

Opportunistic Information Dissemination in Mobile Ad-hoc Networks: Adaptiveness vs. Obliviousness and Randomization vs. Determinism * **

Martín Farach-Colton¹, Antonio Fernández Anta², Alessia Milani³,
Miguel A. Mosteiro⁴, and Shmuel Zaks⁵

¹ Dept. of Computer Science, Rutgers University, Piscataway, NJ, USA &
Tokutek, Inc.

`farach@cs.rutgers.edu`

² Institute IMDEA Networks, Madrid, Spain.

`antonio.fernandez@imdea.org`

³ LABRI, University of Bordeaux 1, Talence, France.

`milani@labri.fr`

⁴ Dept. of Computer Science, Rutgers University, Piscataway, NJ, USA &
LADyR, GSyC, Universidad Rey Juan Carlos, Madrid, Spain.

`mosteiro@cs.rutgers.edu`

⁵ Dept. of Computer Science, Technion - Israel Institute of Technology, Haifa, Israel.

`zaks@cs.technion.ac.il`

Abstract. In this paper the problem of information dissemination in Mobile Ad-hoc Networks (MANETs) is studied. The problem is to disseminate a piece of information, initially held by a distinguished source node, to all nodes in a target set. We assume a weak set of restrictions on the mobility of nodes, parameterized by α , the *disconnection time*, and β , the *link stability time*, such that the MANETs considered are connected enough for dissemination. Such a connectivity model generalizes previous models in that we assume much less connectivity, or make explicit the assumptions in previous papers.

In MANETs, nodes are embedded in the plane and can move with bounded speed. Communication between nodes occurs over a collision-prone single channel. We show upper and lower bounds for different types of randomized protocols, parameterized by α and β .

This problem has been extensively studied in static networks and for deterministic protocols. We show tight bounds on the randomized complexity of information dissemination in MANETs, for reasonable choices of α and β . We show that randomization reduces the time complexity of the problem by a logarithmic or linear factor, depending on the class of randomized protocol considered.

* This research was supported in part by the National Science Foundation (CCF-0937829, CCF-1114930), Comunidad de Madrid grant S2009TIC-1692, Spanish MICINN grant TIN2008-06735-C02-01.

** We thank Seth Gilbert for triggering the development of this work asking a question at DISC 2010.

1 Introduction

A *Mobile Ad-hoc Network (MANET)* is a radio network model defined by a set of processing nodes that move in an environment. The communication is node-to-node by fixed-radius radio transmission, and does not rely on any underlying communication infrastructure (such as routers, access points, etc.). Therefore, MANETs are modeled as graphs where edges are defined between nodes within communication radius. Nodes are assumed to have a synchronous clock, in that they have a uniform time step, but they may fail and reboot at any time, so there is no global notion of a fixed time step. We study the complexity of the *Dissemination* problem [12] in MANETs. The Dissemination problem is that of distributing a piece of information initially held by some source node to a given set of nodes in the network. The Dissemination problem is a generalization of several fundamental problems in distributed systems, e.g. Broadcast, Multicast, Routing, etc., depending on the choice of receiving nodes.

The complexity of information dissemination problems depends on the mobility assumptions one makes, becoming more tractable as more assumptions are made about the mobility [8, 12, 16]. However, overly strong assumptions about node mobility may not be realistic. Several mobility models have been proposed as a compromise, attempting to be as general as possible while still tractable. These include *delay-tolerant networks* [10] and *opportunistic networks* [20]. Any mobility model translates into a model for edge formation/dissolution since edge connectivity depends on Euclidean distance. Thus, mobility models typically have an edge-connectivity flavor, as follows: For every pair of nodes u, v , a message can reach node v from node u in any window of t steps. That is, there may not be any point in time when there is a path from u to v , but as edges appear and disappear over time, there is a sequence of edges that can be taken to get from u to v (a *chrono-path*). Any network with mobility model that requires such connectivity is said to allow *opportunistic communication*. Models differ, for example, in how long they require individual edges to exist, since radio communication cannot be guaranteed in any particular time step, due to collisions.

In this paper we define a class of MANETs, where nodes can move with bounded speed in a plane, communication between nodes occurs over a collision-prone single channel, and which enforces opportunistic communication, as follows. Let $S \subset V$ be any non-empty subset of nodes, and let $\bar{S} = V - S$. An edge is *k-stable* at time t if it exists for k consecutive steps $[t, t + k - 1]$. The opportunistic communication constraint is parameterized by parameters α and β : for every partition (S, \bar{S}) , there are at most α consecutive steps without a β -stable edge between S and \bar{S} . We call such a network (α, β) -*connected*. Note, for example, that (α, β) -connectivity requires that any node that fails reboots (and reestablish connectivity) within α steps.

The aim of this α, β parameterization is twofold. First, in some previous papers, the connectivity depended explicitly [12] or implicitly [20] on the protocol (notably, who has and does not have information at any given time). Here, we remove the dependency on the protocol, while maintaining a similar flavor to

the connectivity constraints. Second, we substantially weaken the connectivity assumptions compared with some previous papers which assumed, for example, a stable spanning tree of edges [16, 19], or that the underlying graph may be non-geometric [4, 7, 8].

Our Results. The main contribution of this work concerns negative results. We prove various bounds for probabilistic algorithms for the Dissemination problem, parametrized by α and β , for bounded speed of movement. Our lower bounds show that the (relatively easy) upper bounds shown are asymptotically tight, for any $\alpha/\beta \in O(1)$. In the classic paper of Kowalski and Pelc [14], the authors studied the impact of randomization and adaptiveness on the time complexity of solving the Broadcast problem in static Radio Networks. We follow that fruitful line of research for the MANET model of dynamic Radio Networks and follow up on our work on deterministic algorithms [12] for this problem.

We classify our results according to the type of randomized protocol considered. (Formal definitions are left to the full version of this paper for brevity.) In **oblivious** randomized protocols [5, 14], the protocol has access to a sequence of random variables at each node which are independent of the execution of the algorithm and mutually independent. We define **locally adaptive** randomized algorithms as a generalization of oblivious protocols in that the random variables at each node are independent of the execution but may be mutually dependent. Orthogonally, in **fair** randomized protocols [5, 11] all nodes transmit with the same probability during each time step.

Let s be the source node of the information to be distributed, and let T be the set of target nodes for this information. We say that a node is **informed** if it has the information. Let I_t be the set of informed nodes at time t , and $\hat{I}_t = \cup_{i \leq t} I_i$. Let $C_t = T \cap \hat{I}_t$ be the set of **covered** nodes. The algorithm terminates at time step t if $T = C_t$.

Here, we show an existential lower bound of $\Omega((n \log(1/\varepsilon))/\log n)$ on the time to cover a new node with probability $1 - \varepsilon$ for fair protocols. We show an existential lower bound of $\Omega(n/\log n)$ on the time to cover a new node, for oblivious protocols with high probability, and for locally adaptive protocols in expectation.

A second collection of results lower bound the time required to complete Dissemination. We show that it takes $\Omega(\alpha n + n^2/\log n)$ time steps to solve the problem, with probability $p \geq 2^{-n/2}$ for fair and oblivious protocols, and in expectation for locally adaptive protocols.

Finally, we show that a straightforward protocol solves Dissemination in $O(\alpha n + (1 + \alpha/\beta)n^2/\log n)$ time, with probability $p \geq 1 - e^{-(n-1)/4}$. Surprisingly, this bound holds for any value of $\alpha \geq 0$, $\beta > 0$, arbitrarily large speed of movement, and arbitrary activation schedules. In this protocol, informed nodes always transmit with the same probability. Observe that, when $\alpha/\beta = O(1)$, the time bound is asymptotically optimal for any of the protocol classes studied, which is rather surprising, given the simplicity of the protocol.

Table 1 summarizes the Dissemination results of this paper, together with the deterministic bounds from [12]. The first important observation is that there

		randomized	deterministic [12]
lower bounds	oblivious	$\Omega(\alpha n + n^2 / \log n)$	$\Omega(\alpha n + n^3 / \log n)$
	adaptive	$\Omega(\alpha n + n^2 / \log n)$ exp.	$\Omega(\alpha n + n^2)$
	fair	$\Omega(\alpha n + n^2 / \log n)$	–
upper bounds	oblivious	$O(\alpha n + (1 + \alpha/\beta) n^2 / \log n)$	$O(\alpha n + n^3 \log n)$
	adaptive	–	$O(\alpha n + n^2)$
	fair	$O(\alpha n + (1 + \alpha/\beta) n^2 / \log n)$	–

Table 1. Randomized lower bounds are to achieve success probability $p \geq 2^{-n/2}$, except as noted. Randomized upper bounds are with probability $p \geq 1 - e^{-(n-1)/4}$.

is no gap between oblivious and locally adaptive protocols. The lower bounds derived match the upper bound shown using a fair oblivious protocol, for any $\alpha/\beta \in O(1)$. The second observation is that randomization reduces the time complexity of the problem by a linear factor in the oblivious case and by a logarithmic factor in the adaptive case (for reasonably small values of α).

It is important to note that all lower bounds in this paper hold even without node failures and with simultaneous node activation, whereas non-simultaneous activation was crucial in showing a separation between obliviousness and adaptiveness for deterministic protocols [12]. On the other hand, it is fair to notice that the adaptive class of protocols considered in [12] is more general, and the constraints on speed assumed to prove the lower bounds are more restrictive.

Roadmap. The rest of this paper is structured as follows. In Section 2 we overview related work. The model used is given in Section 3. In Sections 4 and 5 we present our lower bounds, and the upper bound is proved in Section 6.

2 Previous Work

The closest work to this paper is [12] where we studied deterministic Dissemination in MANETs. The impact of randomization and adaptiveness on the time complexity of Dissemination was studied for the Broadcast problem in static Radio Networks in the paper of Kowalski and Pelc [14]. Here, we follow that fruitful line of research for MANETs. Our results also contribute to understand the impact of mobility in the solvability of problems in Radio Networks. The results in [14] show that in static networks an adaptive algorithm performs better than an oblivious one, independently of randomization and provided that nodes are synchronized. In [12], we show that adaptiveness helps for deterministic Dissemination in MANETs, but only if nodes are not synchronized. We prove here that for randomized algorithms, adaptiveness does not help (for any $\alpha/\beta \in O(1)$) if the algorithm is fair or if we restrict to local adaptiveness. We structure the rest of the related work in terms of model, problem and bounds.

Model. MANETs are frequently assumed to be always connected, but the topology may vary in time provided that a link remains stable sufficiently long [2]. Less restrictive models have been proposed for Dynamic Networks, but the topology may be non-embeddable in \mathbb{R}^2 [8, 15, 16]. In [15, 16] a model for Dynamic Networks is proposed in which, for any T consecutive rounds, there is a *stable connected spanning graph*. Our connectivity assumptions are weaker given that the network may be always partitioned and communication is only opportunistic. In [8] the communication is opportunistic, but for an arbitrary adversary limited (in terms of our model) to $\alpha = 0$ and $\beta = 1$. In [12] communication is opportunistic and the connectivity model is weak, although depends on the protocol.

Bounds. Protocols that do not rely on knowledge of topology may allow node mobility. Topology-independent randomized protocols with theoretical guarantees have been studied, for Gossiping in [6] showing $O(n \log^4 n)$ expected time and for Broadcast in [18] showing $O(n \log n)$ expected rounds. Both protocols require strong connectivity during all the execution and the latter additionally requires collision detection. The first randomized Broadcast protocol for Radio Networks was presented in [3]. The protocol works in $O((D + \log n/\varepsilon) \log n)$ time with probability at least $1 - \varepsilon$, where D is the diameter of the network. More recently, the expectation to solve Broadcast was upper bounded by $O(D \log(n/D) + \log^2 n)$ in [9, 13] for adaptive algorithms, and by $O(n \min\{D, \log n\})$ in [14] for oblivious protocols. None of these protocols is resilient to mobility. On the negative results side, a lower bound of $\Omega(D \log(n/D))$ was derived in [17] for the expected time for any randomized Broadcast protocol and a lower bound of $\Omega(\log^2 n)$ in [1], matching together the upper bound in [9, 13]. Later on, it was shown in [14] that, for every oblivious randomized Broadcast protocol, there exists a static network such that the protocol takes time in $\Omega(n)$ with probability at least $1/2$ to complete Broadcast. Given that a static network is a particular case of a MANET, these lower bounds apply to our setting. Nevertheless, we improve these bounds here by exploiting mobility.

Communication techniques to solve Dissemination, such as flooding, have been studied for Dynamic Networks. Only upper bounds apply to MANETs since the topology in Dynamic Networks may be non-geometric. A straightforward protocol was proved [8] to complete Broadcast with high probability in $O(n^2/\log n)$ steps in Dynamic Networks where $\alpha = 0$ and $\beta = 1$. Here we show the same upper bound for MANETs but parameterized in arbitrary $\alpha \geq 0$ and $\beta > 0$. Since MANET is a sub-class of Dynamic Networks, our lower bounds are valid for the more general class.

3 Model

We consider a MANET formed by a set V of n mobile nodes deployed in \mathbb{R}^2 with the Euclidean distance metric. It is assumed that each node has data-processing and radio-communication capabilities, and a unique identity.

Time. Time is slotted as a sequence of intervals, each of length t long enough to transmit (resp. receive) the information to be disseminated. Computations are assumed to take negligible time. Each node is equipped with a clock that ticks at the same uniform rate $1/t$. All node's ticks are assumed to be in phase (as in [21]). Throughout the paper, we use the term *time slot* to refer to a global time of reference, and the term *time step* to refer to the time reference local to a node.

Node Activation. We say that a node is *active* if it is up and running, and *inactive* otherwise. Due to various events (failures, recharging batteries, etc.) nodes may become inactive and be re-activated. We call the temporal sequence of activation and deactivations of a node the *activation schedule*. We assume that a node is activated in the boundary between two consecutive time slots. If a node is activated between $t - 1$ and t we say that it is activated at slot t , and it is active in t .

Node Memory. Upon activation, a node immediately starts running an algorithm previously stored in its hardware. We consider two scenarios regarding node memory. If an inactive node does not lose its memory, we assume that it continues the execution of the algorithm from the failing step upon reactivation. If memory is lost, we assume that it restarts the algorithm from scratch.

Radio Communication. Nodes communicate via a collision-prone single radio channel without collision detection and fixed radius of transmission r . An active node u receives a transmission from a *neighbor* (a node within radius) v at time slot j if and only if v is the only node in u 's neighborhood (including u itself) transmitting at time slot j . We say in this case that the transmission was *successful*, or *unique* among the neighbors of u .

Adversary. We assume the presence of an adversary that controls each node movement restricted to the speed and the (α, β) -connectivity constraints. For the upper bound only, we further assume that the adversary controls the activation schedule (including failures). This adversary is adaptive, in the sense that it makes decisions at the end of each time slot with access to all nodes' internal state, but without access to future random bits.

4 Link Stability Lower Bounds

Link stability assumptions are crucial in solving any Dissemination problem. In our model, such characteristic is parameterized with β . In this section, we relate link stability to the worst-case running time and the probability of solving Geocast, an instance of Dissemination. In Geocast, a node x is a target of the information if and only if, at the initial time, x is up and running, and it is located within a parametric distance $d > 0$ (called *eccentricity*) from the position of the source node at that time. Notably, exploiting fairness, it is possible to quantify such relation no matter how strong are the stability guarantees (that is for any $\beta > 0$), which we do in the first theorem.

Theorem 1. *For any $v_{\max} > 0$, $d > r$, $0 < \varepsilon < 1$, $\alpha \geq 0$, $\beta > 0$, any $k, c \in \mathbb{N}^+$, such that $45 \leq k < n$, and for any fair randomized protocol Π for Geocast, there*

exists an (α, β) -connected MANET formed by a set V of n nodes running Π such that, if $T \leq k \log_4(1/\varepsilon)/(4 \log k)$, within cT steps after k nodes are covered, with probability at least ε^c , no new node is covered, even if all nodes are activated simultaneously and do not fail.

Proof. Consider three sets of nodes A , B , and B' deployed in the plane, each set deployed in an area of size ξ arbitrarily small, such that $0 < \xi < r$ and $d \geq r + \xi$. (An illustration is left to the full version of this paper for brevity.) The invariant in this configuration is that nodes in each set form a clique, every node in A is placed within distance r from every node in B , every node in B is placed at most at distance ξ from every node in B' , and every node in A is placed at some distance $r < \delta \leq r + \xi$ from every node in B' . Also, ξ is set appropriately so that a node can move ξ distance in one time slot without exceeding v_{\max} .

Initially, the adversary places $k - 1$ nodes in the set B' including the source, an arbitrary node y in the set B , and the remaining $n - k$ nodes in set A . Node y preserves (α, β) -connectivity forever. At the beginning of the first time slot all nodes are activated. Let t be the first time slot when the informed source transmits and, hence, all the nodes in $B \cup B'$ become informed. (If this event never happens the claim of the theorem holds trivially.) The adversary does not move any nodes until then. Given that $d \geq r + \xi$, the nodes in A must become informed to solve the problem.

After time slot t , the adversary moves the nodes according to protocol Π as follows. For each time slot $t' > t$ where nodes transmit with probability at least $4 \log k/k$, the adversary moves all nodes in B' to B . At the end of time slot t' the adversary moves all nodes in B but y back to B' , and the procedure is repeated.

We show now that within the first $T \leq \ln(1/\varepsilon)k/(2 \ln k)$ steps the nodes in A are not informed with probability at least ε . Given that the protocol is fair, in a given time slot all nodes use the same probability of transmission. Let $\pi_{t'}$ be the probability of transmission used in step t' . The probability of not achieving a successful transmission in a step t' is $Pr_{fail} = 1 - \sum_{i \in B} \pi_{t'} \prod_{j \in B, j \neq i} (1 - \pi_{t'}) = 1 - |B| \pi_{t'} (1 - \pi_{t'})^{|B|-1}$. We consider two cases depending on whether $\pi_{t'}$ is larger or smaller than $4 \log k/k$. For both cases it can be shown that $Pr_{fail} \geq 4^{-4 \log k/k}$. The details are omitted in this extended abstract for brevity. Then, the probability of failing to inform the nodes in set A within the interval $[t + 1, t + T]$ is $Pr_{fail}(T \text{ steps}) \geq 4^{-4T \log k/k} \geq \varepsilon$, if $T \leq k \log_4(1/\varepsilon)/(4 \log k)$. Conditioned on this event, the same analysis can be applied to the subsequent interval of T steps, and inductively to each subsequent interval of T steps. \square

For oblivious protocols, possibly not fair, it is shown in the following theorem that, if the link stability guarantees are not strong enough, there exists some configuration where new nodes are not covered with positive probability. The proof, left to the full version of this paper for brevity, uses the adversarial configuration of Theorem 1. However, given that nodes running oblivious protocols may use different probabilities of transmission in a given time slot, in order to preserve the (α, β) -property the node y has to be chosen more carefully. Roughly, the intuition on how y is chosen is the following. We consider each interval of β time

slots separately. It is useful to picture one interval as a matrix of probabilities of transmission, one column for each slot of the interval and one row for each node in $B \cup B'$. Let the columns where the sum over all rows is “small enough” be called *quiet*, and *noisy* otherwise. There exists a row where the sum along quiet columns is also “small enough”. The adversary chooses y to be the node that corresponds to that row because: (a) it is unlikely that y transmits in a quiet slot; and (b) it is unlikely that exactly one of all nodes transmits in a noisy slot.

Theorem 2. *For any $v_{\max} > 0$, $d > r$, $\alpha \geq 0$, $\beta > 0$, any $k, c \in \mathbb{N}^+$, such that $e^3 \leq k < n$, and any oblivious randomized protocol Π for Geocast, there exists an (α, β) -connected MANET formed by a set V of n nodes running Π such that if $\beta < k/(2(1 + \ln k))$, within $c\beta$ steps after k nodes are covered, no new node is covered with probability at least $(2e^e)^{-c}$, even if all nodes are activated simultaneously and do not fail.*

We move now to locally adaptive protocols possibly not fair. Recall that the adversary is adaptive, making decisions at the end of each step with access to all the nodes’ internal state, but without access to their future random bits. For this class, it is shown that if β is not large enough, there exists some configuration where it is not expected to inform a new node within the first β steps. The proof, left to the full version of this paper for brevity, uses the same adversarial configuration of Theorems 1 and 2. The intuition on how y is now chosen is similar to the previous theorem, that is, for each interval of β slots choose a node that is likely to be silent in the quiet slots of that interval. However, now the adversary does not know in advance the probabilities of transmission (the protocol is adaptive). Nevertheless, it can still compute which is the node that most likely is the correct choice for y , as we show.

Theorem 3. *For any locally adaptive randomized protocol Π for Geocast, such that uninformed nodes never transmit, any $v_{\max} > 0$, $d > r$, $\alpha \geq 0$, $\beta > 0$, and any $k \in \mathbb{N}^+$, such that $(2/(1 - 1/e))^{\xi/e} < k < n$, $\xi \triangleq 2/(1 - 1/e)^2$, there exists an (α, β) -connected MANET formed by a set V of n nodes such that if $\beta < k/(2e\Gamma)$, for $\Gamma \triangleq \xi \ln \delta$, $\delta \triangleq \beta^2 k^{e/\xi}$, within β steps after k nodes were covered, in expectation no new node is covered, even if all nodes were activated simultaneously and do not fail.*

5 Dissemination Lower Bounds

In this section, we show that there exists an instance of Dissemination, namely Geocast, for which adversarial configurations of nodes that require a minimum number of steps exist.

Theorem 4. *For any $n > 24$, $d > r$, $\alpha \geq 0$, $\beta > 0$, $v_{\max} > \pi r/(6\alpha)$, and any fair randomized Geocast protocol Π , there exists an (α, β) -connected MANET of n nodes for which, in order to solve the problem with probability at least $2^{-n/2}$, Π takes at least $\alpha n/2 + n^2/(96 \ln(n/2))$ time slots, even if all nodes are activated simultaneously and do not fail.*

Proof. We prove the claim for fair protocols where the probability of transmission of a node is independent among time slots. For non-independent fair protocols, the same techniques can be used to prove the same bound, with a more detailed analysis.

The following adversarial configuration and movement of nodes show the claimed lower bound. Consider four sets of nodes A , B , B' , and C , each deployed in an area of size ε arbitrarily small, such that $0 < \varepsilon < r$ and $d \geq r + \varepsilon$, and four points x, x', y, y' all placed in the following configuration. (An illustration is left to the full version of this paper for brevity.) The invariant in these sets is the following: all nodes in each set form a clique; each node in A is placed at some distance $> r$ and $\leq r + \varepsilon$ from the point x and each node in B , and within distance ε from points x', y' ; each node in B is placed within distance r of point x and the set B contains point y ; each node in B' is placed within distance ε of each node in B and $> r$ from points x, x', y' and each node in A ; each node in C is placed at some distance $> r$ and $\leq r + \varepsilon$ from the point x and distance $> r$ from each node in A and points x', y' . The pairs of points x, x' and y, y' are located at distance r each.

For clarity, assume that n is even. At the beginning of the first time slot, the adversary places $n/2 - 1$ nodes, including the source node s , in set B' , an arbitrary node y in point y (i.e. in set B), an arbitrary node y' in point y' , and the remaining $n/2 - 1$ nodes in the set A , and starts up all nodes. Set C is initially empty. Given that $d \geq r + \varepsilon$, all nodes must be covered to solve the problem. Also, ε is set appropriately so that a node can be moved ε distance in one time slot without exceeding v_{\max} , and so that a node can be moved from set A to point x through the curved part of the dotted line, of length less than $\pi(r + \varepsilon)/6$, in α time slots without exceeding v_{\max} . (To see why the length bound is that, it is useful to notice that the distance between each pair of singular points along the circular dotted line is upper bounded by $(r + \varepsilon)/2$.)

Let t be the first time slot when the informed source transmits. Until time slot t , nodes are not moved, hence, (α, β) -connectivity has been preserved, and only the nodes in $B \cup B'$ become informed at time slot t . At the end of time slot t , node y is moved to B' and node y' is moved to A . Starting at time slot $t + 1$, the adversary moves the nodes so that only one new node at a time becomes informed while preserving (α, β) -connectivity.

First, we give a rough description of the nodes' movements, while the details will be presented later. Some of the nodes in B' are moved back and forth to B to produce contention. Nodes in A are moved one by one following the dotted lines in two phases, first up to point x , and afterwards to the set C . Right when a node moves to point x , another node moves from A to point x' so that (α, β) -connectivity is preserved. While a node moves from point x to C the node in point x' moves to x . The procedure is repeated until all nodes in A are covered.

The movement of each node u moved from A to C is carried out in two phases of α time slots each separated by an interlude as follows.

- *Phase 1.* During the first $\alpha - 1$ time slots, u is moved from A towards the point x maintaining a distance $> r$ and $\leq r + \varepsilon$ with respect to every node

- in B . Nodes in B' stay static during this interval. Given that only nodes in B' are informed and the distance between them and u is larger than r , u does not become covered during this interval. In the α -th time slot of this phase, u is moved to the point x , some node $v \in A$ is moved to the point x' , and any node $y \in B'$ is moved to B , preserving (α, β) -connectivity. Upon reaching point x , all the nodes that are not in $B \cup B'$ remain static until Phase 2. Phase 1 lasts only α slots, hence (α, β) -connectivity is preserved.
- *Interlude.* During this interval, nodes in B' are moved back and forth to B according to protocol II to produce contention as follows. For each time slot where nodes transmit with probability at least $8 \ln(n/2)/n$, the adversary moves all nodes in B' to B . At the end of the time slot the adversary moves all nodes in $B - \{y\}$ back to B' , and the procedure is repeated until u is covered. At the end of such time slot all nodes in B are moved to B' and the interlude ends. During the interlude, all partitions are connected, preserving (α, β) -connectivity.
 - *Phase 2.* In the first time slot of this phase, node u moves ε distance away from node v towards set C so that v is not informed by u . During the following $\alpha - 1$ slots, u is moved towards the set C while v is moved from x' towards the point x maintaining a distance $> r$ between them, while u maintains a distance $\leq r$ from $B \cup B'$. That is, Phase 2 of node u is executed concurrently with Phase 1 of node v (hence, nodes in B' stay static during this interval). At the end of this phase u is placed in set C and stays static forever. At this point node v has reached point x , but u cannot cover v because all nodes in C are at distance greater than r from x . Phase 2 lasts only α slots, hence (α, β) -connectivity is preserved.

The movement detailed above is produced for each node initially in A , overlapping the phases as described, until all nodes have become covered. In each phase of at least α time slots every node is moved a distance at most $\pi(r+\varepsilon)/6+\varepsilon$. Thus, setting ε appropriately, the adversarial movement described does not violate v_{\max} .

We prove now the time bound. For any time slot t in the interludes, the probability of covering the node in x is $P = \sum_{i \in B} \pi_t \prod_{j \in B, j \neq i} (1 - \pi_t) = |B| \pi_t (1 - \pi_t)^{|B|-1}$. For any t when $\pi_t < 8 \ln(n/2)/n$, we have $P < 8 \ln(n/2)/n$ because in this case the adversary puts just a single node in B , which is y . On the other hand, for any t when $\pi_t \geq 8 \ln(n/2)/n$, we also have $P \leq (n/2) (1 - 8 \ln(n/2)/n)^{n/2-1}$ since $8 \ln(n/2)/n < 1$ for $n > 24$. Using that $1 - x \leq e^{-x}$ for any $0 \leq x < 1$, we have $P \leq n / (2e^{8(n/2-1) \ln(n/2)/n}) \leq 8 \ln(n/2)/n$.

Let X be a random variable denoting the number of successful transmissions along $T = n^2 / (96 \ln(n/2))$ interlude steps. The expected number of successful transmissions is $E[X] = TP \leq n/12$. Given that X is the sum of independent Poisson trials, using Chernoff bounds, For $n/2 \geq 6E[X]$, $Pr(X \geq n/2) \leq 2^{-n/2}$. We conclude that T interlude steps are necessary to cover all nodes in A with probability at least $2^{-n/2}$. On the other hand, Phase 1 of all nodes in A adds $\alpha n/2$ steps to the overall time. Thus, the claim follows. \square

We move now to prove an existential lower bound for Dissemination oblivious protocols, possibly not fair. The proof, left to the full version of this paper for brevity, uses the same adversarial configuration of Theorem 4, but given that nodes running oblivious protocols may use different probabilities of transmission in a given time slot, node y has to be chosen more carefully as in Theorem 2.

Theorem 5. *For any $n > 3$, $d > r$, $\alpha \geq 0$, $\beta > 0$, $v_{\max} > \pi r/(6\alpha)$, and any oblivious randomized Geocast protocol Π , there exists an (α, β) -connected MANET of n nodes for which, Π takes at least $\alpha n/2 + n^2/(48e \ln(n/2))$ time slots in order to solve the problem with probability at least $2^{-n/2}$, even if all nodes are activated simultaneously and do not fail.*

The following theorem for locally adaptive protocols, can be proved as a straightforward repeated application of Theorem 3 to the configuration and movement of nodes described in Theorem 5, changing the minimum contention under which informed nodes are moved appropriately. The complete proof is omitted for brevity.

Theorem 6. *For any $n > 17$, $d > r$, $\alpha \geq 0$, $\beta > 0$, $v_{\max} > \pi r/(6\alpha)$, and any locally adaptive randomized Geocast protocol Π , there exists an (α, β) -connected MANET of n nodes for which, Π takes on expectation at least $\alpha n/2 + e^2(e + 1)^2 n^2 / (2(e - 1)^2 \ln(n/2))$ time slots in order to solve the problem, even if all nodes are activated simultaneously and do not fail.*

6 Upper Bound

The Dissemination protocol analyzed in this section is fair and oblivious. The algorithm is straightforward: every informed node transmits the information at each time step with probability $p = \ln n/n$. The same protocol was applied to Broadcast in Dynamic Networks in [8] yielding $O(n^2/\log n)$ with high probability. Here, the following theorem shows an upper bound parameterized in α and β for Dissemination in MANETs. For the upper bound we assume that failing nodes do not lose the information. This is needed in order to make Dissemination solvable because, otherwise, e.g., the adversary may turn off the source node in the first time step. The proof is left to the full version of this paper for brevity.

Theorem 7. *For any (α, β) -connected MANET where $\beta \geq 1$, $n > 2$, and any $v_{\max} > 0$, the fair-oblivious randomized protocol described above solves Dissemination in time $O(\alpha n + (1 + \alpha/\beta)n^2/\log n)$ time slots with probability at least $1 - e^{-(n-1)/4}$, even if nodes are activated at different times and fail.*

References

1. N. Alon, A. Bar-Noy, N. Linial, and D. Peleg. A lower bound for radio broadcast. *J. Comput. Syst. Sci.*, 43:290–298, 1991.

2. R. Baldoni, A. F. Anta, K. Ioannidou, and A. Milani. The impact of mobility on the geocasting problem in mobile ad-hoc networks: Solvability and cost. *Theor. Comput. Sci.*, 412(12-14):1066–1080, 2011.
3. R. Bar-Yehuda, O. Goldreich, and A. Itai. On the time-complexity of broadcast in multi-hop radio networks: An exponential gap between determinism and randomization. *Journal of Computer and System Sciences*, 45:104–126, 1992.
4. H. Baumann, P. Crescenzi, and P. Fraigniaud. Parsimonious flooding in dynamic graphs. In *Proc. Principles of Distributed Computing*, pages 260–269, 2009.
5. B. Chlebus. Randomized communication in radio networks. In *Handbook on Randomized Computing*, volume 1, pages 401–456. 2001.
6. M. Chrobak, L. Gąsieniec, and W. Rytter. Fast broadcasting and gossiping in radio networks. In *Proc. of Symp. on Foundations of Computer Science*, 2000.
7. A. Clementi, C. Macci, F. Pasquale, A. Monti, and R. Silvestri. Flooding time in edge-markovian dynamic graphs. In *Proc. of PODC*, pages 213–222, 2008.
8. A. Clementi, F. Pasquale, A. Monti, and R. Silvestri. Communication in dynamic radio networks. In *Proc. of PODC*, pages 205–214, 2007.
9. A. Czumaj and W. Rytter. Broadcasting algorithms in radio networks with unknown topology. In *Proc. of Symp. on Foundations of Computer Science*, 2003.
10. K. Fall. A delay-tolerant network architecture for challenged internets. In *Proc. of Conf. on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pages 27–34, 2003.
11. M. Farach-Colton, R. J. Fernandes, and M. A. Mosteiro. Lower bounds for clear transmissions in radio networks. In *Proc. of the 7th Latin American Theoretical Informatics Symp.*, volume 3887 of *Lecture Notes in Computer Science*, pages 447–454. Springer-Verlag, Berlin, 2006.
12. A. Fernández Anta, A. Milani, M. A. Mosteiro, and S. Zaks. Opportunistic information dissemination in mobile ad-hoc networks: The profit of global synchrony. In *Proc. of the 24th Symp. on Distributed Computing*, volume 6343 of *Lecture Notes in Computer Science*, pages 374–388. Springer-Verlag, Berlin, 2010.
13. D. R. Kowalski and A. Pelc. Broadcasting in undirected ad hoc radio networks. In *Proc. 22nd Ann. ACM Symp. on Principles of Distributed Computing*, 2003.
14. D. R. Kowalski and A. Pelc. Time complexity of radio broadcasting: adaptiveness vs. obliviousness and randomization vs. determinism. *Theoretical Computer Science*, 333:355–371, 2005.
15. F. Kuhn, T. Locher, and R. Oshman. Gradient clock synchronization in dynamic networks. In *Proc. of 21st Symp. on Parallelism in Algorithms and Architectures*, pages 270–279, 2009.
16. F. Kuhn, N. Lynch, and R. Oshman. Distributed computation in dynamic networks. In *Proc. of Symp. on Theory of Computing*, pages 513–522, 2010.
17. E. Kushilevitz and Y. Mansour. An $\Omega(D \log(N/D))$ lower bound for broadcast in radio networks. *SIAM Journal on Computing*, 27(3):702–712, 1998.
18. R. Mellier and J.-F. Myoupo. A mobility-management mechanism for broadcasting in unknown mobile ad hoc networks. In *Proc. of Wksp. on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks*, pages 242–245, 2005.
19. R. O’Dell and R. Wattenhofer. Information dissemination in highly dynamic graphs. In *Proc. of DIAL-M/POMC*, pages 104–110, 2005.
20. L. Pelusi, A. Passarella, and M. Conti. Opportunistic networking: data forwarding in disconnected mobile ad hoc networks. *Communications Magazine, IEEE*, 44(11):134–141, 2006.
21. L. G. Roberts. Aloha packet system with and without slots and capture. *Computer Communication Review*, 5(2):28–42, 1975.