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

No Institute Given

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 a broadcast storm in which the number of collisions is so large that we get system failure. To prevent such storms, mechanisms that reduce unnecessary transmissions must be adopted. Such mechanisms will also improve the efficiency and reduce power requirements. Several variants of the flooding protocol have been proposed that reduce the 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. 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 and distance 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 which exploit local connectivity and adaptively correct local propagation failures. These use a cell-grid approach and preserve cell-distance. Our last two gossip protocols are non probabilistic 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, 19, 12]. 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 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 also competing properties, in that improving efficiency is likely to

[*] 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.

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 [6, 7, 20, 18, 10, 19]. The de facto route discovery algorithm for such networks is broadcast flooding [2–4, 18]. With flooding, each node that receives a message retransmits that message exactly once. Flooding [18] has many positive properties for ad hoc networks including maximal coverage, distance preservation and redundancy. Maximal coverage means that if a time-relevant path[1] 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 less secure 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 [14, 11, 21]. Indeed in dense networks, the redundancy may be catastrophic if a broadcast storm [16] is triggered.

A solution to the broadcast storm problem is to reduce message propagation. This is the approach taken with probabilistic retransmission protocols, also referred to as *gossip* protocols [10, 21, 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 may forego participation in accordance with the protocol. Thus, nodes 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.

The rest of this paper is organized as follows. In Section 2 we discuss our security model and malicious link faults. In Section 3 we define our cell-grid and the concept of cell-to-cell propagation and 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 not change. (see *e.g.*, [5]). 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 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.

[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.

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 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 link faults. Such faults are caused 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 (X, Y). Unlike ad hoc link faults which may occur with a predictable frequency, malicious link 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 message propagation while retaining some of the basic features of flooding: maximal coverage, distance information and redundancy. 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 protocols. For our gossip protocols our approach will be to have at least one node from each cell to be active and propagate the message, resulting in cell-to-cell propagation. Thus effectively we reduce node-to-node flooding to cell-to-cell gossiping. To minimize the number of propagations we must choose the size of each cell to be maximal subject to faid-out. Therefore, for cell-to-cell propagation, we choose the grid size so that the maximum distance between any two points of two adjacent 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:

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

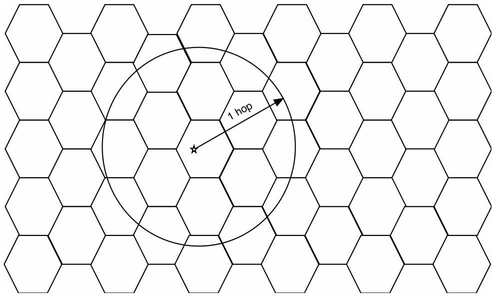

Fig. 1. The cell-grid and a node with its broadcast range.

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

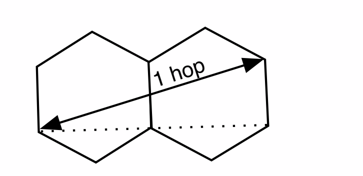

Fig. 2. The maximum distance between points of an adjacent pair of cells is 1 hop.

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 all nodes are faulty, we shall assume that cells with at least one non-faulty node form a connected region.

3.1 Cell-based gossip

Our first gossip protocol aims at ensuring that at least one node from each cell will retransmit the received message m. If there are r nodes in a cell, then this would be achieved, on average, if each cell node were to retransmit m with probability $p = 1/r$. Obviously if there are no nodes in a cell then the network is locally sparse and our protocols will fail. While this may appear to restrict the applicability of our protocols, we contend that the primary target of ad hoc network applications is dense networks, for which our protocols are designed.

Since the selection of retransmitting nodes, or *gossip* nodes, is probabilistic, there is a possibility of *propagation failure*. This is roughly: $(1 - \frac{1}{r})^r \sim e^{-1}$, for large r. To reduce this we can use a larger message propagation probability, say $p = k/r$, where k, $1 \leq k \leq r$, is a small integer. In this case the probability of propagation failure will be approximately e^{-k}. We call k the *security parameter*. Observe that for a cell of the cell-grid to be excluded, there must be propagation failure in all six of its neighbor cells. This reduces even further the propagation failure to e^{-6k}.

The easiest way to approximate r is to assume a lower bound for the density of the network. Suppose that n_0 is such a bound for the density of a hop circle. We must have at least one

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

gossip node per cell, so in this minimal configuration, $r = n_0/5\pi$, and therefore $p = 5\pi/n_0$. The first cell-based gossip protocol that we describe, Gossip1, is a variation of that presented by Haas-Halpern-Li [11] in which the propagation probability is based on cell density. The input for Gossip1 is: k, n_0, s, m, where k is the security parameter, s the source and m the message.

```
**Gossip1**$(k, n_0; s, m)$ [11]
 Node $s$ broadcasts $m$.
 FOR EACH node $x$ that receives $m$ for the first time DO
    broadcast $m$ with probability $p = 5k\pi/n_0$.
```

3.2 Dynamic cell-based gossip

A more dynamic way to approximate r is to compute it based on the number of neighbors of a node (its degree). In this case, for node x, $r_x = d_x/5\pi$, where d_x is the degree of x (the area of a cell is $1/5\pi$ of a hop circle). Of course, now nodes need to now the number of their neighbors. This can be achieved by having all nodes make short "hello" calls at regular intervals. We next describe our second gossip protocol, which is a further extension of the protocol in [11]. Let p_0, $0 < p_0 \leq 1$, be a hello density factor and id_x an identifier for node x.

```
**Gossip2** $(k, p_0; s, m)$
 [ FOR EACH node $x$ DO:
    periodically broadcast "hello, $id_x$" with probability $p_0$
    compute $d_x^* = \#\{(\text{hello}, id_x)\text{'s received during a one time-period}\}/p_0$ ]^3.
 Node $s$ broadcasts $m$.
 FOR EACH node $x$ that receives $m$ for the first time DO
    broadcast $m$ with probability $p_x = 5k\pi/d_x^*$.
```

In this protocol the number of hello calls has been reduced (from d_x) by using the hello density factor p_0. This protocol takes into account the local node density and therefore will reduce the propagation failure in networks where density varies, or for which the given lower density bound is too low.

Theorem 1. *Let s be the number of nonempty cells in a network. If there are no malicious link faults then:*

1. *The propagation failure probability of **Gossip1** is at most se^{-k}.*
2. *The propagation failure probability of **Gossip2** is at most $(1+s)e^{-k/(1+\epsilon)}$, where $\epsilon = \sqrt[3]{k/2n_0p_0^2}$.*

Proof. Let n_x be the number of neighbor nodes of node x. Because nodes are distributed evenly in the local area, the number of nodes in the cell of x is approximately $n_x/5\pi$. The probability that a neighbor node y of x will retransmit a gossip of m is $p = 5k\pi/n_0$. So in Gossip1 the probability that no neighbor node y in the cell of x will retransmit a gossip is:

$$(1-p)^{\frac{n_x}{5\pi}} = \left[\left(1 - \frac{5k\pi}{n_0}\right)^{\frac{n_0}{5k\pi}}\right]^{\frac{5k\pi \cdot n_x}{n_0 \cdot 5\pi}} < \left(\frac{1}{e}\right)^{\frac{kn_x}{n_0}} < e^{-k}.$$

[3] Hello calls are sent and received in the background.

Thus the probability that a cell has no gossip node is at most $se^{-k} = e^{-k+\ln s}$. So with probability at least $1 - e^{-k+\ln s}$, every cell has a gossip node. By our assumption, all the cells are connected. Thus, a message m will be propagated to all cells with this probability. Consequently Gossip1 fails with negligible probability, at most $e^{-k+\ln s}$.

In Gossip2, the probability that $d_x^* > n_x(1+\epsilon)$ is at most $e^{-2n_x(p_0\epsilon)^2}$ for all $\epsilon > 0$ (Okamoto's bound [17]). Using the same argument as above, we can show that Gossip2 will fail with probability at most $e^{-k/(1+\epsilon)+\ln s} + e^{-2n_0(p_0\epsilon)^2}$, where $n_x \geq n_0$. Taking $\epsilon_0 > 0$ such that $\epsilon_0^2(1+\epsilon_0) = k/2n_0p_0^2$, we see that the propagation failure probability of Gossip2 is at most

$$(1+s)e^{-k/(1+\epsilon_0)} < (1+s)e^{-k/(1+\sqrt[3]{k/2n_0p_0^2})},$$

where $\epsilon_0 < \sqrt[3]{k/2n_0p_0^2}$. □

While Gossip2 is able to adapt to global node density variations, it requires that all nodes adhere to the protocol. If there are malicious link failures, then the propagation failure will be affected (raised), and indeed may limit propagation to only a few cells. Similarly the propagation failure of Gossip1 will be affected. For example, suppose that there are f_0 faulty nodes in the neighborhood of a node x that do not respond to the protocol calls of Gossip1. Then the expected number of gossip nodes in the cell of x will be reduced by f_0k/n_0 and the propagation failure of that cell will be raised to $e^{-k(1-f_0/n_0)}$. In the following section we shall show how to deal with malicious link faults by using adaptive calls.

4 Adaptive gossip

We present three gossip protocols that adaptively correct propagation failures due to malicious behavior by using readily available information. The first protocol is a variation of Gossip2, in which we avoid hello calls by using *random broadcast delay.* 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.

4.1 A basic adaptive gossip protocol

The first adaptive gossip variation relies on probabilistic delay to serialize message retransmissions and extends the protocols in [16, 11]. 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 à 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 g_m. If the counter does not meet the threshold then x will retransmit m and disregard all further gossips g_m.

```
AdaptiveGossip (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.
    IF number of received gossips of $m$ is less than $5k\pi$ THEN
        broadcast $m$.
```

Theorem 2. *The propagation failure probability of AdaptiveGossip is at most se^{-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 there are no gossip nodes in the cell of x is at most

$$(1 - \frac{1}{5\pi})^{5k\pi} = ((1 - \frac{1}{5\pi})^{5\pi})^k < e^{-k}.$$

Consequently, the propagation failure probability of AdaptiveGossip is at most se^{-k}. □

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 3, 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 3 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 $\frac{2h}{3\ell}(2 - \frac{h}{\ell})$. □

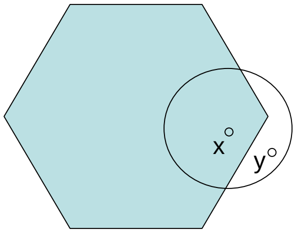

Fig. 3. Node x has distance at most h from the cell boundary.

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. The propagation failure probability of AdaptiveSignalStrengthGossip is at most*

$$s\left[\frac{2h_0}{3\ell}\left(2-\frac{h_0}{\ell}\right)\right]^k \;<\; s\left(\frac{4h_0}{3\ell}\right)^k .$$

In particular when $h_0 = \ell$, the propagation failure probability is at most $s(3/2)^{-k}$.

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 overall failure probability is at most $s\,(\frac{2h_0}{3\ell}\,(2-\frac{h_0}{\ell}))^k$. For $h_0 = \ell$ this is $s(3/2)^{-k}$. In this case the *gossip rate per hop circle*, *i.e.*, the number of nodes in a hop circle that retransmit, is $13k$ (there are k gossip nodes in a circle of radius ℓ, and $\ell = 1\sqrt{13}$ of a hop). We can reduce the failure rate by taking h_0 to be a fraction of ℓ, but then the gossip rate will increase. □

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 $p\,(sigstrength(g_m))$. 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 $v_x = \sum_{g_m} - \ln\,[p\,(sigstrength(g_m))]$, where $g_m$ is any received gossip of $m$.
    IF $v_x < k$ THEN broadcast $m$.
```

Theorem 4. *The propagation failure probability of AdaptiveVariableSignalStrengthGossip is at most se^{-k}.*

Proof. For each node x, the probability that a received g_m was broadcast from outside the cell of x is $p\,(sigstrength(g_m))$. So the probability that no node in the cell of x retransmits the message is at most $\prod_{g_m} p\,(sigstrength(g_m)) = e^{-v_x}$. If x does not retransmit m itself then $v_x \geq k$. Thus the propagation failure probability is at most $se^{-v_x} \leq se^{-k}$. □

4.3 Summary

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, 21, 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.

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.

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^o 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 gossip 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 $(k; s, m)$

```
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 stop.
```

Theorem 5. *Protocol AdaptiveSignalDirectionGossip always succeeds.*

Proof. Let x be a non-faulty node, A^i, $i \in [1..6]$, be a sector and $A^i_x \subset A^i$ be the part of A^i that is within distance one hop from x, *i.e.*, A^i_x is one-sixth of the hop circle of x. If x does not retransmit then there is at least one node y in A^i_x that retransmits. Clearly every node in A^i_x 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 4. This argument applies to all sectors A^i_x, $i \in [1..6]$, and all nodes x. Therefore, if cells with at least one non-faulty node form a connected region then we get complete network coverage. □

6 Cell location gossip

In this section, we present a cell location gossip protocol. In contrast to location-aware networks [8, 15], our protocol is lightweight. It uses only local location information and preserves

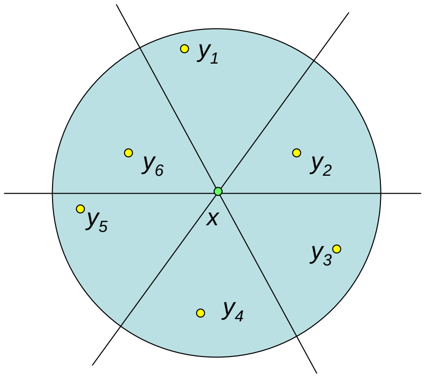

Fig. 4. The six direction sectors of x. The broadcast area is completely covered by the nodes $y_1, \ldots, y_6$.

the privacy of mobile nodes. The protocol assumes that each node x can obtain a cell identifier cid_x. To avoid malicious attacks, we also 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. To avoid replay attacks, the location information should include the time of measurement. An authenticated cid_y will not be valid if it is older than τ time units, where τ is an upper time bound for a one-hop transmission. If privacy is a concern then encryption mechanisms should be used. These requirements can be addressed by using a keyed hash function.[4]

AdaptiveLocationGossip (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 $(cid_x, m)$.
```

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. □

7 Security and Efficiency issues

We have proposed five new adaptive gossip protocols, namely

- AdaptiveGossip
- AdaptiveSignalStrengthGossip

[4] For example, $cid_x := \mathrm{AES}_{sk}(cellLocation(x), time)$ where $time$ is appropriately rounded. AES is the Advanced Encryption Standard [1].

- AdaptiveVariableSignalStrengthGossip
- AdaptiveSignalDirectionGossip, and
- AdaptiveGeodesicGossip.

The first three protocols use redundancy to probabilistically guarantee message transmission. These have an exponentially small failure probability in the security parameter k, by using linear redundancy in k. In the last two protocols, message delivery is *guaranteed* while minimizing redundancy. In Figure 5 we compare the performance of 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 failure rate for the first two protocols assumes no malicious link faults, whereas the failure rates of the adaptive gossip protocols apply even when there are malicious link 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 [13] in this paper. Rather, we assume that DoS is handled by choke points at the physical level.

Protocol	Propagation failure	Robustness against malicious faults	Gossip rate	Notes
Gossip1 [11]	se^{-k}	no	$5k\pi/n_0$	$n_0 \leq \max_x n_x$
Gossip2	$(s+1)e^{-k/(1+\epsilon)}$	no	$5k\pi/n_x + \lambda p_0$	$\epsilon = \sqrt[3]{k/2n_0p_0^2}$
AdaptiveGossip	se^{-k}	yes	$5k\pi/n_x$	
SignalStrengthGossip	$s(3/2)^{-k}$	yes	$13k/n_x$	
VariableSignalStrengthGossip	se^{-k}	yes	$13k/n_x$	
SignalDirectionGossip	0	yes	$6/n_x$	
GeodesicGossip	0	yes	$5\pi/n_x$	

Fig. 5. A comparison of the security features (propagation failure and robustness) and the redundancy (gossip rate) of the proposed protocols. Here: s is the number of cells in the network, k the security parameter, n_x the number of nodes in the neighborhood of x, λ the relative frequency of hellos per broadcast request and p_0 the hello density.

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.[5]

References

1. Advanced Encryption Standard (AES). Federal Information Processing Standards Publications (FIPS PUBS) 197, 2001.
2. 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.
3. 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.
4. J. Broch et al, *A performance comparison of multi-hop wireless ad hoc network routing protocols.* Proc. ACM MOBICOM, pp. 85-97, 1998.
5. 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.
6. 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.
7. 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.
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. Z.J. Haas, J.Y. Halpern and L. Li. *Gossip-based ad hoc routing.* Proc. INFOCOM'02, 2002, pp. 1707–1716.
12. 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.
13. V. Karpijoki, *Signalling and Routing Security in Mobile Ad Hoc Networks.* Proc. Helsinki University of Technology, Seminar on Internetworking - Ad Hoc Networks, May 2000.
14. 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.
15. Y.B. Ko and N.H. Vaidya, *Location-Aided Routing (LAR) in mobile ad hoc networks.* Wireless Networks **6**, pp. 307–321, 2000.
16. 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
17. M. Okamoto, *Some inequalities relating to the partial sum of binomial probabilities.* Annals of the Institute of Statistical Mathematics, **10**, pp. 29–35, 1958.
18. 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.
19. 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.
20. E. Royer and C-K. Toh, *A Review of Current Routing Protocols for Ad-Hoc Mobile Wireless Networks.* IEEE Personal Communications Magazine, pp. 46-55, 1999.
21. Y. Sasson, D. Cavin and A. Schiper, *Probabilistic Broadcast for Flooding in Wireless Mobile Ad hoc Networks.* Proc. of IEEE WCNC 2003.

[5] 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.