Dynamic Probabilistic Retransmission in Ad hoc Networks

[1]Donald J. Scott
Computer Science Department
Florida State University
Tallahassee, FL 32306-4530

[2]Alec Yasinsac (Contact Author)
Computer Science Department
Florida State University
Tallahassee, FL 32306-4530

Abstract

Ad hoc networks are becoming more common. In order for these networks to be effective, they need a proper routing protocol. Although many probabilistic protocols have been proposed to optimize the communication between nodes, none of these protocols takes into account the dynamic nature of ad hoc networks. In this paper we propose a routing protocol that dynamically adjusts itself to the local topography of a network in order to provide reliable and efficient interaction between nodes in an ad hoc network. This approach can prevent broadcast storms during flooding in dense networks and can enhance comprehensive delivery in sparse networks

Keywords: ad hoc, gossip, broadcast, routing, MANET

1.0 Introduction

Ad Hoc networks are self-organizing wireless networks, absent any fixed infrastructure. Nodes in such networks communicate through radio transmissions of limited range, sometimes requiring the use of intermediate nodes to reach a destination. Nodes in ad hoc networks are also limited in their power supply and wireless bandwidth. Node mobility further complicates the environment. A communication protocol must take into account not only the restrictions on resources imposed by the nodes in ad hoc networks, but must also be robust enough to deal with the dynamic network topology.

One well-known solution to routing in ad hoc networks is flooding [1]. The flooding protocol works as follows: upon receiving the message for the first time, each node retransmits the message to all neighbors. Although a very simplistic protocol, flooding has the virtue of being reliable, while requiring minimal state retention. Unfortunately, flooding often results in redundant messages, consuming valuable bandwidth and power as well as causing contention, collision, and packet loss. Probabilistic flooding [2, 3, 4.] was introduced to address these problems, but has many problems of its own.

In this paper, we discuss approaches to address problems of broadcast storms in dense networks and message delivery in sparse networks of ad hoc organization. We begin by reviewing probabilistic rebroadcast, or gossip. In the next section, we address density and distribution as they relate to broadcast storms. We follow with a theoretical discussion of optimal retransmission and introduce dynamic, probabilistic rebroadcast in Sections 3 and 4. We summarize and conclude the paper in the final section.

1.1 Probabilistic Broadcast

While flooding is a reliable transmission protocol, it is prone to the broadcast storm problem in dense networks [5, 6]. A simple solution to the broadcast storm problem is to send fewer redundant messages. This is the approach taken in probabilistic retransmission protocols, also referred to as *gossip* protocols. The difficulty in gossip is to know which messages are necessary and which are redundant. In the absence of better information, gossip assumes that all messages are equally important.

The most straightforward version of gossip behaves in much the same way as flooding, beginning with the source node broadcasting the message. Since the source node always sends out this message, we say that the source transmits this

[1]Funded in part by the Department of Defense Information Assurance Scholarship Program (IASP) under Sponsor ID# MDA904-03-1-0205

[2]This material is based upon work supported in part by the U.S. Army Research Laboratory and the U.S. Army Research Office under grant number DAAD19-02-1-0235

message with probability *p* of 1. In gossip, the retransmission probability is set to some number *p* from 0 to 1. When a node receives a message for the first time, it only broadcasts that message with probability *p*, thus reducing the total number of messages sent through the network.

Although gossip results in fewer messages than true flooding, the retransmission probability should be designed to retain the reliability that is inherent in flooding. As discussed in [2], percolation theory has long been discussed in relation to basic gossip protocols where even the source node broadcasts with probability *p*. Clearly, the lower *p* value selected, the fewer redundant messages will be sent. Unfortunately, reducing *p* also reduces delivery reliability.

Haas, et al. detail a theory that gossip exhibits bimodal behavior in sufficiently large networks. This means that for any network and any arbitrary value of *p*, either almost every node in the network will receive the message, or hardly any nodes will receive it. Of course, the goal is to select *p* with the property that all (or nearly all) nodes receive the messages and the number of redundant messages is minimized. For example, work done in [2] identified an optimum *p* value as .65 for their test network. This value allowed for a reduction in total messages sent while still getting the message to most of the nodes in the network.

1.2 Determination of *p*

The critical question thus becomes how to optimally select *p* in the general case. The probability *p* must be great enough that it allows the message to reliably propagate to all regions of the network. Yet, it cannot be too large or the broadcast storm problem resurfaces. The fundamental characteristic of ad hoc networks causes this problem. Since there is no concrete structure, the dynamic topography can be quite varied.

In a fixed network with a moderately uniform distribution, *p* may be selected through a process of trial and error to determine its optimal value. In an ad hoc network, some areas may be dense while others are sparse. Worse yet, the distribution may change constantly when the nodes are mobile. It may be reasonable to use a static value for *p* when gossiping in a static network of a fairly uniform distribution. However, in a mobile network that may experience significant changes in distribution, it makes much more sense to develop a protocol in which *p* can change with the variation in network density. In this paper, we propose a Dynamic Gossip protocol for this purpose.

1.3 Disjoint Broadcast Zones

Research on broadcast storms in ad hoc networks focuses on analyzing broadcast zones for adjacent nodes [5]. The authors suggest mechanisms for reducing redundancy, contention, and collisions in flooding. Section two of that paper motivated some aspects of our work. Specifically, they consider the broadcast areas of two adjacent nodes. Their work focuses on predicting the utility of retransmission by nodes in the disjoint areas after some number of iterations of the message has been received. Though they recognize probabilistic retransmission method, they do not use area analysis in any combination thereof. In the next section, we show how to leverage knowledge about density in the overlap area to establish an effective retransmission probability.

2 Node Density and Distribution

Many factors affect broadcast storms, with one of the most important being network density. It is easy to envision a network where flooding is optimal in the sense that there are very few redundant messages. We show such a network in Figure 1, where the dotted circles represent the transmission range of the nodes. Notice that the only nodes that fall within any transmission overlap are adjacent nodes, and redundant messages could occur only in the end nodes.

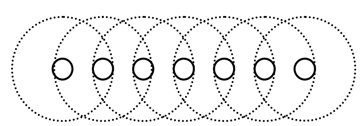

Figure 1. A Flooding-optimal Network

Adding a node within the range of any existing node in Figure 1 would result in redundant retransmissions under a flooding protocol, as we see from Figure 2. There, messages retransmitted from the additional nodes would not propagate to nodes that would not have otherwise received that message. This demonstrates that increases in density can cause redundancy in flood routing.

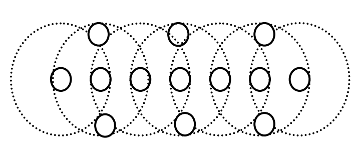

Figure 2. Redundancy in Flooding

While density is a major issue in flooding, it is not the overriding issue; proximity is also important, as we illustrate in Figure 3. There, the density is greater than in Figure 1, and no redundant messages are generated. In fact, only one message is triggered as the network is totally disconnected due to proximity, i.e. though there are more nodes than in Figure 1, but they are positioned just out of the range of one another.

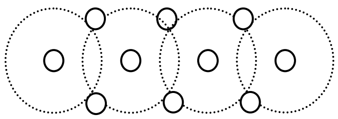

Figure 3. A Disconnected Network

While Figure 1 may be considered a flooding-optimal network, we are more interested in determining the optimal retransmission pattern for an arbitrary network. There is little analytical or deductive foundation for the approaches taken in the literature. Rather, protocol authors test their propositions empirically, most often through computer simulation, as in [2]. While reasoning analytically is difficult in this environment, there are some simple analytic observations that illuminate important properties of retransmission patterns in flood-like protocols.

We begin our analytic argument by making two rigid assumptions that facilitate reasoning about retransmission properties, specifically:

(1) We consider original transmissions and first hop retransmission only, and

(2) That all links are bidirectional (all nodes transmission ranges are the same).

We first note that the maximum coverage area for any first hop retransmission scheme is represented in Figure 4a by the difference between the concentric circles A and B, where A represents the broadcast coverage area of the originating node with broadcast range r, and the radius of B (we term the one-hop broadcast reach of A) is 2r.

We then ask, for the sake of our argument: If only a fixed number of retransmissions are to occur, can we compute the optimal placement for these nodes? Clearly, if there are n retransmissions, the optimum placement of retransmitting nodes is distributed uniformly, equidistant apart (circumference divided by n) around the circumference of the broadcast area of the originating node as illustrated in Figure 4.b-d. In parts b and c, the unreached area is highlighted with darker gray color. Of course, as n increases, the total uncovered area is reduced and as n approaches infinity, the uncovered area disappears.

The utility of this observation lays in our ability to determine the optimal number of desired retransmissions that messages should trigger. If there are no retransmissions, the uncovered area is the difference in the area of circles A and B, which is easily computed as $3\pi r^2$. When there are multiple overlapping areas to consider, the computation becomes more challenging. Rather than engage a precise, but highly complex formula, we chose to approximate and create an upper bound on the uncovered area.

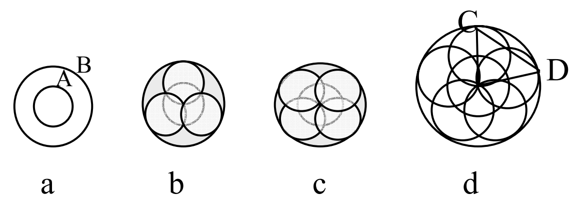

Figure 4

We achieve an upper bound on uncovered area by recognizing that the area bounded by the arc and line ((C,D) in Figure 4.d) between the tangential intersection of the one hop range and two adjacent retransmission areas, is larger than the amount of uncovered area between the two adjacent rebroadcast areas. The area between the arc and line is the difference between the cone with apexes at C and D and the triangle with apexes at C and D. Since there are n such areas, we can derive an upper bound on the total uncovered, one-hop retransmission area (UA) with optimal node placement that is dependent only on r and n. That formula is given in Equation 1.

$$UA = 4r^2*(\pi/n - \sin(\pi/n)*\cos(\pi/n)) \qquad \text{Eq 1}$$

As shown in Table 1, if optimally placed, ten retransmitting nodes cover over 93% of the possible area. Increasing n to fifteen results in 97% coverage, and adding five more retransmitting nodes bumps the coverage to 98.5%. Under optimal conditions, we could use these numbers to determine how many retransmissions need to be generated in order to guarantee comprehensive delivery to all two hop neighbors with high probability. We give a mechanism for statistically ensuring a certain number of retransmissions as we describe Dynamic Gossip in Section 3.

Table 1. Uncovered Area		
r	n	Approximate Percent Uncovered
1	3	58.65
1	4	36.34
1	5	24.32
1	10	6.45
1	15	2.90
1	20	1.64
1	25	1.05
1	100	0.07

3 Dynamic Gossip

While a static retransmission probability may reduce redundancy resulting from flooding in static networks, node mobility renders such static mechanisms ineffective. We propose an extension to static gossip that leverages the positive properties of probabilistic retransmission to reduce flooding redundancy in mobile networks. In Dynamic Gossip, each node determines the rebroadcast probability of each neighbor at transmission time based on local node density. Thus, as nodes move about and the local node density changes, nodes may change the retransmission frequency of their messages accordingly.

3.1 Density Within the Broadcast Overlap

We begin our node distribution discussion by considering networks where nodes are reasonably uniformly distributed. Under this assumption, we can use area measures to statistically derive the number of nodes in the broadcast overlap between any two communicating nodes. The number of overlapping nodes may then be used to determine an optimal gossip retransmission probability. That value may be generated and transmitted with the message by the originator.

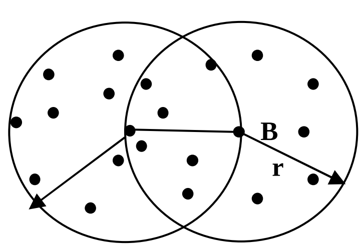

Figure 5.

Consider Figure 5, where nodes A and B are at the limit of each other’s transmission range. There are seven nodes in the transmission overlap of nodes A and B. If we assume the network uses an originator-generated retransmission probability, nodes within the overlap area should retransmit with a probability of .143 (one in seven) in order to expect that the message will be propagated by at least one node within the overlap area.

The one-of-seven ratio is obvious by inspection, but could also be generated mathematically, based on the total area within the broadcast range of the source node, the size of the overlap area of the source and destination nodes, and realization of local density of the source, all of which are easily computed. The computations are founded on the assumption that uniform node distribution results in area computation being proportional to node density, i.e. if there are ten nodes in ten square units of area, then there is exactly one node in one area unit.

We compute the proportional ratio (PR) of the overlap between the source and destination nodes to the total broadcast area of the source for any two identical range, bi-directional nodes using the integral given in Equation 2, where the two nodes lie at the limit of their ranges. For any round broadcast area, that ratio is $.41\pi r^2$ [5].

$$PR = 4\int_{r/2}^{r}\sqrt{r^2 x^2}\,dx \qquad \text{Eq 2}$$

Local node density (LND) can also be easily computed. By measuring the transmission range, nodes can compute their broadcast area, and can determine the total number of nodes in the broadcast area by identifying the number of one hop neighbors that they have. One hop neighbors may be identified in several ways depending on the network needs. For example, nodes may establish density awareness through a "hello" process, where nodes transmit a short message (density ping) on time intervals selected based on node mobility and community need. Alternatively, nodes may acquire on-demand density awareness through a "ping-acknowledge" protocol. Regardless of the counting method, the number of one hop neighbors divided by the broadcast area is the LND.

The number of nodes in the overlap area (NNOA) is the node density multiplied by the proportional relationship between the broadcast area and the overlap area. Thus, from Equation 2, NNOA = LND * .41. As illustrated earlier, using originator-generated gossip probability, the originator can use NNOA to generate the retransmission probability that will give a reasonable expectation of the number of nodes in an overlap area that will retransmit the message. Specifically, the originator generates p = DR / NNOA, where DR is the desired number of retransmissions per overlap area. Thus, if each of NNOA recipients retransmits with probability p, on average, the desired number of retransmissions will be generated.

3.2 Retransmission Distribution and Volume Under Dynamic Gossip

There are two remaining questions for Dynamic Gossip in a uniformly distributed network.

1. What about nodes outside of the overlap area?
2. How can we best optimize retransmissions balancing reliability against redundancy?

The first is a bit subtle, and our assumption of uniform distribution is critical in this analysis. This is because uniform distribution guarantees that the number of overlapping nodes within ANY overlap area of the source will be the same, i.e. if nodes are uniformly distributed, any overlap area created by rotating node A around the circumference of B's transmission range will contain the same number of nodes. Thus with this method, uniformity statistically ensures that the message is retransmitted at the same rate in every direction as long as the transmission is not directionally bounded.

The second question highlights the power of Dynamic Gossip. Using Dynamic Gossip, the total number of first hop retransmissions can be statistically controlled by the message originator, allowing them to decide the tradeoff between reliability and redundancy. When high reliability is needed, more retransmissions can be triggered, yet in even extremely dense networks, the number of retransmissions can be effectively bounded by the originator. This process can be repeated at every retransmission step, reducing message growth from potentially exponential to a statistical expectation that the number of retransmissions is bounded by a constant. These results are attained without requiring nodes to retain state beyond that required in normal flooding.

We illustrate our computations in Table 2. The table shows that for less dense networks, a higher gossip probability is required to generate the same number of retransmissions. For example, in the first entry, if one retransmission is desired with only two nodes in the overlap area, each node must retransmit with probability .5 to expect that one retransmission will occur. That expectation can be strengthened by increasing p.

Table 2. Density-Driven Gossip

LND	nnoa	p = Retrans Probability		
		DR = 1	DR = 5	DR = 10
10	4	0.2500	1	1
100	39	0.0256	0.1282	0.2564
1000	390	0.00256	0.01282	0.02564
10000	3900	0.00026	0.00128	0.00256

In a denser network, retransmission probability is lower, e.g. where there are twelve nodes in the overlap area, to generate five retransmissions, each node need only retransmit with probability .4167. Regardless, by being LND-aware, an originator can statistically control the total number of neighbor retransmissions that occur.

Finally, we note that because the overlap area is approximately 40% of the broadcast area, the total expected number of one hop neighbor retransmissions (TENR) generated by the message is TENR = DR * 2.5.

3.3 Proportional Perimeter Retransmissions

We of course recognize that the nature of ad hoc networks specifically prevents us from selecting optimal placement of retransmitting nodes. Fortunately, geometry can give us additional information regarding proximity. In a uniformly distributed network where we know the network density, we can calculate how many nodes reside within some delta of the perimeter of a node's broadcast area.

Because probabilistic retransmission is location un-aware, nodes near the originator are as likely to retransmit as are those along the broadcast perimeter. Originators can ignore the nearby nodes in their computations by computing the desired retransmission volume based on the occupancy of a rim of some width epsilon around the perimeter. As with our earlier computations, the area in the rim that we select reflects the number of nodes within that area. Thus, we can use the proportion of the area of the rim to the area of the circle to compute the number of nodes in the rim. Given the Local Node Density (LND), broadcast range r, number of desired retransmissions along the rim n, and rim width r*(1 - ε), the formula for computing p is given in Equation 3.

$$p = n / (1 - \varepsilon^2) * LND, \; 0 < \varepsilon < 1 \qquad \text{Eq. 3}$$

We give some illustrative computations in Table 3. When a larger number of retransmissions is desired, p must increase. Notice in the third row, the desired number of retransmissions in the rim cannot be achieved by a single transmission of each node in the broadcast area. Similarly, as the rim collapses (ε is larger), the harder it is to get retransmissions to occur in the rim, so higher values of p must be selected.

Table 3. P For Near-rim Transmissions

p	n	LND	ε^2	ε
0.263158	5	100	0.81	0.9
0.526316	10	100	0.81	0.9
1.052632	20	100	0.81	0.9
0.026316	5	1000	0.81	0.9
0.052632	10	1000	0.81	0.9
0.105263	20	1000	0.81	0.9
0.098039	5	100	0.49	0.7
0.196078	10	100	0.49	0.7
0.392157	20	100	0.49	0.7
0.009804	5	1000	0.49	0.7
0.019608	10	1000	0.49	0.7
0.039216	20	1000	0.49	0.7

The impact of these observations is that an originating node can control many aspects of the reach of one-hop retransmissions. Beyond establishing a statistical bound on the number of one hop retransmissions, given some assumptions regarding node distribution, they can estimate the proximity for a number of re-transmissions.

3.4 State Retention in Dynamic Gossip

All flooding requires state retention to allow nodes to recognize previously received messages. This can be done by utilizing message sequence number or other coordinating information included within the transmission protocol. In the same way as flooding, nodes employing Dynamic Gossip discard all except the first version of any message, thus must maintain state information to allow retransmission recognition. We do not specify which mechanism should be used to detect retransmission, since any mechanism suitable to flooding is also suitable for Dynamic Gossip.

A major advantage of Dynamic Gossip is that, outside the normal flooding state retention, it is essentially stateless. The only information that a node needs for Dynamic Gossip is its transmission range and its local density. Transmission range is inherent to the device and can be retained in any of many simple non-volatile mechanisms. Local density can be efficiently computed on demand when transmission is necessary, though a very low volume of storage is necessary should local density computation be designed for constant retention with intermittent updates.

4 Dynamic Gossip, Density & Distribution

The uniform node distribution assumption is convenient for describing the core principal of Dynamic Gossip. It is unreasonable to expect ad hoc networks in general to have uniform node distribution. Instead, nodes may be distributed randomly, with areas of particularly dense node occupation, other sparse areas, and other areas where node density is near the average. These conditions demand a slightly different approach with Dynamic Gossip.

Consider the complete network shown in Figure 6. Node A identifies their LND as fifteen. If a conservative approach is taken and A desires to limit the number of first hop retransmissions to ten (a one third reduction from flooding), node A will transmit and request the one-hop neighbors retransmit with a probability of .667. In this scenario, it is likely (89%) that node B will receive the message since two nodes within A's range retransmit with probability .667. Node C, however, will not receive the message one third of the time.

This level of delivery reliability may be suitable for some applications, unsuitable in others.

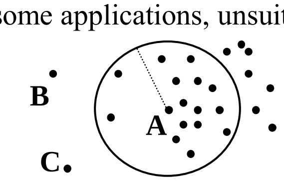

Figure 6. Skewed

Clearly, probabilistic retransmission such as gossip [2] in relatively sparse networks with skewed distribution cannot affect consistent, reliable delivery. Further, it is well known that the probability that each node will receive a retransmission under any probabilistic scheme with nodes randomly distributed is difficult to express theoretically. While simulation can provide important insights regarding behavior of gossip-type schemes, many applications need solid performance and reliability predictors with strong theoretical foundations. In sparse, skewed networks, such computations are elusive.

Recall that probabilistic retransmission evolved as an overhead reduction mechanism. Its behavior in sparse networks is a negative side-effect that may be dealt with in other ways. We propose mechanisms for increasing delivery performance for probabilistic retransmission protocols in sparse networks in the sequel, but focus in this paper on mechanisms to limit overhead while generating a high rate of message delivery in more dense networks. We establish the first theoretical foundations for generating and controlling probabilistic retransmissions to prevent broadcast storms and ensure delivery.

In this paper, we have focused on networks with uniform distribution. In the sequel, we will extend these ideas to networks where other distributions are considered. Specifically, we will propose that our computations are accurate estimates under a random distribution if the skew is bounded, i.e. that minimum connectivity is bounded by a constant relative to the specific network density.

5 Conclusions

In this paper, we consider routing issues presented by the dynamic nature of ad hoc networks. Routing is well understood in fixed networks. These decisions are complicated in mobile ad hoc networks, where portions of the network may be densely populated, while other areas may be sparse or vacant and these attributes are constantly changing.

We present a probabilistic solution that is appropriate to solving problems of broadcast storms in dense mobile networks. We contend that in order to have a truly effective probabilistic protocol for mobile ad hoc networks the protocol must take into account the dynamic nature of the network. Dynamic gossip utilizes a relay ping method to give a local awareness of the density of the network. This information is used to adjust the probability p used in the retransmission decision of the one-hop neighbors.

There is little theoretical foundation for probabilistic retransmission protocols. We provide a framework for reasoning about gossip-type protocols. Our method enables the message originator to compute and implement a statistical bound on the number of retransmissions, essentially eliminating the broadcast storm problem in mobile ad hoc networks that employ this protocol. The method is mathematically founded, though with some restrictive assumptions, it provides a platform to extend discussion to more nearly ad hoc environments.

6 Acknowledgements

Thanks to Sam Huckaba for his insights into mathematical approximations and to Todd McDonald and Willard Thompson for many helpful discussions.

7 References

[1] Charles E. Perkins, Elizabeth M. Royer, and Samir R. Das. IP Flooding in ad hoc networks. Internet draft (draft-ietf-manet-bcast-00.txt), Nov 2001. Work in progress.

[2] Zygmunt J. Haas, Josep Y. Halpern, and Li Li. Gossip-based ad hoc routing. In *IEEE INFOCOM*, Jun 2002.

[3]Yoav Sasson, David Cavin, and Adnre Schiper. Probabilistic Broadcast for Flooding in Wireless Mobile Ad hoc Networks. In Proc. of IEEE WCNC 2003.

[4]J. Cartigny and D. Simplot, "Border Node Retransmission Based Probabilistic Broadcast Protocols in Ad-Hoc Networks." In *Proc. 36th International Hawaii International Conference on System Sciences (HICSS'03)*, Hawaii, USA. 2003.

[5] Sze-Yao Ni, Yu-Chee Tseng, Yuh-Shyan Chen, and Jang-Ping Sheu. "The broadcast storm problem in a mobile ad hoc network", In *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking*, pages 151-162, Aug 1999

[6] B. Karp and H. Kung, Greedy Perimeter Stateless Routing for Wireless Networks, *Proceedings of the 6th International Conference on Mobile Computing and Networking*, Boston, USA, 2000, 243-254.