST PATH optimizes fraction delivered and minimizes load, time-limited WANDERER achieves low load and is insensitive to noise. Optimizing all three goals is a balancing act, and we suspect that no algorithm fully achieving all three goals simultaneously can exist. A great advantage of the tunable methods is that they allow a system designer to set the parameter so that the resulting performance meets the need of the application. As long as such a need lies on the curve plotted in the graph, the trade-off between performance in load and fraction delivered can be controlled. Among the tunable methods, DIRECTED TRANSMISSION achieves the best balance for these three goals.

6.2. Random node placement with average degree 13.7

Figure 12 shows how the different protocols perform with 30% noise level on randomly generated networks of average degree 13.7. When compared to figure 9 (left panel) for the low density network, figure 12 shows that DIRECTED




Figure 12. Load vs. Fraction Delivered for random node placement with average degree 13.7 for each algorithm for noise level 30%.

TRANSMISSION and DESTINATION ATTRACTOR essentially remain at the same level of performance with slight improvements for both protocols. The curves for PURE GOSSIP also look similar in both figures, however in the denser network a gossiping probability of about 30% already suffices to reach a delivery fraction of almost one. WANDERER and SHORTEST PATH slightly improve their position in a dense network, but remain uncompetitive. SHORTEST PATH COUNTING and SHORT PATH on the other hand highly profit from the increased density of the network delivering 100% of the messages without incurring any significant network load, and thus they both beat the Parametric Probabilistic Sensor Network Routing protocols. The reason for this is that in a dense network there are many paths leading to the destination, and it is therefore more difficult for the packet to get lost. Also, the paths are shorter which again results in a smaller chance of misforwarding the packet.

We do not show figures illustrating parameter values required to achieve a certain QoS-level for dense networks (corresponding to figure 11), as figure 12 already illustrates the main point: the parameter values to achieve the same QoS-level remain almost constant for DESTINATION ATTRACTOR and DIRECTED TRANSMISSION; however, they change dramatically for PURE GOSSIP, which suggests that




Figure 13. Load vs. Fraction delivered for noise levels of 0%, 3%, 10%, 30%, 100% and 300%, for three different algorithms on random networks with average degree 13.7.

the Parametric Probabilistic Sensor Network Routing protocols are more stable in this sense as well.

The plots depicting the dependence of protocol performance on noise level are shown in figure 13 for DIRECTED TRANSMISSION, DESTINATION ATTRACTOR, and SHORT PATH. When compared to their counterparts in sparser networks of figure 10, it is again remarkable how little DIRECTED TRANSMISSION and DESTINATION ATTRACTOR seem to be affected by the density of the network. SHORT PATH on the other hand clearly improves as density increases.

6.3. Realistic node placement

Figure 14 shows the performance of the different protocols for realistic networks with average degree 6.7 (left panel) and 13.7 (right panel). When compared to their corresponding plots for random placement (i.e., figure 9 left panel and figure 12), the relative ranking of the different protocols remains roughly the same for both average degrees. The Parametric Probabilistic Sensor Network Routing protocols are very strong for small average degree, but are beaten by SHORT PATH and SHORTEST PATH COUNTING for large degree. In fact, DIRECTED TRANSMISSION and DESTINATION ATTRACTOR seem to improve their performance relative to randomly placed nodes, thus suggesting that they efficiently exploit the network structure. The opposite is true for their

strongest contenders: SHORT PATH and SHORTEST PATH COUNTING both seem to perform weaker in the realistic setting than in the random setting.

Overall, the results are qualitatively similar to the random placement results. The principal difference lies in the ability of the freeway to guide packets through the network, partly because most nodes are along the freeways.

Figure 15 shows how DIRECTED TRANSMISSION, DESTINATION ATTRACTOR, and SHORT PATH depend on different noise levels in these realistic networks with average degree 6.7. The plots look very similar to their counterparts in random networks in figure 10. All protocols seem to have a slightly worse performance in realistic networks, but the Parametric Probabilistic Sensor Network Routing protocols are still a lot less dependent on the noise level compared to SHORT PATH.

Figure 16 shows how the two best protocols, SHORT PATH (two top panels) and DIRECTED TRANSMISSION (two bottom panels) depend on the average degree in the random node placement setting (left two panels) and in the realistic setting (right two panels). DIRECTED TRANSMISSION is clearly less affected by different average degrees in both the realistic and the random setting. The stability of DIRECTED TRANSMISSION with respect to changes in average degree, network topology (random vs. realistic) and noise level is remarkable. SHORT PATH on the other hand does not cope well

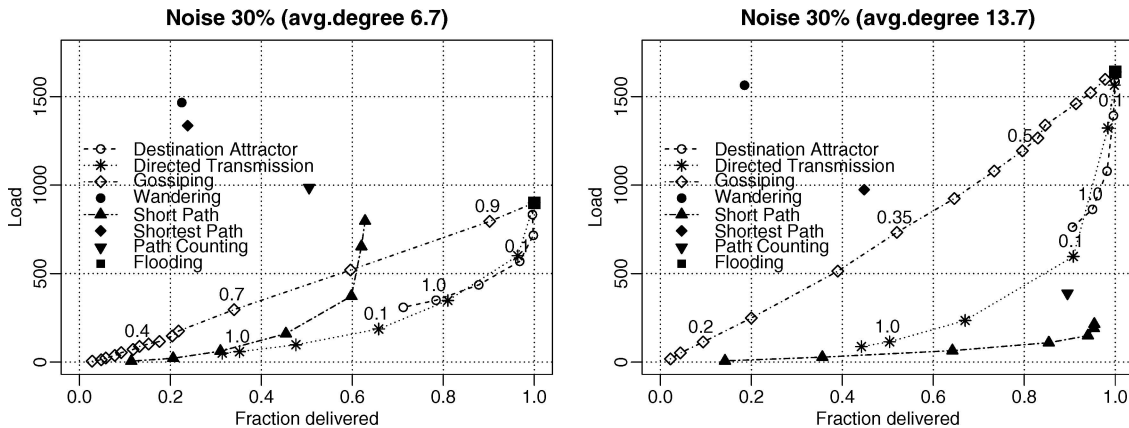


Figure 14. Load vs. Fraction Delivered for noise level 30% on realistic networks with average degree 6.7 (left panel) and average degree 13.7 (right panel).




Figure 15. Load vs. Fraction delivered for noise levels of 0%, 3%, 10%, 30%, 100% and 300%, for three different algorithms on realistic networks with average degree 6.7.

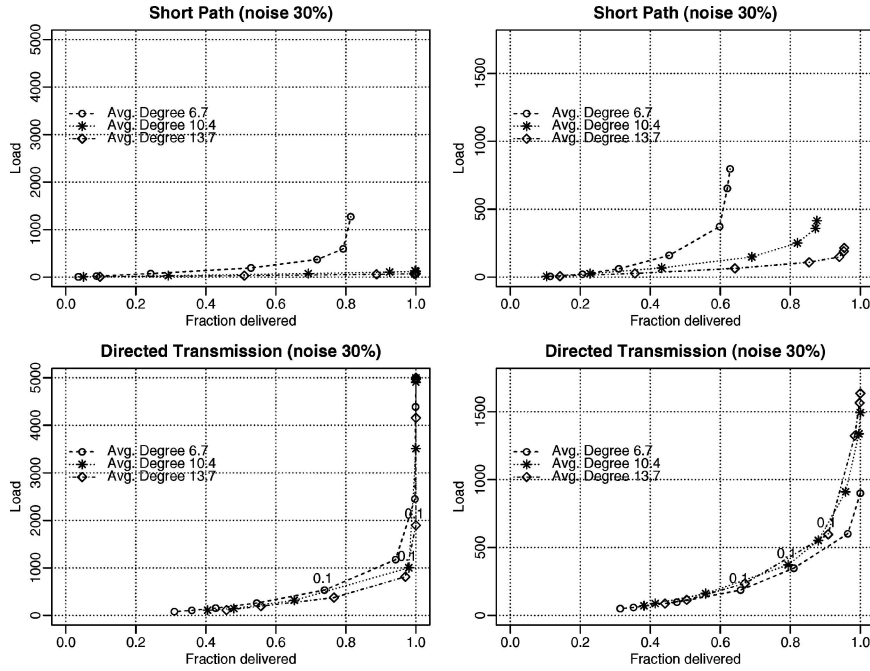


Figure 16. Average degree dependence for SHORT PATH (top half) and DIRECTED TRANSMISSION (bottom half) in random distribution (left half) and realistic distribution (right half).

with smaller average degrees, which is due to the fact that the single path, on which the packet gets routed, is more likely to be the only path that exists in a sparse network: if the packet moves off this path due to mis-information it is more likely to run into a dead-end. In dense networks, many paths exist from a source to a destination and thus misinformation is not as fatal.

7. Conclusion

We have proposed the routing protocol family Parametric Probabilistic Sensor Network Routing, whose members are characterized by a description of their retransmission probability function that a sensor uses to decide whether it forwards a received packet. Our simulations have shown that the two Parametric Probabilistic Sensor Network Routing protocols DESTINATION ATTRACTOR and DIRECTED TRANS-

MISSION can guarantee any quality of service level (measured as fraction delivered) even in the presence of highly noised network information. Our results also suggest that multi-path methods outperform single-path methods in the presence of noise. Our results serve as a proof-of-concept for the basic idea of Parametric Probabilistic Sensor Network Routing: making the retransmission probability depend on network information parameters is a promising step towards more robust routing protocols in sensor networks.

Figure 17 illustrates how the performance measures load, fraction-delivered, and sensitivity to noise of misinformation form a three-dimensional optimization space with the optimum point at the origin. It is straight-forward to find a protocol that optimizes two of the three measures (and these are shown in the figure). The Parametric Probabilistic Sensor Network Routing protocols allow us to trade-off these measures against each other and thus move around on or even underneath the triangle depicted in the figure.




Figure 17. Trade-offs among load, fraction-delivered and sensitivity to misinformation.

There is no reason to believe that DIRECTED TRANSMISSION is the best protocol from the Parametric Probabilistic Sensor Network Routing family. In fact, various points remain open for future research in the area of parametric sensor network routing protocols. Designing a protocol variant that outperforms DIRECTED TRANSMISSION is the most prominent challenge. Changing model assumptions to make the scenarios even more realistic is another research direction that has several aspects. For instance, there are many different noise models that can be tested, especially absolute bounds instead of relative bounds, or an information propagation model, in which noise on individual nodes is correlated to noise on its neighboring nodes. Although mobility can be considered to be a type of mis-information, this should be investigated separately. While we assume that these results apply to arbitrary geometric random graphs, there are some peculiarities of grids that may warrant special attention.

We are currently building a full-fledged simulator for large-scale networks; we plan to implement Parametric Probabilistic Sensor Network Routing protocols in this tool, which will allow us to take into account interactions among different protocol stack layers as well as provide a more realistic framework for comparing methods to estimate global information.

References

- [1] I.F. Akyildiz, W. Su, Y. Sankarasubramanian and E. Cayirci, *A Survey on Sensor Networks*, IEEE Communications Magazine (August 2002).
- [2] C.L. Barrett, M. Drozda, M.V. Marathe, S.S. Ravi and J.P. Smith, *A Mobility and Traffic Generation Framework for Modeling and Simulating Ad Hoc communication networks*, in press, Scientific Programming (2003).
- [3] J. Broch, D. Maltz, D. Johnson, Y. Hu and J. Jetcheva, Performance comparison of multi-hop wireless Ad Hoc network routing protocols, in: *Proc. 4th Annual ACM/IEEE International Conference on Mobile Computing and Networking*, ACM, Dallas, TX (October 1998).
- [4] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart and D. Terry, Epidemic algorithms for replicated database maintenance, in: *Proceedings ACM Symposium on Principles of Distributed Computing (PODC)* (1987) pp. 1–12.
- [5] D. Ganesan, R. Govindan, S. Shenker and D. Estrin, Highly resilient, energy efficient multipath routing in wireless sensor networks, *Mobile Computing and Communications Review* 1(2).
- [6] Z. Haas, J.Y. Halpern and Li Li, Gossip-based Ad Hoc routing, in: *Proceedings of INFOCOM 2002* (2002).
- [7] D. Kempe, J. Kleinberg and A. Demers, Spatial gossip and resource location protocols, in: *Proceedings of 33rd ACM Symposium on the Theory of Computing (STOC)* (2001).
- [8] L. Kleinrock and J. Sylvester, Optimum transmission radii for packet radio networks or why six is a magic number, in: *Proc. IEEE National Telecommunications Conference*, Birmingham, AL (Dec. 1978) pp. 4.3.1–4.3.5.
- [9] <http://www.ms.mff.cuni.cz/~lkro8496/LANL/PPSNR/>.
- [10] F. Kuhn, R. Wattenhofer and A. Zollinger, Worst case optimal and average case efficient geometric Ad-hoc routing, in: *Proc. 4th ACM International Symposium on Mobile Networking and Computing* (2003) pp. 267–278.
- [11] J. Luo, P.T. Eugster and J.-P. Hubaux, Route driven gossip: Probabilistic reliable multicast in Ad Hoc networks, in: *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] S. Ramanathan and M. Steenstrup, A survey of routing techniques for mobile communication networks, *Mobile Networks and Applications* 1/2 (1996) 89–104.
- [14] S.D. Servetto and G. Barranechea, Constrained random walks on random graphs: Routing algorithms for large scale wireless sensor networks, in: *Proceedings of 1st ACM International Workshop and Wireless Sensor Networks and Applications (WSNA 2002)* (2002).
- [15] *Routing in Communication Networks*, ed. M. Steenstrup, Prentice Hall, New Jersey, 1995.



Christopher L. Barrett is leader of the Basic and Applied Simulation Science Group of the Computing and Computational Sciences Division at Los Alamos National Laboratory. His Group is a simulation science and technology (S&T) invention organization of 30 scientists devoted to providing large-scale, high performance methods for systems analysis and simulation-based assisted reasoning. His Group engages in fundamental mathematical, algorithmic, and complex systems analysis research. Current applied research is focused on interdependent simulation and analysis tools for complex, socio-technical systems like transportation, communications, public health and other critical infrastructure areas. His scientific experience is in simulation, scientific computation, algorithm theory and development, system science and control, engineering science, bio-systems analysis, decision science, cognitive human factors, testing and training. His applied science and engineering achievements include, for example, development of large-scale, high performance simulation systems (e.g., Transportation Analysis Simulation System, TRANSIMS) and development of a distributed computing approach for detailed simulation-based study of mobile, packet switched digital communications systems (Self Organizing Stochastic Rebroadcast Relay, SORSRER). He has a M.S. and Ph.D. in Bio-information Systems from California Institute of Technology. He is a decorated Navy veteran having served in both the submarine service and as a pilot. He has been awarded three Distinguished Service Awards from Los Alamos National Laboratory, one from the Alliance for Transportation Research, one from the Royal Institute of Technology, Stockholm, and one from Artificial Life and Robotics, Oita University, Japan.
E-mail: barrett@lanl.gov



Stephan J. Eidenbenz is a technical staff member in the Basic and Applied Simulation Science group (CCS-5) at Los Alamos National Laboratory (LANL). He received an M.Sc. in Computer Science from the Swiss Federal Institute of Technology (ETH) in Zurich in 1997 and a Ph.D. in Computer Science from ETH in 2000; he also obtained a Bachelor's degree in business administration from GSBA in Zurich in 1999. Stephan has worked for McKinsey & Co. in Switzerland, where he received

training in business administration and microeconomics. He has held a post-doctoral position at ETH and he has been a postdoctoral fellow at LANL. Stephan's more than 30 publications cover a wide range of subjects such as approximability and inapproximability properties of visibility problems in polygons and terrains, error modeling in sequencing problems for computation biology, and designing communication protocols robust against selfish behavior. His current research interests include selfish networking, algorithmic game theory, network modeling and simulation, network design, and network optimization.

E-mail: eidenben@lanl.gov



Lukas Kroc is a student of M.Sc. program in Computer Science at Charles University in Prague. In 2003, he was a Graduate Research Assistant at the Basic and Applied Simulation Science group (CCS-5) at Los Alamos National Laboratory. His research interests include simulation, wireless networking and artificial intelligence.

E-mail: kroc@lanl.gov



Madhav V. Marathe is a Team Leader for Mathematics and Computer Science in the Basic and Applied Simulation Science group, Computing and Computational Sciences (CCS-5) at the Los Alamos National Laboratory. He obtained his B.Tech in 1989 in Computer Science and Engg. from IIT Madras, India and his Ph.D. in 1994 in Computer Science, from University at Albany. His team focuses on developing mathematical and computational tools for design and analysis of large

scale simulations of socio-technical and critical infrastructure systems. His research interests are in modeling and simulations of large socio-technical systems, design and analysis of algorithms, computational complexity theory, theory of parallel, distributed and mobile computing and communication systems. He has published over 100 research articles in peer reviewed journals and conferences. He is an adjunct faculty in the Computer Science Department at the University of New Mexico.

E-mail: marathe@lanl.gov



James P. Smith is a technical staff member in the Basic and Applied Simulation Science Group of the Computing and Computational Sciences Division at Los Alamos National Laboratory. His principal interest is in high performance computing applied to modeling, simulation and analysis of socio-technical systems. His current research applies to national infrastructure, especially telecommunication/computing, public health, and transportation.

He has scientific experience in high performance computing and parallel processing applied to large-scale microscopic simulations, including original software design and debugging of very large, evolving systems of inter-operable computational systems, and efficient analysis and synthesis of massive data produced by multi-scale complex environments. Before attending graduate school he worked for a short time in nuclear theory, and had several publications in experimental biophysics from the Pennsylvania Muscle Institute and Bockus Research Institute. During graduate school he took a one year hiatus to start a company to work in analytic finance, and then spent time doing theoretical space physics at LANL. His graduate work eventually included theoretical and experimental fusion research, but concentrated on computational space plasma physics. He has publications in biophysics, analytic finance, education, space plasma physics and computer science, and is a co-inventor on the TRANSIMS patent. He has a Ph.D. in Theoretical Plasma Physics from the University of Texas at Austin.

E-mail: jpsmith@lanl.gov