

Adaptive gossip protocols: managing security and redundancy in dense ad hoc networks^{*}

Mike Burmester, Tri Van Le and Alec Yasinsac

Department of Computer Science, Florida State University
Tallahassee, Florida 323204-4530
{burmester,levan,yasinsac}@cs.fsu.edu

Contact author: Mike Burmester, Tel: (001) (850) 644-6410, Fax: (001) (850) 644-0058

Abstract. Many ad hoc routing algorithms rely on broadcast flooding for location discovery or, more generally, for secure routing applications. Flooding is a robust algorithm but because of its extreme redundancy, it is impractical in dense networks. Indeed in large wireless networks, the use of flooding algorithms may lead to broadcast storms where the number of collisions is so large that it causes system failure. To prevent broadcast storms, many mechanisms that reduce redundant transmissions have been proposed that reduce retransmission overhead either deterministically or probabilistically.

Gossip is a probabilistic algorithm in which packet retransmission is based on the outcome of coin tosses. The retransmission probability can be fixed, dynamic or adaptive. With dynamic gossip, local information is used to decide on retransmission probability. With adaptive gossip, the decision to relay is adjusted adaptively based on the outcome of coin tosses, the local network structure, and the local response to the flooding call. The goal of gossip is to minimize the number of retransmissions, while retaining the main benefits of flooding, *e.g.*, universal coverage, minimal state retention, and path length preservation.

In this paper we consider ways to reduce the number of redundant transmissions in flooding while guaranteeing security. We present several new gossip protocols that exploit local connectivity to adaptively correct propagation failures and protect against Byzantine attacks. A main contribution of this work is that we introduce a cell-grid approach that allows us to analytically prove performance and security protocol properties. The last two gossip protocols that we give are fully adaptive, *i.e.*, they automatically correct all faults and guarantee delivery, the first such protocols to the best of our knowledge.

Keywords Ad hoc networks, secure MANETs, flooding, gossip, broadcast redundancy, broadcast storms, secure routing.

1 Introduction

Ad hoc networks are self-organizing wireless networks, absent of any fixed infrastructure [10, 21, 14]. Nodes in such networks communicate through wireless transmissions of limited range, sometimes requiring the use of intermediate nodes to reach a destination. Also, nodes are usually limited in their power supply and bandwidth. The mobility of the system further complicates

^{*} A preliminary version of this paper appeared in the Proceedings of ADHOCNOW'05, Vancouver. This material is based on work supported in part by the U.S. Army Research Laboratory and the U.S. Research Office under grant number DAAD19-02-1-0235.

the situation. Two primary issues in ad hoc network research are efficiency and security. Because of their nature and restricted resources, efficiency is essential in ad hoc networks. Also, naturally, ad hoc networks are more vulnerable to security threats than fixed, wired networks. Unfortunately, efficiency and security are competing properties, in that improving efficiency is likely to reduce security and efforts to increase security are likely to negatively impact efficiency. The security and efficiency of ad hoc networks is the focus of this paper.

Routing in ad hoc networks is an active area of research [5, 6, 7, 20, 10, 21]. The de facto route discovery algorithm for such networks is broadcast flooding [1–3, 20]. With flooding, each node that receives a message retransmits that message exactly once. Flooding [20] has many positive properties for ad hoc networks including maximal coverage, distance preservation and redundancy. Maximal coverage means that if a time-relevant path¹ exists between a source and any destination, flooding will discover that path. Flooding will also find the shortest path between the source and destination. We call this property distance preservation. Redundancy is a positive attribute in ad hoc networks because these networks are naturally less reliable and more vulnerable than their static counterparts. Conversely, many seek to replace flooding as the ad hoc routing algorithm of choice because of its inefficiency that is directly related to redundancy [16, 13, 22]. Indeed in dense networks, the redundancy may be catastrophic if a broadcast storm [19] is triggered.

A solution to the broadcast storm problem is to reduce message redundancy. This is the approach taken with probabilistic retransmission protocols, also referred to as *gossip* protocols [10, 22, 9]. Gossip is similar to flooding, with one important distinction. In gossip, when a node receives a message for the first time, rather than immediately retransmitting it as in flooding, it engages a probabilistic process to determine whether or not to retransmit. Essentially, it retransmits each message with probability p . From a security point of view, this approach may have undesirable properties. Chief among them is that malicious (Byzantine) nodes are given undue influence in the propagation process, while non-faulty nodes which adhere to the protocol may forego participation. Thus, protocols that may be highly reliable and efficient in a fair environment, will be inactive in the face of a malicious attack. In this paper we describe several subtle adaptations of gossip that, combined with other available information, can offer substantial security enhancement with improved efficiency.

A central contribution of our work is to present a model that supports the analytical argument of protocol properties, addressing both security and performance issues. This model allows us to describe the protocols in a unique way, where properties can be isolated and analyzed. In this paper, for each case, we intuitively describe the proposed protocols and exercise the model by giving mathematical proofs of the protocols' security and performance properties, rather than relying on simulations that cannot capture security properties, and where hidden assumptions can impact performance analysis.

We also note that our methods are founded on the well-understood flooding protocols, which have been extensively simulated in both wired and wireless networks, with mixed results. For example in [7], Gavin *et al.* simulated performance of the simple flooding protocol using each of the three most prominent simulators (OPNET Modeler, NS-2 and GLoMoSim). They showed how each simulator achieved significant variations in results even though the assumptions were held constant. With the extensive simulation analysis that flooding and gossip protocols have undergone, there is little to be gained by further simulation.

¹ Since ad hoc networks are dynamic, a path may form or dissolve during the flooding process. Whether the flooded message finds nodes involved is time dependent.

The rest of this paper is organized as follows. In Section 2 we discuss our security model and malicious faults. In Section 3 we define our cell-grid, the concept of cell-to-cell propagation and connected vertex coverings. We then present two basic gossip protocols. In Section 4 we present four adaptive gossip protocols that correct propagation failures by using local neighbor information. In Section 5 we present a gossip protocol that uses directional information and in Section 6 a gossip protocol that uses cell location. Finally, in Section 7 we discuss security and efficiency issues and conclude in Section 8.

2 Ad hoc faults and malicious faults

There are several ways in which one can model the unpredictable nature of an ad hoc network. For a stochastic approach one may use a Bayesian model in which the status of links tends to be stable (see *e.g.*, [4]). With such an approach one should allow for Markov interdependencies between some links. For example, if A, B are nodes that are close to each other and are on the hop boundary of a node X , then it is more likely that the status of the links (A, X) and (B, X) will be affected in the same way. Such Bayesian models can be used to describe the stochastic aspects of the network and formulate some of the basic properties of ad hoc networks (in particular, for a formal security analysis), but are too general for simulation purposes.

Whatever model is used one must allow for malicious behavior. The traditional Byzantine threats model allows for an adversary who coordinates the malicious nodes according to some plan. The task of the adversary is to frustrate the normal operation of the network. When a link is broken we say that a *fault* occurs. *Ad hoc* faults are random faults that are caused by the mobility of the network and Nature. Such faults are typically independent, although one must allow for certain weak dependencies, *e.g.*, links to nodes that are close to each other are more likely to brake together. Also Nature may cause faults that are dependent. However such dependencies are not part of a coordinated plan, and are usually addressed by using reliability mechanisms and intrusion detection mechanisms (for traceability). For example, in a low mobility network with only ad hoc faults, routes have a high probability of remaining connected, and when they are disconnected they can be rebuilt locally.

Malicious or *Byzantine* faults are caused by the adversary, and are usually strongly dependent. The adversary can be *passive* or *active*. Passive attacks are essentially eavesdropping attacks. Active attacks involve action by the adversary which can take different forms. The adversary can corrupt communicated data, fabricate data, or impersonate other nodes. In the extreme case, we may have to deal with one-time, all-out attacks such as terrorist attacks. Malicious faults affect the *robustness* of the network and are usually addressed by using a combination of cryptographic mechanisms and redundancy.

In this paper we are mainly concerned with malicious faults. Such faults occur when a node fails to respond to protocol calls in the prescribed way. For example, when a node X does not respond to protocol calls from a one-hop neighbor Y , thus effectively breaking the link (Y, X) . Unlike ad hoc faults which may occur with a predictable frequency, malicious faults are unpredictable and cannot be addressed by using statistical analysis tools. Intrusion detection tools may also fail to detect such faults.

3 The cell-grid and two basic gossip protocols

Our goal in this section is to find gossip protocols that minimize redundancy while retaining some of the positive features of flooding, such as maximal coverage, minimal state retention

and minimal path length delivery. We are only concerned with large dense networks for which the redundancy in flooding may cause a broadcast storm. We assume that there is generally a locally uniform node density, in particular that no parts of the network are sparse. Finally, for simplicity, we assume that all nodes of the ad hoc network have the same broadcast range: one hop. This will be our unit of measurement. We start by defining the cell-grid and show how it can be used for gossiping.

A *cell-grid* is a covering (or tiling) of the Euclidean plane with regular hexagons, or *cells* – as shown in Figure 1. The cells are the basis for message propagation in our gossip protocols. Our

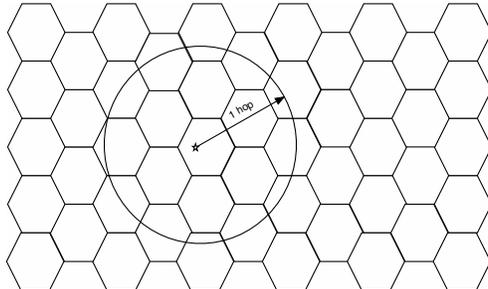


Fig. 1. The cell-grid and a node with its broadcast range.

approach in general will be to have at least one node from each cell to be active and propagate the message (a *gossip* node), resulting in cell-to-cell propagation (although as we shall see, some cells can be silent). Thus effectively we reduce node-to-node flooding to cell-to-cell *gossiping*. To minimize the redundancy we must choose the size of each cell to be maximal subject to fade-out. Therefore, for cell-to-cell propagation, we choose the grid size so that the maximum distance between any two points of an adjacent pair of cells is no more than one hop, since in the worst case, there may only be nodes on the boundary of the cells. Let ℓ be the length of an edge of the regular hexagon cell in hops. The maximum distance between two adjacent cells is:

$$\ell \sqrt{(2\sqrt{3})^2 + 1} = \ell\sqrt{13},$$

as shown in Figure 2. Since we want this distance to be bounded by one hop, we take $\ell = \frac{1}{\sqrt{13}}$

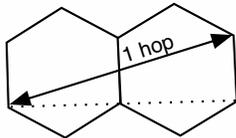


Fig. 2. The maximum distance between points of an adjacent pair of cells is 1 hop.

of a hop. Then the area of a cell is: $\frac{3}{2}\ell^2\sqrt{3} = \frac{3}{26}\sqrt{3}$ of a hop square, which is roughly $\frac{1}{5}$, or $\frac{1}{5\pi}$ of a hop circle.² We will apply this observation shortly. Since one cannot control cells in which

² A hop circle is the area of a circle with radius one hop.

all nodes are faulty, we are only concerned with cells that have non-faulty nodes. Furthermore, we shall assume that these cells are connected, that is:

Connectivity Assumption: All cells with at least one non-faulty node form a connected region. This will be our basic assumption and will be used with our gossip protocols, when dealing with network coverage issues. This is not as restrictive as it may seem: if the network is partitioned then our results still hold, but are restricted to partitions.

3.1 Connected vertex coverings of the network cell graph

Our gossip protocols promote message propagation by cell-to-cell gossiping. One way to achieve this is by ensuring that at least one node from each cell will retransmit the received message. However, as we shall see, this is a worst case solution. In general it is not necessary for *all* cells of the network to be gossipy. We only need to consider connected subsets whose one-hop range will cover the network.

To analyze the structure of such sets we view the set V_{cell} of all cells with at least one non-faulty node, as the node set of a *network cell graph* $G_{cell} = (V_{cell}, E_{cell})$, whose edges $(c_i, c_j) \in E_{cell}$ are adjacent pairs of cells in the cell-grid. This graph may be regarded as a communication graph for the network, since we get network coverage when its cells are gossipy. A *connected vertex covering* (CVC) of G_{cell} is a connected subset C_{cell} of V_{cell} for which at least one of the cells c_i, c_j of every edge $(c_i, c_j) \in E_{cell}$ belongs to C_{cell} . From our discussion in Section 3 (and Figure 2) it is clear that we get network coverage if, and only if, the cells of some CVC of G_{cell} are gossipy. The number of cells of a minimal CVC of G_{cell} ranges from approximately $\frac{1}{3}|V_{cell}| + \frac{1}{2}\sqrt{|V_{cell}|} - 1$ to $|V_{cell}| - 2$. We get the former with square cell graphs and the later with linear cell graphs –see Figure 3. Note that the Vertex Cover problem³ is

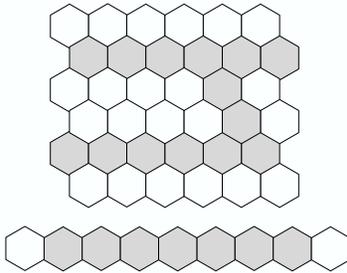


Fig. 3. Minimal connected vertex coverings of a square and a linear cell graph.

NP-complete (for general networks).

Remark 1. There are several advantages in using cell based protocols and a cell based analysis for mobile environments, the main one being that the cells are static. This is particularly relevant for dense networks. For sparse networks, a cell based approach is inappropriate because some

³ Given a graph $G = (V, E)$ and a constant $K \leq |V|$, is there a subset $V' \subset V$ such that: $|V'| \leq K$ and, for each edge (u, v) in E , at least one of u, v belongs to V' ?

cells may become empty and cell based analysis becomes sub-optimal. In this paper we use a cell based approach to get provable (analytical) results (such as Theorems 1–6) as opposed to empirical results that are based on simulations.

3.2 Cell-based gossip

A cell is gossipy if at least one of its nodes x will retransmit a received message m . If there are c_x nodes in the cell of node x , then this would be achieved, on average, if each cell node were to retransmit m with probability $p = 1/c_x$. Since the selection of retransmitting nodes, or gossip nodes, is probabilistic, there is a probability that there will be no gossip node in a cell, *i.e.*, the cell will be *silent*. This is roughly: $(1 - \frac{1}{c_x})^{c_x} \sim e^{-1}$, for large c_x . To reduce this we can use a larger message propagation probability, say $p = k/c_x$, where k , $0 < k \leq c_x$, is a positive *real* number (since some cells can be silent, we can allow for $k < 1$). In this case the probability of a cell being silent is $(1 - k/c_x)^{c_x} \sim e^{-k}$. The probability of a cell being gossipy is then roughly $\rho(k) := 1 - e^{-k}$. We call k the *propagation parameter* and $\rho(k)$ the *gossip cell probability*. The parameter k controls the tradeoff between efficiency and security. Propagation failure occurs when a connected node does not receive a transmitted message.

The easiest way to approximate c_x is to assume a lower bound for the density of the network. Suppose that c_{min} is such a bound for the density of a cell. Then if we take $p = 1/c_{min}$, we should expect to get at least one gossip node per cell.

The first cell-based gossip protocol that we consider, Gossip1, was presented by Haas-Halpern-Li [13]. The input for Gossip1 is: k, c_{min}, s, m , with k the propagation parameter, c_{min} a lower bound for the cell density, $0 < k \leq c_{min}$, s the source and m the message.

Gossip1($k, c_{min}; s, m$) [13]

Node s broadcasts m
for each node x that receives m for the first time **do**
 broadcast m with probability $p = k/c_{min}$

Observe that if we choose $k = c_{min}$, then we have ordinary flooding.

3.3 Dynamic cell-based gossip

When the information on the number of neighbor nodes of a node x , *i.e.* its degree n_x , is available [13], we can select p dynamically for each node x by computing $p_x = \frac{5k\pi}{n_x}$, $0 < k \leq \frac{n_x}{5\pi}$. Our second gossip protocol, which is a further extension of the protocol in [13], uses this approach.

Gossip2 ($k; s, m$)

Node s broadcasts m
for each node x that receives m for the first time **do**
 broadcast m with probability $p_x = 5k\pi/n_x$

In this protocol the expected number of gossips per cell has been reduced from $c_x(\frac{k}{c_{min}})$ to k , when the local node density is uniform. This protocol takes into account the local density and therefore will reduce the propagation failure in networks where density varies, or for which the given lower density bound is too low.

Definition. Let c be the number of cells of the network cell graph G_{cell} and $Bin_{cvc}(z) := \sum_{C \in \mathcal{C}_{cvc}} z^{|C|} (1-z)^{c-|C|}$, where \mathcal{C}_{cvc} is the set of all connected vertex coverings of G_{cell} .

Theorem 1. *If there are no malicious faults then the probability that we get complete coverage with Gossip1 and Gossip2 is at least:*

$$P_{Gossip1} = P_{Gossip2} = Bin_{cvc}(\rho(k)).$$

Proof. In Gossip1 the probability that a neighbor node y of x will retransmit a message m is $p = k/c_{min}$. Assuming that the nodes are distributed at random, the probability that a neighbor node y is in the same cell as node x is $1/5\pi$. So the probability that no neighbor node y in the cell of x will retransmit m is $(1 - p/5\pi)^{n_x} = (1 - k/5\pi c_{min})^{n_x} \leq e^{-k}$ ($n_x \geq 5\pi c_{min}$). Therefore with probability at least $\rho(k) = (1 - e^{-k})$, a cell will retransmit a received message. Since the cells of a CVC C are connected and cover G_{cell} , we get complete network coverage with probability: $\rho(k)^{|C|} (1 - \rho(k))^{c-|C|}$. The proof for Gossip2 is similar, except that we replace all occurrences of c_{min} in the proof by c_x . \square

Estimating the probability in Theorem 1 can be quite complex even if the network cell graph is only a few hundred cells. There are several ways of approximating this probability. One way is to compute the expected coverage (reachability) of a random cell graph C confined to a given area, by approximating the summation in the probability. In Figure 4 we illustrate this for a 5×5 hop square ($\approx 10 \times 10$ cells), with 100-250 random nodes (uniformly distributed). The y -axis shows the expected coverage (reachability) and the x -axis the cell gossip probability $\rho(k)$. It can be seen that for 250 nodes: when $\rho(k) = 0.8$, the expected coverage is 90%, and when $\rho(k) = 0.9$ the expected coverage is 98%. We get similar results for larger dense graphs. For example, for a 10×10 hop square ($\approx 20 \times 20$ cells) with 1000 nodes, the expected coverage is at least 86% when $\rho(k) \geq 0.8$, and 94% when $\rho(k) \geq 0.9$. These results are confirmed empirically by using ad hoc network simulations [13].

The coverage probability of Gossip1 and Gossip2 is restricted to networks with no malicious faults because of the following attack.

The silent attack: A malicious node may fail to respond to the gossip protocol. This will distort the distribution of gossip nodes, resulting in propagation failure. For example, suppose that there are f_0 faulty nodes in the cell of node x that do not respond to the protocol calls of Gossip2. Then the expected number of gossip nodes in the cell of x will be reduced by $f_0 k/c_{min}$ and the coverage failure of that cell raised to $e^{-k(1-f_0/c_{min})}$. The effect of this on the expected coverage will be to shift the graph in Figure 4 to the right (we have to discount the silent nodes).

In Sections 5 and 6 we shall show how to deal with malicious faults.

4 Adaptive gossip

We next present three gossip protocols that adaptively correct probabilistic propagation failures. The first protocol avoids redundant gossips by using *random broadcast delay*, first presented by Ni-Tseng-Chen-Sheu [19]. In the second protocol we assume that nodes can measure the strength of received signals. Signal strength is used to decide whether a node should retransmit the message in order to maximize coverage. In the third protocol nodes can locate the relative direction of the broadcast source. This additional information helps reduce propagation failure and network congestion while maintaining coverage. Notice that each of these approaches is essentially stateless, which is a primary feature of flooding.

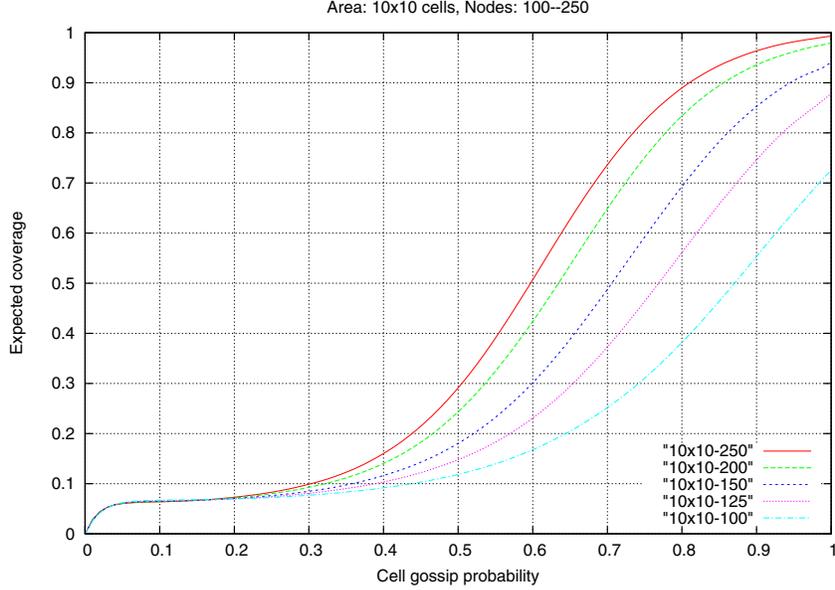


Fig. 4. The expected coverage (reachability) of a cell graph confined to a 5×5 hop square with 150-250 nodes.

4.1 A basic adaptive gossip protocol

Our first adaptive gossip variation relies on probabilistic delay to serialize message retransmissions and extends the protocols in [19, 13]. In this protocol, when a node x receives the gossip g_m of a message m for the first time, it generates a wait-period, randomized within an a priori selected range, *e.g.*, between one and five milliseconds. The node waits during the selected period, counting the number of received gossips g_m . If the counter meets a retransmission threshold before the wait-period ends, then x will not retransmit and disregard all further gossips of m . If the counter does not meet the threshold then x will retransmit m and disregard all further gossips of m .

AdaptiveGossip ($k; c, m$) [19]

Node s broadcasts m

for each node x that receives m for the first time **do**

delay at random within the contention time

if the number of received gossips of m is less than $5k\pi$ **then**

broadcast m

Theorem 2. *Suppose there are no malicious faults. Then, the probability of complete coverage of AdaptiveGossip is at least $\text{Bin}_{cvc}(\rho(k))$.*

Proof. In any neighborhood of a node x , there will be at least $5k\pi$ gossip nodes (provided $5k\pi \leq n_x$). Since we assume that the local density is uniform, the probability that a given neighbor y of x will be in the same cell as x is $\frac{k}{5k\pi} = \frac{1}{5\pi}$. Therefore the probability that a cell is gossipy is at least

$$1 - \left(1 - \frac{1}{5\pi}\right)^{5k\pi} \geq 1 - e^{-k} = \rho(k).$$

The rest is as in Theorem 1. □

The chatterbox attack: A malicious node can retransmit repeatedly the same message. This will distort the distribution of the gossip nodes (some nodes that would normally be gossipy will now be silent), resulting in propagation failure. In particular, if the hop envelope of the chatterbox nodes encloses the hop circle of the source node, the coverage is restricted to the envelope.

Both the silent attack and the chatterbox attack belong to a general family of attacks on gossip protocols in which the adversary tries to distort the distribution of the gossipy nodes to cause propagation failure.

4.2 A signal strength gossip protocol

In this protocol, signal strength information is used to estimate whether the sender and receiver are in the same cell. If this is the case then propagation is not needed. The estimation is obtained by calculating the probability that the sender is not in the receiver's cell given his signal strength. The protocol is given below.

Let g_m be a received gossip of m , $sigstrength(g_m)$ be the signal strength of g_m , with value in the range $[0, 1]$, and $S_0 \in [0, 1]$ be a signal strength threshold.

AdaptiveSignalStrengthGossip ($k, S_0; s, m$)

Node s broadcasts m

for each node x that receives m for the first time **do**

delay at random within the contention time

if the number of gossips g_m of m with $sigstrength(g_m) < S_0$ is less than k **then**

broadcast m

Lemma 1. *Let $h \in [0, \ell]$, where ℓ is the cell radius. The probability $p(x; h)$ that a randomly selected node with distance at most h (of a hop) from x is not in the cell of x is at most $\frac{2h}{3\ell}(2 - \frac{h}{\ell})$.*

Proof. We consider two cases: when the distance of x from its cell boundary is at most h —see Figure 5, and when the distance is greater than h . In the first case the probability $p(x; h)$ is at most $\frac{2}{3}$, with the maximum occurring when x is at one of the six corners of the cell (in Figure 5 take y to be the random node outside the cell of x). This case occurs with probability $1 - (1 - \frac{h}{\ell})^2 = \frac{h}{\ell}(2 - \frac{h}{\ell})$. In the second case, when the distance of x from its cell boundary is greater than h , we have $p(x; h) = 0$. Therefore overall, the probability that a randomly selected node is not in the same cell as x is at most $\frac{2h}{3\ell}(2 - \frac{h}{\ell})$. □

Theorem 3. *Let $S_0 \in [0, 1]$ be a signal strength threshold and h_0 the distance (in hops) over which the signal strength reduces to S_0 , with $h_0 \leq \ell$, the cell radius. Suppose there are no malicious faults. Then, the probability of complete coverage of AdaptiveSignalStrengthGossip is at least $\text{Bin}_{cvc}(\rho(k\beta))$, where $e^{-\beta} = \frac{2h_0}{3\ell}(2 - \frac{h_0}{\ell})$. For $h_0 = \ell$ this is $\text{Bin}_{cvc}(\rho(k \ln 1.5))$.*

Proof. From the protocol, for every non-gossip node x , there will be at least k gossip nodes y with distance at most h_0 from x . By Lemma 1, the probability that each such node y is not in the same cell as x is at most $\frac{2h_0}{3\ell}(2 - \frac{h_0}{\ell})$. In total, the probability that all these nodes y are not in the same cell as x is at most $(\frac{2h_0}{3\ell}(2 - \frac{h_0}{\ell}))^k$. Therefore, the probability that a cell is gossipy is at least $(1 - (\frac{2h_0}{3\ell}(2 - \frac{h_0}{\ell}))^k) = 1 - e^{-k\beta} = \rho(k\beta)$. The rest is as in Theorem 1. □

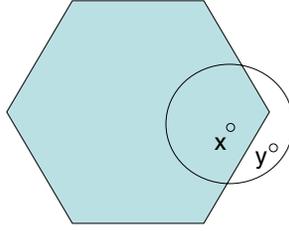


Fig. 5. Node x has distance at most h from the cell boundary.

In the following variation, we use signal strength information to estimate whether the sender and receiver are in the same cell. If this is the case then retransmission is not needed. The estimate is obtained by calculating the probability that the sender is not in the receiver’s cell given his signal strength. Denote this probability by $p(\text{sigstrength}(g_m))$. The protocol is given below.

AdaptiveVariableSignalStrengthGossip ($k; s, m$)

Node s broadcasts m

for each node x that receives m for the first time **do**

delay at random within the contention time

let $v_x = \sum_{g_m} -\ln[p(\text{sigstrength}(g_m))]$, where g_m is any received gossip of m .

if $v_x < k$ **then** broadcast m

Theorem 4. *Suppose there are no malicious faults. Then, the probability of complete coverage for AdaptiveVariableSignalStrengthGossip is at least $\text{Bin}_{cvc}(\rho(k))$.*

Proof. For each node x , the probability that a received g_m was broadcast from outside the cell of x is $p(\text{sigstrength}(g_m))$. So the probability that no node in the cell of x retransmits the message is at most $\prod_{g_m} p(\text{sigstrength}(g_m)) = e^{-v_x}$. Then the probability that a cell is gossipy is at least:

$$1 - e^{-v_x} = \rho(v_x) \leq \rho(k).$$

The rest is as in Theorem 1. □

4.3 Remarks

With adaptive gossip protocols, nodes can make the retransmission decision non-deterministically based on local information. This allows ad hoc networks to avoid failures from broadcast storms and to significantly improve their energy efficiency. Rather than use coin flips to reduce collisions as done in classic gossip protocols [10, 22, 9], our protocols use random contention time. Additionally, by taking advantage of signal strength, our protocols guarantee full network coverage with an exponentially small chance of failure. In the next sections, we show how to guarantee delivery.

The power of our model is illustrated when dealing with signal strength detection error. For example, as long as we are given an upper bound E on the error, we can easily replace signal strength used in the protocol by $\text{sigstrength}(g_m) + E$.

5 Signal Direction

In the preceding section there was a small probability that propagation may fail to reach some nodes. While this may be acceptable in certain cases, for other cases, involving route discovery and broadcasting of control information, a guaranteed broadcast protocol is often desired. In this section, and in the following section, we analyze possible solutions for this problem that avoid broadcast storms. Further, our protocol is robust against malicious faults.

In our next gossip protocol we use direction information to eliminate propagation failure. We assume that each node can distinguish the direction sector from which a signal is received, by using a directional antenna. The area around a node x is divided into 6 sectors, 60° each –see Figure 4, which we label A^i , $i \in [1..6]$. In the protocol, nodes first check to see if the target message has sufficiently propagated without their participation; if not, they will retransmit. More specifically, each intermediate node will retransmit if, and only if, after a random time period, it has not received gossips from all six sectors A^i . The nodes perform the following protocol.

AdaptiveSignalDirectionGossip (s, m)

Node s broadcasts m

for each node x that receives m for the first time **do**

delay at random within contention time

for each direction t in $[1..6]$ **do**

if no gossip g_m of m is received from the direction sector t **then**

broadcast m and halt

Theorem 5. *Protocol AdaptiveSignalDirectionGossip always succeeds.*

Proof. Let x be a non-faulty node, A^i , $i \in [1..6]$, be a sector and $A_x^i \subset A^i$ be the part of A^i that is within distance one hop from x , *i.e.*, A_x^i is one-sixth of the hop circle of x . If x does not retransmit then there is at least one node y in A_x^i that retransmits. Clearly every node in A_x^i can cover this sector completely. So y 's retransmission will cover this sector and x does not need to transmit. The broadcast area of x is illustrated in Figure 6. This argument applies to all sectors A_x^i , $i \in [1..6]$, and all nodes x . Therefore, if cells with at least one non-faulty node form a connected region then we get complete network coverage. \square

5.1 Remarks

The focus of our model is to provide provable deliverability properties, even in the worst case scenario. For example, we prove that our signal direction protocol guarantees coverage in our model where directional detection is precise. This works because of our selection of six sectors which guarantees that in the worst case, received messages cover at least two sectors of area (120 degrees of area relative to the receiving/deciding node (RDN)). Thus, in the worst possible grouping arrangement, where two transmitting nodes are side by side, but in different sectors, both at the limit of the transmission range of the RDN, their transmissions are still guaranteed to cover all of the nodes in the sector in which they reside, in addition to nodes in an area equivalent to one overlapping sector in the two adjacent sectors.

This redundancy is a natural property of our protocols. In the best case, where all six received messages originate from nodes very close to the RDN, all nodes in the RDN's transmission range will receive the received message six times without retransmission by the RDN.

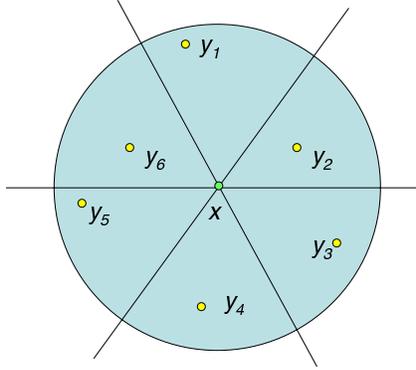


Fig. 6. The six direction sectors of x . The broadcast area is completely covered by the nodes y_1, \dots, y_6 .

The further away the transmitting nodes are from the RDN, the less redundancy is encountered. Still, the redundancy guarantees complete coverage by at least one retransmission even in the absolute worst case.

The power of our model is illustrated when dealing with signal direction detection error. For example, as long as we are given an upper bound on the error, we can easily select an appropriate number of sectors that: (i) allows us to guarantee delivery even in the worst case and, (ii) is optimal in its redundancy while maintaining delivery guarantee. To illustrate this consider a scenario where the maximum signal direction error is 15 degrees. In the six sector model, this error allows an upper bound of a worst case uncovered area of thirty degrees. If we reduce our sector angle by 7.5 degrees (practically, we would reduce it by 8.7 degrees to have seven equivalent sectors), we again can guarantee complete coverage even in the worst case.

6 Geodesic gossip

In this section, we present a geodesic-based gossip protocol. In contrast to location-aware networks [8, 17], our protocol uses only local location information and preserves the location privacy of mobile nodes. The protocol assumes that each node x can obtain a cell identifier cid_x .

AdaptiveGeodesicGossip (s, m)

Node s broadcasts (cid_s, m)

for each node x **do**

delay at random within contention time

if no valid gossip $g_m = (cid_x, m)$ of m is received **then** broadcast (cid_x, m)

When robustness against malicious attacks is desirable, we shall require that the location cid_x is obtained from a *tamper-proof* device attached to x (in a tamper-proof way) that authenticates the location information cid_x . This requirement can be addressed by using a message authentication code. For example, $cid_x = MAC_K(location, time)$, where *location* is the cell location and *time* is rounded to account for time differences and transmission time.

An alternative way to obtain this cell identification information, without using a message authentication code and location sensors, is by combining signal strength and direction. In this

case, gossip nodes do not broadcast their cid_x . This is determined by the receiving nodes by combining the signal strength and direction.

Theorem 6. *Protocol AdaptiveGeodesicGossip always succeeds.*

Proof. By assumption, all cells which have at least one non-faulty node, are connected. Furthermore each cell, with at least one non-faulty node, will have at least one gossip node, since a non-faulty node will retransmit if no other node within its cell transmits. So we get complete coverage. \square

7 Security and Efficiency issues

We have considered five adaptive gossip protocols, namely: Adaptive Gossip, Adaptive Signal Strength Gossip, Adaptive Variable Signal Strength Gossip, Adaptive Signal Direction Gossip, and Adaptive Geodesic Gossip.

The first three protocols use redundancy to probabilistically guarantee message transmission. These have an exponentially small failure probability in the propagation parameter k , by using linear redundancy in k . The last two protocols tolerate malicious faults. For these, message delivery is *guaranteed* while keeping redundancy minimal.

In Figure 7 we compare the performance of all our gossip protocols. The gossip rate is the fraction of nodes in the (one hop) neighborhood of x that retransmit (the fraction of gossip nodes). Note that the coverage probability for the first five protocols assumes no malicious

Protocol	Full coverage probability	Byzantine robustness	Gossip rate	Interface requirement
Gossip1(k) [13]	$Bin_{cvc}(\rho(k))$	no	k/c_{min}	None
Gossip2(k) [19]	$Bin_{cvc}(\rho(k))$	no	$\min(1, 5k\pi/n_x)$	None
ADT-Gossip(k) [19]	$Bin_{cvc}(\rho(k))$	no	$\min(1, 5k\pi/n_x)$	None
ASS-Gossip(k, t)	$Bin_{cvc}(\rho(k\beta))$	no	$\min(1, 13k/n_x)$	Strength sensor
AVSS-Gossip(k)	$Bin_{cvc}(\rho(k))$	no	$\min(1, 13k/n_x)$	Strength sensor
ASD-Gossip	1	yes	$\min(1, 6/n_x)$	Direction sensor
AGEO-Gossip	1	yes	$\min(1, 5\pi/n_x)$	Location sensor

Fig. 7. A comparison of the security features (probability of getting complete coverage and robustness against Byzantine faults) and the gossip rate (the fraction of gossipy nodes over all nodes in a cell – the redundancy) of the proposed protocols compared to flooding. In this table: $\rho(z) = (1 - e^{-z})$; k is the propagation parameter; c is the number of cells of the cell network; β is as in Theorem 3; n_x is the number of nodes in the neighborhood of x ; c_{min} is a lower bound on for the number of nodes in a non-empty cell.

faults, whereas the coverage probability of the last two protocols applies even when there are malicious faults. Finally, note that while our protocols are designed to prevent broadcast storms under reasonable circumstances, we do not consider all-out denial of service (DoS) attacks [15] in this paper. Rather, we assume that DoS is handled by choke points at the physical level.

8 Conclusion

In this paper, we have identified an approach for managing redundancy and security of flooding protocols in dense ad hoc networks. We have mathematically shown the negative impacts of redundancy on ad hoc network bandwidth and how the redundancy can be controlled. Specifically, we give mechanisms that allow network managers the ability to trade off redundancy and its resulting overhead, to provide delivery reliability, and we show how security issues are addressed with this controlled redundancy. Our approach is founded on the concept of cell-grid propagation. Essentially, by tiling the network area with regular hexagons, message redundancy and volume can be tuned to meet the demands for reliability and security. We give protocols to accomplish these objectives and proofs of theorems related to the security properties of those protocols. We show how density is the dominant factor in controlling redundancy in dynamic networks. Finally we note that it is possible to use other regular tilings such as triangular or square tilings, but these do not improve the tradeoff between the propagation failure and the gossip rate.⁴

References

1. B. Awerbuch, D. Holmer, C. Nita-Rotaru and H. Rubens, *An On-Demand Secure Routing Protocol Resilient to Byzantine Failures*. ACM Workshop on Wireless Security (WiSe'02), 2002, pp.21–30.
2. E.M. Belding-Royer and C.-K. Toh, *A review of current routing protocols for ad-hoc mobile wireless networks*. IEEE Personal Communications Magazine, 1999, pp. 46-55.
3. J. Broch et al, *A performance comparison of multi-hop wireless ad hoc network routing protocols*. Proc. ACM MOBICOM, pp. 85-97, 1998.
4. M. Burmester and T. van Le. *Secure Communications in Ad hoc Networks*. Proc. 2004 IEEE Workshop on Information Assurance and Security, West Point, NY, pp. 234–241, 2004.
5. M. Burmester and T. van Le. *Tracing Byzantine faults in ad hoc networks*. Proc. Computer, Network and Information Security 2003, New York, pp. 43–46, 2003.
6. M. Burmester and T. van Le. *Secure Multipath Communication in Mobile Ad hoc Networks*. Proc. International Conference on Information Technology, Coding and Computing, Las Vegas, pp. 405–409, 2004.
7. D. Cavin, Y. Sasson and A. Schiper, *On the Accuracy of MANET Simulators*. Proc. of the 2nd ACM international workshop on Principles of mobile computing, Toulouse, France, pp.38–43, 2002.
8. S. Capkun, M. Hamdi and J. Hubaux, *Gps-free positioning in mobile ad hoc networks*. Proc. of Hawaii Int. Conf. on System Sciences, page 908, 2001.
9. J. Cartigny and D. Simplot, *Border node betransmission based probabilistic broadcast protocols in ad-hoc networks*. Telecommunication Systems **22**(1-4), pp. 189–204, 2003.
10. S. Corson and J. Macker, *Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations*. Memo RFC2501, January 1999.
11. A. Goel, S. Rai and B. Krishnamachari, *Sharp thresholds for monotone properties in random geometric graphs*, Proceedings of the thirty-sixth annual ACM Symposium on Theory of Computing (STOC'04), pp. 580–586, 2004
12. P. Gupta and P.R. Kumar, *Stochastic analysis, control, optimization and applications*. Critical power for asymptotic connectivity in wireless networks, a Volume in Honor of W.H Fleming. Birkhauser, Boston, pp. 547-566, 1998.

⁴ For cell-to-cell propagation the edge of an equilateral triangular tile must be $\ell = 1/2$ of a hop and the density per hop circle roughly 9.2π tiles; similarly, the edge of a square tile must be $\ell = 1/2\sqrt{2}$ of a hop and the density per hop circle 8π tiles. The density of hexagonal tiles per hop circle is 5π which is less in both cases.

13. Z.J. Haas, J.Y. Halpern and L. Li. *Gossip-based ad hoc routing*. Proc. INFOCOM'02, 2002, pp. 1707–1716.
14. D.B. Johnson and D.A. Maltz, *Dynamic Source Routing in Ad-Hoc Wireless Networks*. Mobile Computing. Ed. T. Imielinski and H. Korth, Kluwer Academic Publisher, pp. 152–181, 1996.
15. V. Karpijoki, *Signaling and Routing Security in Mobile Ad Hoc Networks*. Proc. Helsinki University of Technology, Seminar on Internetworking - Ad Hoc Networks, May 2000.
16. B. Karp and H. Kung, *Greedy Perimeter Stateless Routing for Wireless Networks*. Proc. 6th International Conference on Mobile Computing and Networking, Boston, pp. 243-254, 2000.
17. Y.B. Ko and N.H. Vaidya, *Location-Aided Routing (LAR) in mobile ad hoc networks*. Wireless Networks **6**, pp. 307–321, 2000.
18. J. Luo, P. Eugster and J. P. Hubaux, *Route Driven Gossip: Probabilistic Reliable Multicast in Ad Hoc Networks*, Proc. Infocom 2003, 2003.
19. S-Y. Ni, Y-C. Tseng, Y-S. Chen, and J-P. Sheu, *The broadcast storm problem in a mobile ad hoc network*. Proc. 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking, pp. 151-162, 1999
20. C.E. Perkins, E.M. Royer and S.R. Das, *IP Flooding in ad hoc networks*. Internet draft (draft-ietf-manet-bcast-00.txt), Nov 2001. Work in progress.
21. C.E. Perkins and E.M. Royer, *Ad hoc on-demand distance vector routing*, IEEE Workshop on Mobile Computing Systems and Applications, pp. 90-100, 1999.
22. Y. Sasson, D. Cavin and A. Schiper, *Probabilistic Broadcast for Flooding in Wireless Mobile Ad hoc Networks*. Proc. of IEEE WCNC 2003.