VOLUME X, NUMBER 1

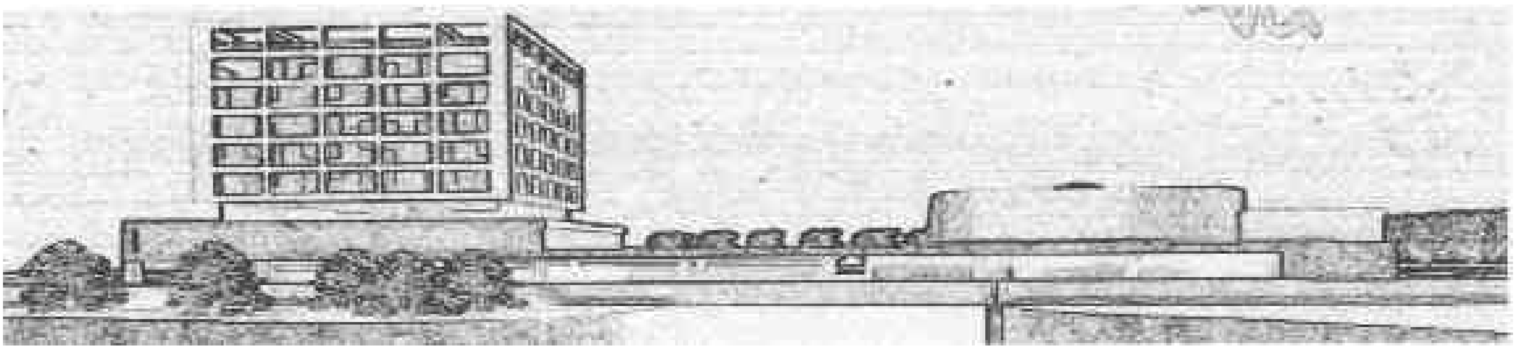

REPORTS ON SYSTEMS AND COMMUNICATIONS

Opportunistic Information Dissemination in Mobile Ad-hoc Networks: The Profit of Global Synchrony

Antonio Fernández Anta, Alessia Milani, Miguel A. Mosteiro, and Shmuel Zaks

Móstoles (Madrid), February 2010
Depósito Legal: M-50653-2004
ISSN: 1698-7489

Table of Contents

Opportunistic Information Dissemination in Mobile Ad-hoc Networks: The Profit of Global Synchrony

Antonio Fernández Anta⋆[1], Alessia Milani⋆⋆[2], Miguel A. Mosteiro⋆⋆⋆[1], and Shmuel Zaks[3]

[1] LADyR, GSyC, Universidad Rey Juan Carlos, 28933 Móstoles, Madrid, Spain
anto@gsyc.es, miguel.mosteiro@urjc.es
[2] LIP6, Universit Pierre et Marie Curie - Paris 6, Paris, France
alessia.milani@lip6.fr
[3] Dept. of Computer Science, Technion - Israel Inst. of Technology, Haifa, Israel
zaks@cs.technion.ac.il

Abstract. The topic of this paper is the study of *Information Dissemination* in Mobile Ad-hoc Networks by means of deterministic protocols. In this problem, a piece of information initially held by a source node has to be propagated to a given set of destination nodes. We consider multihop dynamic networks, in which n nodes move in the plane, join, and leave. Our connectivity model only requires that, within α time slots, some node that has the information must be connected to some node without it for at least β time slots. In this context, the protocols considered make use of *opportunistic communication* during this connectivity interval to make progress in the dissemination. The protocols studied are classified into three classes: *oblivious* (the transmission schedule of a node is only a function of its ID), *quasi-oblivious* (the transmission schedule may also depend on a global time), and *adaptive*.

A collection of interesting complexity gaps among protocol-classes is observed in this work. In order to guarantee any progress towards solving the problem, it is shown that β must be at least $n-1$ in general, but that this bound becomes $\beta \in \Omega(n^2/\log n)$ if an oblivious protocol is used. Since quasi-oblivious protocols can guarantee progress with $\beta \in O(n)$, this represents a significant gap, almost linear on β, between oblivious and quasi-oblivious protocols. Regarding the time to complete the dissemination, a lower bound of $\Omega(n\alpha+n^3/\log n)$ is proved for oblivious protocols, while a constructive $O(n\alpha+n^3\log n)$ upper bound is shown for the same class (which is tight up to a polylogarithmic factor). It is also proved that quasi-oblivious protocols achieve $O(n\alpha+n^2)$, which is shown to be optimal by giving a matching lower bound for adaptive protocols.

These results show that the gap in time complexity between oblivious and quasi-oblivious, and hence adaptive, protocols is almost linear. This gap is what we call the *profit of global synchrony*, since it represents the gain the network obtains from global synchrony with respect to not having it. We note that the global synchrony required by the efficient quasi-oblivious protocol proposed is simply achieved by piggybacking in the messages sent the time at the source node, as a global reference.

⋆ Research partially supported by Comunidad de Madrid grant S2009TIC-1692 and Spanish MICINN grant TIN2008-06735-C02-01.
⋆⋆ Research partially supported by ANR VERSO Shaman project.
⋆⋆⋆ Research partially supported by EU Marie Curie European Reintegration Grant IRG 210021; NSF grant CCF 0632838; and Comunidad de Madrid grant S-0505/TIC/0285.

1 Introduction

A Mobile Ad-hoc Network (aka MANET) is a set of mobile nodes which communicate over a multihop radio network, without relying on a stable infrastructure. In these networks, nodes are usually battery-operated devices that can communicate via radio with other devices that are in range. Due to unreliable power supply and mobility, nodes may have a continuously changing set of neighbors in that range. This dynamic nature makes challenging to solve even the simplest communication problems in general. Hence, proposed protocols often have strong synchronization and stability requirements, like having a stable connected network for long enough time.

Current trends in networking-architecture developments, like *delay and disruption tolerant networks*, and *opportunistic networking* [17,37], aim to deal with the disconnections that naturally and frequently arise in wireless environments. Their objective is to allow communication in dynamic networks, like a MANET, even if a route between sender and receiver never exists in the network. The result is that multi-hop communication is provided through *opportunistic communication*, in which the *online route* of a message is followed one link at a time, as links in the route become available. While the next link is not available, the message is held in a node. With opportunistic communication, strong connectivity requirements are no longer needed. Furthermore, in some cases mobility is the key to allow communication (e.g., consider two disconnected static nodes, where communication between them is provided by a device that, due to mobility, sometimes is in range of one and sometimes of the other).

In this paper, we formally define a particular class of MANET which is suited for opportunistic communication, and which we call *potentially epidemic*. A MANET is potentially epidemic if the changes in the communication topology are such that an online route exists among any two nodes that wish to communicate. The network is potentially epidemic because the actual propagation of the information on the online routes depends on the stability of each of its communication links.

In this context, we define and study the deterministic solvability of a problem that we call *Dissemination*. In this problem, at a given time a source node holds an information that must be disseminated to a given set of nodes belonging to the MANET. The nodes elected to eventually receive the information are the ones that satisfy a given predicate. Depending on this predicate, the Dissemination problem can instantiate most of the common communication problems in distributed systems, such as Broadcast, Multicast, Geocast, Routing, etc.

In particular, we determine assumptions on link stability and speed of nodes under which a distributed deterministic protocol exists that solves Dissemination in potentially epidemic networks. Moreover, we relate the time complexity of the solution to the speed of movement and to the information that protocols may use.

1.1 The Dissemination Problem

We study the problem of disseminating a piece of information, initially held by a distinguished source node, to all nodes of a given set in the network. Formally,

Definition 1. *Given a MANET formed by a set V of n nodes, let $\mathcal{P}$ be a predicate on V and $s \in V$ a node that acquires a piece of information I at time t_1 (s is the source of dissemination). The* Dissemination *problem consists of distributing I to the set of nodes $V_{\mathcal{P}} = \{x \in V :: \mathcal{P}(x)\}$. A node that has received I is termed* covered*, and otherwise it is* uncovered*. The Dissemination problem is solved at time slot $t_2 \geq t_1$ if, for every node $v \in V_{\mathcal{P}}$, v is covered by time slot t_2.*

The Dissemination problem abstracts several common problems in distributed systems. E.g. Broadcast, Multicast, Geocast, Routing etc, are all instances of this problem for a particular predicate $\mathcal{P}$. In order to prove lower bounds, we will use one of these instances: the Geocast problem. The predicate $\mathcal{P}$ for Geocast is $\mathcal{P}(x) = \mathrm{true}$ if and only if, at time t_1, x is up and running, and it is located within a parametric distance $d > 0$ (called *eccentricity*) from the source node.

1.2 Model

We consider a MANET formed by a set V of n mobile nodes deployed in $\mathbb{R}^2$, where no pair of nodes can occupy the same point in the plane simultaneously. It is assumed that each node has data-processing and radio-communication capabilities, and a unique identificator number (ID) in $[n] \triangleq \{1, \ldots, n\}$.

Time. Each node is equipped with a clock that ticks at the same uniform rate ρ but, given the asynchronous activation, the clocks of different nodes may start at different times. A time interval of duration $1/\rho$ is long enough to transmit (resp. receive) a message. Computations in each node are assumed to take no time. Starting from a time instant used as reference, the global time is slotted as a sequence of time intervals or *time slots* $1, 2, \ldots$, where slot $i > 0$ corresponds to the time interval $[(i-1)/\rho, i/\rho)$. Without loss of generality [39] all node's ticks are assumed to be in phase with this global tick.

Node Activation. We say that a node is *active* if it is powered up, and *inactive* otherwise. It is assumed that, due to lack of power supply or other unwanted events that we call *failures*, active nodes may become inactive. Likewise, due also to arbitrary events such as replenishing their batteries, nodes may be re-activated. We call the temporal sequence of activation and failures of a node the *activation schedule*. The activation schedule for each node is assumed to be chosen by an adversary, in order to obtain worst-case bounds. Most of the lower-bound arguments included in this paper hold, even if all nodes are activated simultaneously and never fail (which readily provide a global time), making the results obtained stronger.

We assume that a node is activated in the boundary between two consecutive time slots. If a node is activated between slots $t-1$ and t we say that it is activated at slot t, and it is active in that slot. Upon activation, a node immediately starts running from scratch an algorithm previously stored in its hardware, but no other information or status is preserved while a node is inactive. Consequently, it is possible that a covered node does not hold the information I, because it has been inactive after receiving it. To distinguish a covered node that does not hold the information from one that holds it, we introduce the following additional terminology: we say that a node p is *informed* at a given time t if it holds the information I at time t, otherwise p is said to be *uninformed*.

Radio Communication. Nodes communicate via a collision-prone single radio channel. A node v can receive a transmission of another node u in time slot t only if their distance is at most the *range of transmission* r during the whole slot t. The range of transmission is assumed to be the same for all nodes and all time slots. If two nodes u and v are separated by a distance at most r, we say that they are *neighbors*. In this paper, no collision detection mechanism is assumed, and a node cannot receive and transmit at the same time slot. Therefore, an active node u receives a transmission from a neighboring node v at time slot j if and only if v is the only node in u's neighborhood transmitting at time slot j. Also, a node cannot distinguish between a collision and no transmission.

Link stability. We assume that nodes may move on the plane they are deployed. Thus, the topology of the network is time dependent. For simplicity, we assume that the topology only changes in the boundaries between time slots. Then, at time slot t nodes u and v are connected by a link in the network topology iff they are neighbors during the whole slot t. An online route between two nodes u and v is a path $u = w_0, w_1, \ldots, w_k = v$ and a sequence of time slots $t_1 < t_2 < \cdots < t_k$ such that the network has a link between w_{i-1} and w_i at time slot t_i. Observe that in order to be able to solve an instance of Dissemination, it is necessary that the network is potentially epidemic. I.e. after the initial time t_1, there is an online route from the source s to every node in $V_{\mathcal{P}}$. However, as argued in [14], worst-case adversarial choice of topologies for a dynamic network precludes any deterministic protocol from completing Broadcast, even if connectivity is guaranteed. Note that Broadcast is an instance of Dissemination, and that if there is connectivity then there are online routes between all nodes. Thus, the property that the network is potentially epidemic as described is not sufficient to solve Dissemination, and further limitations to the adversarial movement and activation schedule are in order. While respecting a bound on the maximum speed $v_{\max}$, which is a parameter, the adversarial movement and activation schedule is limited by the following connectivity property

Definition 2. *Given a Mobile Ad-hoc Network, an instance of the Dissemination problem that starts at time* t_1*, and two integers* $\alpha \geq 0$ *and* $\beta > 0$*, the network is* (α, β)-connected *if, for every time slot* $t \geq t_1$ *at which the problem has not yet been solved, there is a time slot* t' *such that the following conditions hold:*

- *the intersection of time intervals* $[t, t + \alpha)$ *and* $[t', t' + \beta)$ *is not empty, and*
- *there is a pair of nodes* p, p'*, such that at* t' p *is informed and* p' *is uncovered, and they are active and neighbors the whole time interval* $[t', t' + \beta)$*.*

Observe that the (α, β)-connectivity property by itself does not guarantee that the network is epidemic, i.e. that the information is eventually disseminated. This can be seen, for instance, if the source s never gets to transmit the information I to any node. In fact, thanks to the (α, β)-connectivity, at most every α slots, the source s is connected to other nodes of the network for at least β time slots. But, we have progress only if the protocol to solve Dissemination is able to use the β slots of connectivity to cover some uncovered node. This is why, we call *potentially* epidemic all networks that ensure the (α, β)-connectivity property.

1.3 Protocols for Dissemination

We consider distributed deterministic protocols, i.e., we assume that each node in the network is preloaded with its own and possibly different deterministic algorithm that defines a schedule of transmissions for it. Even if a transmission is scheduled for a given node at a given time, that node will not transmit if it is uninformed.

Following the literature on various communication primitives [29,30], a protocol is called *oblivious* if, at each node, the algorithm's decision on whether or not to schedule a transmission at a given time slot depends only on the identifier of the node, and on the number of time slots that the node has been active. Whereas, if no restriction is put on the information that a node may use to decide its communication schedule, the protocol is called *adaptive*. Additionally, in this paper, we distinguish a third class of protocols that we call *quasi-oblivious*. In a quasi-oblivious protocol the sequence of scheduled transmissions of a node depends only on its ID and a global time. Quasi-oblivious protocols have been sometimes called oblivious, since the model assumed simultaneous activation, and hence a global time was readily available. However we prefer to make the difference explicit, as done in [36], because we found a drastic gap between this class and fully oblivious protocols.

1.4 Previous Work

A survey of the vast literature related to Dissemination is beyond the scope of this article. We overview in this section the most relevant previous work. Additionally, a review of relevant related work for static and dynamic networks beyond MANETs can be found in Section 1.5.

The Dissemination problem abstracts several common problems in Radio Networks. When some number $1 \leq k \leq n$ of active nodes hold an information that must be disseminated to all nodes in the network, the problem is called *k-Selection* [29] or *Many-to-all* [10]. If $k = 1$ the problem is called *Broadcast* [5, 31], whereas if $k = n$ the problem is known as *Gossiping* [8, 18]. Upper bounds for these problems in mobile networks may be used for Dissemination, and even those for static networks may apply if the movement of nodes does not preclude the algorithm from completing the task (e.g., round-robin). On the other hand, if only the subset of k nodes have to receive the information, the problem is known as *Multicast* [10, 23], and if only nodes initially located at a parametric distance from the source node must receive the information the problem is called Geocast [27], defined in Section 1.1.

Deterministic solutions for the problems above have been studied for MANETs. Their correctness rely on strong synchronization or stability assumptions. In [33], deterministic Broadcast in MANETs was studied under the assumption that nodes move in a one-dimensional grid knowing their position. Two deterministic Multicast protocols for MANETs are presented in [25, 35]. The solutions provided require the network topology to globally stabilize for long enough periods to ensure delivery of messages, and they assume a fixed number of nodes arranged in some logical or physical structure. Leaving aside channel contention, a lower bound of $\Omega(n)$ rounds of communication was proved in [38] for Broadcast in MANETs, even if nodes are allowed to move only in a two-dimensional grid, improving over the $\Omega(D \log n)$ bound of [7], where D is the diameter of the network. This bound was improved to $\Omega(n \log n)$ in [16] without using the movement of nodes, but the diameter of the network in the latter is linear. Recently, deterministic solutions for Geocast were proposed in [4] for a one-dimensional setting and in [19] for the plane. In the latter work, the authors concentrate in the structure of the Geocast problem itself, leaving aside communication issues such as the contention for the communication channel.

1.5 Related Work

Static networks Given that a static network is just an instance of a network where nodes are allowed to move, lower bounds for the Broadcast problem also may apply to Dissemination, but only if the network is geometric (i.e., the adjacencies of the topology assumed can be embedded in $\mathbb{R}^2$) and the target of the dissemination are all the nodes in the network. The following bounds are proved for models similar to the one assumed in this paper. Exploiting a lower bound on the size of a combinatorial problem called *selective families*, a lower bound for deterministic Broadcast of $\Omega(n \log D)$, where D is the diameter of the layered topology used, was shown in [12], yielding a linear lower bound if $D \in O(1)$. In [30], adaptive and oblivious deterministic lower bounds of $\Omega(n)$ and $\Omega(n \min\{D, \sqrt{n}\})$ respectively were presented for the Broadcast problem in static Radio Networks. The construction used for the former has diameter at most 4 and at least 7 for the latter. Extending the construction used previously in [7, 9] to geometric Radio Networks, a lower bound of $\Omega(n \log n)$ is shown in [16] for the Broadcast problem even if nodes do not move. The diameter of the network used in this case is linear. In [24], it was shown that, for each oblivious deterministic Broadcast algorithm, there exists a network of diameter 2 such that the running time is at least $\Omega(\sqrt{g})$, where g is the inverse of the minimum Euclidean distance between any pair of nodes, by adversarially placing nodes as the algorithm progress. The above bounds apply also to Geocast, but only if $d \geq D$ where D is also the eccentricity but with respect to all nodes in the constructions used to prove them.

Regarding upper bounds, in [9], Chlebus et al. presented an adaptive protocol that completes deterministic Broadcast in less than $14n$ steps where n is the number of nodes in the system, for a symmetric network without collision detection where nodes do not have any information of the network except their own unique identity (later denoted ID for short). Making use of the simultaneous activation of all nodes, the protocol defines an eulerian cycle over a spanning tree of the network. Hence, this protocol could be used for Dissemination only if the movement of nodes does not change the topology. None of these assumptions are present in our model. Limiting the adversary to changes of topology due to node failures only, it was proved in [13] that the well-known round-robin $O(Dn)$ algorithm, where D is the diameter of the network, is optimal for solving Broadcast restricted to connected non-failing nodes.

Dynamic Networks beyond MANETs Similar problems have been studied recently for dynamic networks. A suitable model for time-dependent topologies is the *Dynamic Graph* [6,11,15]. A dynamic graph is a set of n nodes V and a sequence of edge-sets $E_1, E_2, \ldots$ over V. Mapping each E_i with a time slot i, a dynamic graph models a network with a possibly different topology for each time slot, usually called a *Dynamic Network* [3,14]. If such topology can be embedded in $\mathbb{R}^2$, a dynamic network is a suitable model of a MANET (even under failures since a node that is inactive at time step i may be modeled by not including any of its edges to neighboring nodes in E_i). Dynamic networks may be *eventually connected* [2]. In other words, the temporal sequence of edge-sets may be such that no edge-cut lasts forever, as assumed in *population protocols* [1]. Communication primitives have been studied for dynamic networks [6, 11, 14, 15, 34]. In [6, 11, 14, 15] the results obtained are stochastic and, as argued in [14], worst-case adversarial choice of topologies precludes any deterministic protocol from completing Broadcast, even if connectivity is guaranteed. In [34], they consider deterministic solutions to solve two particular instances of the dissemination problem, e.g., routing and flooding. They do not explicitly deal with the contention for the communication channel, and they request the network to be always connected and to provide some local stability.

1.6 Our Results

In the following, we summarize our results. For a model where nodes may fail, there is no global clock, and nodes may be activated at different times, an upper bound of $n(\alpha + 4n(n-1)\ln(2n))$ for Dissemination is proved in Theorem 8 by means of an oblivious deterministic protocol based on Primed Selection [20]. For the same model, we show in Theorem 6 that any oblivious protocol takes at least $\Omega((\alpha + n^2/\ln n)n)$ steps to solve the Geocast problem if $v_{max} > \pi r/6(\alpha + \lfloor (n/3)(n/3-1)/\ln(n/3(n/3-1)) \rfloor - 2)$. Observe that these results are tight up to a poly-logarithmic factor.

Moreover, in the same model, there is a quasi-oblivious protocol that solves Dissemination in at most $n(\alpha + n)$ steps as shown in Theorem 7. This algorithm mainly consists of a first step where each node synchronizes its clock with the source, and then the execution of the classic Round Robin protocol. The global synchronization signal is carried by the messages exchanged for dissemination and it is the time elapsed since the source started the dissemination. On the other hand, Theorem 5 shows that, even if nodes are activated simultaneously and do not fail, and an adaptive protocol is used, any Geocast protocol takes at least $\Omega(n(\alpha + n))$ if $v_{max} > \pi r/(3(2\alpha + n - 4))$.

The latter results are asymptotically tight and show that full adaptiveness does not help with respect to quasi-obliviousness. The first lower bound and the last upper bound, show an asymptotic separation almost linear between oblivious and quasi-oblivious protocols. In a more restrictive model, where nodes are activated simultaneously, it exists an oblivious protocol, e.g., the well-known *Round Robin* protocol, that solves Dissemination in at most $n(\alpha+n)$. This can be trivially derived from the fact that when nodes are all activated simultaneously, the above quasi-oblivious protocol does not need the synchronization step and thus simply reduces to the *Round Robin* protocol. Hence, the lower bound proved in Theorem 6 shows the additional cost of obliviousness when nodes are not simultaneously activated. This gap is what we call the *profit of global synchrony*, since it represents the gain the network obtains from global synchrony with respect to not having it. Moreover, the quasi-oblivious protocol shows that for the Dissemination problem, the simultaneous activation performance can be achieved by distributing the time elapsed since the source started the dissemination.

We note that the issue of who whether all, or only a subset of the nodes, start an execution, is of a central importance in distributed computing. In rare cases it might affect the solvability of the problem (see Example 1 below), and it certainly might affect the time (see Example 2 below) or the message complexity (see Example 3 below). All these examples refer to asynchronous networks, so the issue of simultaneity is different from the one used in our case. The examples are: (1) In [21] it is shown that no algorithm can solve the consensus problem if at least one node might fail during the execution; however, the result does not hold if the failure can occur only at the beginning of the execution. Actually, the problem can be solved as long as a majority of the processes are non-faulty, but those that fail do so at the beginning of the execution. (2) In [22] an algorithm to find a minimum spanning tree in a network is presented. It is shown that the execution can take $\Omega(n^2)$ time if the processors are allowed to start at different times; however, if they all start simultaneously then the time complexity is $O(n \log n)$. (n is the number of processors). (3) Finding a leader in a network whose topology is a complete graph is of message complexity $\Omega(n \log n)$ if all nodes start the execution [28], but can be solved by at most $O(k \log n)$ messages when only k nodes start [26].

Additionally, it is shown in Theorem 1 that no protocol can solve the Geocast problem (and hence Dissemination) in all (α, β)-connected networks unless $\beta \geq n - 1$. Interestingly, it is shown in Theorem 2 that this bound becomes $\beta > \lfloor (n-1)(n-3)/4 \ln((n-1)(n-3)/4) \rfloor$ if the protocol is oblivious. Comparing these bounds with the requirements of the protocols presented above, the quasi-oblivious protocol required $\beta \geq n$, which is almost optimal, while the oblivious protocol required $\beta \in \Omega(n^2 \log n)$, which is only a polylogarithmic factor larger than the lower bound. These results also expose another aspect of the profit of global synchrony mentioned before: while $\beta = n$ is enough for quasi-oblivious protocols to solve Dissemination, oblivious protocols require a value of β almost a linear factor larger.

Finally, for an arbitrary small bound on node speed, we show in Theorem 3 the existence of an (α, β)-connected network where Geocast takes at least $\alpha(n-1)$ steps, even using randomization; and the existence of an (α, β)-connected network where any deterministic protocol that transmits periodically takes at least $n(n-1)/2$ steps, even if nodes do not move, in Theorem 4.

1.7 Paper Organization

The rest of the paper is organized as follows. In Section 2 we introduce some technical lemmas that will be used to prove our main results; in Section 3 we prove the lower bound on link stability and on the time complexity to solve the Dissemination problem with respect to some important aspects of the system (e.g. speed of movement of nodes and their activation schedule) and of the protocols (e.g., obliviousness versus adaptiveness). We finally present the corresponding upper bounds in Section 4.

2 Auxiliary Lemmas

The following lemmas will be used throughout the analysis. A straightforward consequence of the pigeonhole principle is established in the following lemma.

Lemma 1. *For any time step t of the execution of a Dissemination protocol, where a subset V' of k informed nodes do not fail during the interval $[t, t+k-2]$, there exists some node $v \in V'$ such that v does not transmit uniquely among the nodes in V' during the interval $[t, t+k-2]$.*

In the following lemma, we show the existence of an activation schedule such that, for any *oblivious* deterministic protocol, within any subset of at least 3 nodes, there is one that does not have a unique transmission scheduled within a period roughly quadratic in the size of the subset.

Lemma 2. *For any deterministic oblivious protocol that solves Dissemination in a MANET of n nodes, where nodes are activated possibly at different times, and for any subset of k nodes, $k \geq 3$, there exists a node-activation schedule such that, for any time slot t and letting $m = \lfloor k(k-1)/\ln(k(k-1)) \rfloor$, each of the k nodes is activated during the interval $[t-m+1, t]$, and there is one of the k nodes that is not scheduled to transmit uniquely among those k nodes during the interval $[t, t+m-1]$.*

Proof. Consider any oblivious protocol Π and any subset of k nodes, where k is an arbitrary value such that $3 \leq k \leq n$. If, according to Π, the algorithm of one of the k nodes, call it i, does not include a transmission scheduled within any period of m steps within the first $2m$ steps, the claim holds by activating i so that such period starts at time slot t.

Otherwise, the algorithm of all the k nodes must include at least one transmission scheduled within any period of m steps within the first $2m$ steps. If, according to Π, the algorithm of one of the k nodes, call it i, includes less than k transmissions scheduled within any period of m steps within the first $2m$ steps, the claim holds by activating i so that such period starts at time slot t and the other $k-1$ nodes are activated during the interval $[t-m+1, t]$ so that each of i's transmissions are scheduled at the same time that some other transmission. This is possible because i is scheduled to transmit less than k times during the interval $[t, t+m-1]$ and each of the other $k-1$ nodes is scheduled to transmit at least once within any period of m steps within the first $2m$ steps.

Otherwise, the algorithm of all the k nodes must include at least k scheduled transmissions within any period of m steps within the first $2m$ steps, we say that the *density* of scheduled transmissions in such period is at least k/m. Then, we prove the existence of the claimed activation schedule by the probabilistic method. Consider any of the k nodes, call it i. Let i be activated at time slot t. Choosing uniformly at random, for each of the other $k-1$ nodes, a slot within the interval $[t-m+1, t]$ for activation, and using that the density of scheduled transmissions within the first $2m$ steps of any node is at least k/m, the probability that i has a unique (among the k nodes) transmission scheduled during the interval $[t, t+m-1]$ is less than $m(1-k/m)^{k-1}$. Then, in order to prove the claim, it is enough to show that the probability that i does not have a unique (among the k nodes) transmission scheduled at some time step within i's first m steps is strictly bigger than 0. Replacing, it is enough to prove $1-\lfloor k(k-1)/\ln(k(k-1))\rfloor(1-k\lfloor k(k-1)/\ln(k(k-1))\rfloor^{-1})^{k-1} > 0$. Using the inequality $1-x \leq e^{-x}$, for $0 < x < 1$ [32, §2.68], it can be seen that the inequality holds for any $k \geq 3$. □

3 Solvability of the Dissemination Problem

If there is at least one node in $V_{\mathcal{P}} - \{s\}$ at least one time slot is needed to solve Dissemination, since the source node has to transmit at least once to pass the information. Furthermore, if all nodes in $V_{\mathcal{P}}$ are neighbors of s, one time slot may also be enough if the source node transmits before neighboring nodes are able to move out of its range. On the other hand, if the latter is not possible, more than one time slot may be needed. Let us consider the Geocast problem. Given that the specific technological details of the radio communication devices used determine the minimum running time when the eccentricity is $d \leq r$, all efficiency lower bounds are shown for $d > r$ unless otherwise stated.

3.1 Link Stability Lower Bounds

The following theorem shows a lower bound on the value of β for the Geocast problem.

Theorem 1. *For any $v_{max} > 0$, $d > r$, and $\alpha > 0$, if $\beta < n-1$, there exists an (α, β)-connected MANET of n nodes such that no deterministic protocol for Geocast exists, even if all nodes are activated simultaneously and do not fail.*

Proof. Consider three sets of nodes A, B, and C deployed in the plane, each set deployed in an area of size ε arbitrarily small, such that $0 < \varepsilon < r$ and $d \geq r+\varepsilon$. (See Figure 1.) 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 ε from every node in C, and every node in A is placed at some distance $r < \delta \leq r+\varepsilon$ from every node in C. At the beginning of the first time slot, the adversary places $n-1$ nodes, including the source node s, in the set C, the remaining node x in set A, and activates all nodes. The set B is initially empty. Given that $d \geq r+\varepsilon$, x must become informed to solve the problem . Also, ε is set appropriately so that a node can move ε distance in one time slot without exceeding v_{max}.

For any protocol Π for Geocast, let t be the first time slot where the source node is the only node to transmit in the set C. Adversarially, let t be the first time slot when the source is informed. Thus, (α, β)-connectivity is preserved up to time slot t for any α. At time slot t, all nodes placed in C are informed.

After time slot t, the adversary moves the nodes as follows. Given that the problem was not solved yet and nodes in C do not fail, according to Lemma 1,there exists a node $y \in C$ that does not transmit uniquely among the nodes in C during the interval $[t+1, t+n-2]$. Given that Π is a deterministic protocol, and the adversary knows the protocol and defines the movement of all nodes, the adversary knows which is the node y.

Assume, for the sake of contradiction, that $\beta \leq n-2$. Then, the adversary places y in B for all time slots in the interval $[t+1, t+\beta]$. Additionally, for each time slot $t' \in [t+1, t+\beta]$ where y transmits, the adversary moves to B some node $z \in C$ that transmits at t' to produce a collision. At the end of each time slot t' the adversary moves z back to C. Such a node z exists since y does not transmit uniquely during the interval $[t+1, t+n-2]$ and $n-2 \geq \beta$.

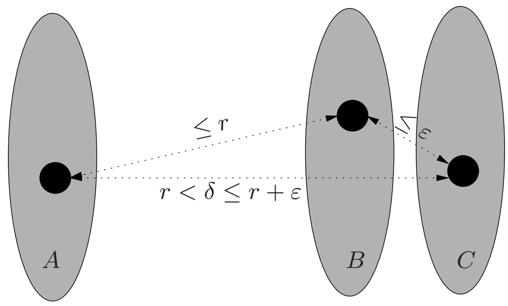

Fig. 1. Illustration of Theorem 1.

At the end of time slot $t + \beta$, the adversary moves y back to C and the above argument can be repeated forever preserving the (α, β)-connectivity and precluding Π from solving the problem. Therefore, β must be at least $n - 1$. □

Building upon the argument used in the previous theorem, but additionally exploiting the adversarial node activation, the following theorem shows a lower bound on the value of β for the Geocast problem if the protocol used is oblivious.

Theorem 2. *For any $v_{max} > 0$, $d > r$, $n \geq 8$ and $\alpha > 0$, if $\beta \leq m = \lfloor (n-1)(n-3)/4 \ln((n-1)(n-3)/4) \rfloor$, there exists an (α, β)-connected MANET of n nodes such that no deterministic oblivious protocol for Geocast exists.*

Proof (Sketch). Consider again the configuration described in Theorem 1. (See Figure 1.) The invariant, the initial location of nodes, and the value of ε are the same as in Theorem 1. The adversarial node activation schedule follows.

The adversary activates x and s at the beginning of some time slot t_0. Let $t_1 \geq t_0$ be the time slot when the source node is scheduled to transmit the information for the first time. Adversarially, let t_1 be the first time slot when the source is informed. Thus, (α, β)-connectivity is preserved up to time slot t_1 for any α.

Consider a partition $\{C_o, C_e\}$ of the nodes in C so that $|C_o| = |C_e| = (n-1)/2$. (For clarity assume that n is odd.) Nodes in $C_o - \{s\}$ are activated during the interval $[t_1 - m + 1, t_1]$ so that some node $y \in C_o$ does not transmit uniquely among the nodes in C_o during the interval $[t_1 + 1, t_1 + m]$. As proved in Lemma 2, such a node exists because the protocol is oblivious and $n \geq 8$. Given that Π is a deterministic oblivious protocol known by the adversary, the adversary knows which is the node y.

Let $t_2 > t_1 + m$ be the time slot when y transmits uniquely among the nodes in C_o for the first time after t_1. The adversary places y in B for all time slots in the interval $[t_1 + 1, t_2)$. Additionally, for each time slot $t' \in [t_1 + 1, t_2)$ where y transmits, the adversary moves to B some node $z \in C_o$ that transmits at t' to produce a collision at x. At the end of each time slot t' the adversary moves z back to C_o. Such a node z exists since y does not transmit uniquely during the interval $[t_1 + 1, t_2)$. Right before time slot t_2, the adversary moves y back to C_o precluding x from becoming covered. Right after time slot t_2 when y transmits the information, the adversary deactivates all nodes in C_o.

For the following interval $[t_2 + 1, t_2 + m]$, the above argument can be repeated using the nodes in the set C_e. At time slot t_2, the transmission of node y will inform all nodes in C_e.

Thus, activating appropriately the nodes in C_e during the interval $[t_2 - m + 1, t_2]$, some node $y' \in C_e$ does not transmit uniquely among the nodes in C_e during the interval $[t_2 + 1, t_2 + m]$. Thus, moving nodes between sets C_e and B as described before once again x is not informed, this time during the interval $[t_2+1, t_3)$, where $t_3 > t_2+m$ is the time slot when y' is scheduled to transmit uniquely among the nodes in C_e.

The above argument can be repeated inductively forever so that the problem is not solved but, if $\beta \leq m$, (α, β)-connectivity is preserved. Therefore, β must be bigger than m. □

3.2 Time Complexity Lower Bounds versus Speed, Activation and Obliviousness

Exploiting the maximum time α that a partition can be disconnected, a lower bound on the time efficiency of any protocol for Geocast, even regardless of the use of randomization and even for arbitrarily slow node-movement, can be proved. The following theorem establishes that bound.

Theorem 3. *For any $v_{max} > 0$, $d > r$, $\alpha > 0$, and $\beta > 0$, there exists an (α, β)-connected MANET of n nodes, for which any Geocast protocol takes at least $\alpha(n-1)$ time slots , even if all nodes are activated simultaneously and do not fail.*

Proof. The following adversarial configuration and movement of nodes shows the claimed lower bound. Initially, nodes are placed in a circle (as depicted in Figure 2), all are activated simultaneously, and let the source be informed immediately upon activation. Let $\{v_1, v_2, \ldots, v_{n-1}, v_n\}$ be the nodes located around the circle in clockwise direction where $v_1 = s$ is the source node. The nodes are located so that the distances between each pair are the following. The distance between v_1 and v_2 is $r + \varepsilon$, where $\varepsilon > 0$ is set appropriately so that a node can move as described below in one time slot without exceeding v_{max}. For each $1 < i < n$, the distance between v_i and v_{i+1} is r. The pair v_1, v_n is placed at distance δ, $r < \delta < d - \varepsilon'$, where $\varepsilon' > 0$ is set appropriately so that node v_n can move as described below in one time slot without exceeding v_{max} and it will be always within distance d from the source $s = v_1$. Every other pair of nodes is separated by distance bigger than r. In order to solve Geocast, node v_n must be covered.

Nodes stay in the position described until time slot $\alpha - 1$ when node v_2 moves so that in time slot α it is placed at distance r of v_1 and at distance $r + \varepsilon$ from v_3. Node v_2 does not move after time slot α. Due to this configuration, the source node is not able to inform v_2 during the first $\alpha - 1$ time slots and, in the best case, v_2 becomes informed in time slot α. After v_2 becomes informed, the argument can be repeated iteratively for $v_3, v_4, \ldots, v_n$. Given that each node, except for the source, becomes informed at least α time slots after its predecessor, the overall time is lower bounded as claimed. Given that a new uncovered node becomes a neighbor of an informed node within α slots and does not move after that, and all other partitions are always connected, (α, β)-connectivity is preserved.

□

The linear lower bound for Geocast proved in Theorem 3 was shown exploiting the maximum time of disconnection between partitions. Exploiting the adversarial schedule of node activation, a quadratic bound can be shown for the important class of *equiperiodic* protocols, even if nodes do not move. The protocol definition and the theorem follows.

Definition 3. *A protocol of communication for a Radio Network is* equiperiodic *if for each node, the transmissions scheduled are such that the number of consecutive time steps without transmitting, call it $T-1$, is always the same. We say that T is the* period *of transmission of such node.*

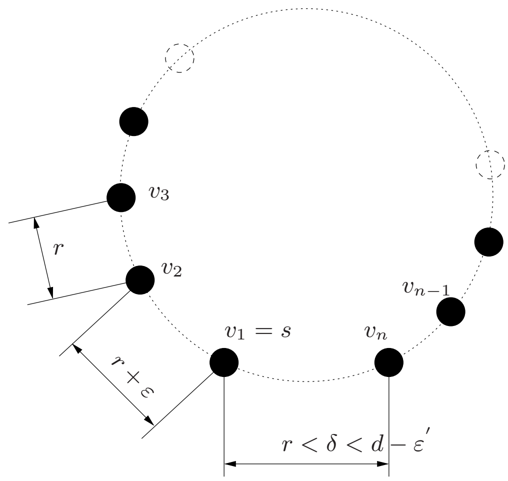

Fig. 2. Illustration of Theorem 3.

Theorem 4. *For any $v_{max} \geq 0$, $d > r$, $\alpha > 0$, and $\beta > 0$, there exists an (α, β)-connected MANET of n nodes, for which any deterministic equiperiodic protocol takes at least $n(n-1)/2$ time slots to solve Geocast .*

Proof. The following adversarial configuration and activation schedule of nodes shows the claimed lower bound. Initially, nodes are placed as in Theorem 3, except that now the distance between v_1 and v_2 is r and nodes will be static. (See Figure 2.) The adversary chooses the IDs of nodes so that the periods of transmission are assigned in increasing order. I.e., the smallest period corresponds to the source node v_1, the second smallest to v_2, and so on. Then, v_1 and v_n are activated by the adversary at the same time, and let the node source v_1 be informed immediately upon activation. Given that they are activated simultaneously, in order to solve Geocast, v_n must become covered. Regarding the remaining nodes, the adversary chooses the activation schedule so that, for each node v_i with assigned period T_i, $1 < i < n$, node v_i transmits $T_i - 1$ steps after becoming covered. In order to preserve (α, β)-connectivity, each node i may need to be activated a number of periods T_i before, eventually even before the source node. Assume that all periods are different. Given that the first period cannot be smaller than 2, and each node i, $1 < i < n$, transmits for the first time the information at least T_i steps after becoming covered, the time bound follows.

If on the other hand the protocol is such that two nodes x, y have the same period of transmission, the following (α, β)-connected network and activation schedule shows that the problem cannot be solved. The adversary places the nodes in a static two-hop topology so that, the nodes x and y are neighbors of all nodes, but the source node is placed within distance d and bigger than r from all nodes except x and y. The adversary activates x and y so their scheduled transmissions coincide in time. The rest of the nodes are activated at the same time that x or y, whichever is the latest. When the source node transmits the information for the first time, both nodes x and y, and no other node, become covered simultaneously. Given that both nodes are scheduled to transmit at the same time, their transmissions collide forever at the rest of the nodes. Thus, Geocast cannot be solved. The network is (α, β)-connected since it is statically connected. $\square$

In Theorems 3 and 4 we showed lower bounds for Geocast for arbitrarily small values of v_{max}. We now show that, by slightly constraining v_{max}, a quadratic lower bound can also be shown for arbitrary deterministic protocols.

Theorem 5. *For any $v_{max} > \pi r/(3(2\alpha + n - 4))$, $d > r$, $\alpha > 0$ and $\beta > 0$, there exists an (α, β)-connected MANET of n nodes, for which any deterministic protocol takes $\Omega((\alpha + n)n)$ time slots to solve Geocast, even if all nodes are activated simultaneously and do not fail.*

Proof. The proof is based on showing a configuration and movement of nodes such that, in order to solve the problem, all nodes have to receive the information but, each of $n/2$ nodes receive the information at least $\Omega(\alpha + n)$ steps after the previous one. In order to achieve that, ε movements among the other $n/2$ nodes, which are informed initially, are used to produce collisions while preserving (α, β)-connectivity. The following adversarial configuration and movement of nodes shows the claimed lower bound. Consider six sets of nodes A, A', B, B', C, 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, y, x', and y' placed in the configuration depicted in Figure 3(a). The invariant in these sets is that nodes in each set form a clique; each node in A' is placed at some distance $> r$ and $\leq r + \varepsilon$ from the points x, y', and each node in B; each node in A is placed at some distance $> r$ and $\leq r + \varepsilon$ from the points x', y, and each node in B; each node in B is placed within distance r of points x, y, x', and y', and each node in C and C'; each node in C is placed at some distance $> r$ and $\leq r + \varepsilon$ from the point x; each node in C' is placed at some distance $> r$ and $\leq r + \varepsilon$ from the point x'; and each node in B' is placed within distance ε of each node in B and within distance r of each node in C and C'. [4]

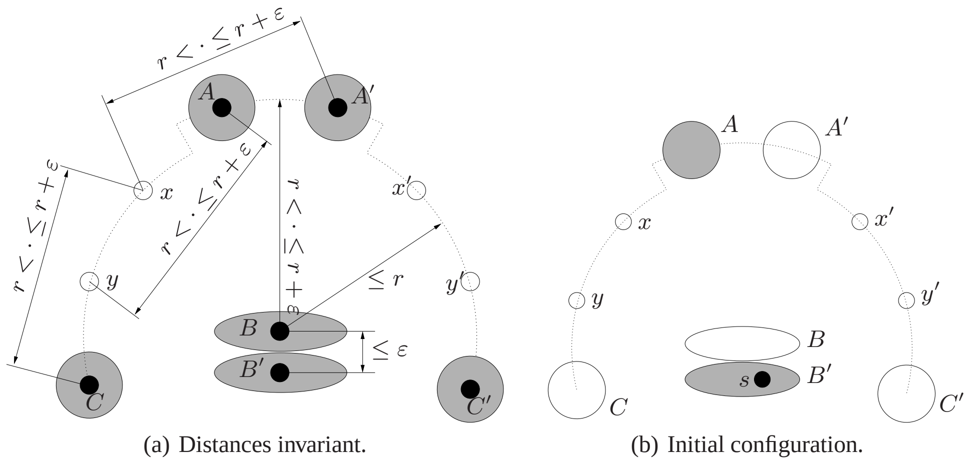

Fig. 3. Illustration of Theorem 5. A small empty circle depicts a point in the plane. A small black circle depicts a node. A big empty circle/ellipse depicts an empty set. A big shaded circle/ellipse depicts a non-empty set.

At the beginning of the first time slot, the adversary places $n/2$ nodes, including the source node s, in set B', the remaining $n/2$ nodes in the set A, and starts up all nodes. (For clarity,

[4] During some periods of time a given set could be empty. we mean that x is separated (within) that distance from any point in the area designated to the set X

assume that n is even.) All the other sets are initially empty. (See Figure 3(b).) 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 parts of the dotted line (see Figure 3(a)), of length less than $\pi(r + \varepsilon)/6$, in $\alpha + n/2 - 2$ 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 each of the circular dotted lines is upper bounded by $(r + \varepsilon)/2$.)

Let t be the first time slot when the source is scheduled to transmit. Adversarially, let t be the first time slot when the source is informed. Nodes stay in the positions described until t and, consequently, all the other $n/2-1$ nodes in set B' receive it. Starting at time slot $t+1$, the adversary moves the nodes so that only one new node every $\alpha + n/2$ steps becomes informed. First we give the intuition of the movements and later the details. (See Figure 3(b).) Nodes that are not in B or B' are moved following the doted lines. Some of the nodes in B' are moved back and forth to B. Nodes initially in A are moved clockwise to A', except for one of them, say u, which is moved simultaneously counter-clockwise to the point x. Upon reaching A' nodes are moved counter-clockwise back to A, except for one of them, say v, which is moved simultaneously clockwise to the point x', while the node u is also moved simultaneously to the point y. Upon reaching A, the remaining nodes repeat the procedure while u keeps moving towards C and v keeps moving towards C' through y' respectively. Nodes initially in A are moved in the above alternating fashion, one to C and the next one to C', until all nodes become informed. Movements are produced so that (α, β)-connectivity is preserved. The details follow.

The movement of each node u moved from A to C is carried out in three phases of at least $\alpha + n/2 - 2$ time slots each as follows. (As explained below, some nodes initially in A will be moved instead to C', but the movement is symmetric. For clarity, we only describe the whole three phases for one node.)

Phase 1 During the first $\alpha - 2$ time slots, u is moved counterclockwise from A towards the point x maintaining a distance $> r$ and $\leq r + \varepsilon$ with respect to every node in B. In the $(\alpha - 1)$-th time slot of this phase, u is moved within distance r of every node in set B preserving (α, β)-connectivity. (See Figure 4(a).) Nodes in B' stay static during this interval. Given that only nodes in B' are informed and the distance between them and u is bigger than r, u does not become covered during this interval.

During the following $n/2-1$ time slots of the first phase, the counterclockwise movement of node u towards the point x continues, but now maintaining a distance at most r with respect to every node in B. In the last time slot of the second phase, u is moved to point x. (See Figure 4(b).) During this interval, nodes in B' are moved back and forth to B as described in Theorem 1 to guarantee that u does not become covered before reaching point x. Upon reaching point x, u and all the other nodes in the network not in B or B' remain static. Phase 1 ends the time slot before u becomes covered.

Simultaneously, along the first $\alpha + n/2 - 2$ time steps of this phase, the remaining nodes initially in A are moved clockwise to A'. Then, even if u becomes informed immediately upon reaching point x, u cannot inform nodes in A' because they are separated by a distance $> r$. (See Figure 4(b).)

Phase 2 During this phase, u is moved counterclockwise towards point y maintaining a distance at most r with respect to every node in B and B'. Simultaneously, nodes that were in A' at the end of the second phase are moved counterclockwise to A except for one node v that moves in its own first phase to x'. (See Figure 4(c).)

Nodes moving from A' to A maintain a distance $> r$ with respect to u. Thus, even if u becomes covered the information cannot be passed to the former. At the end of this phase v

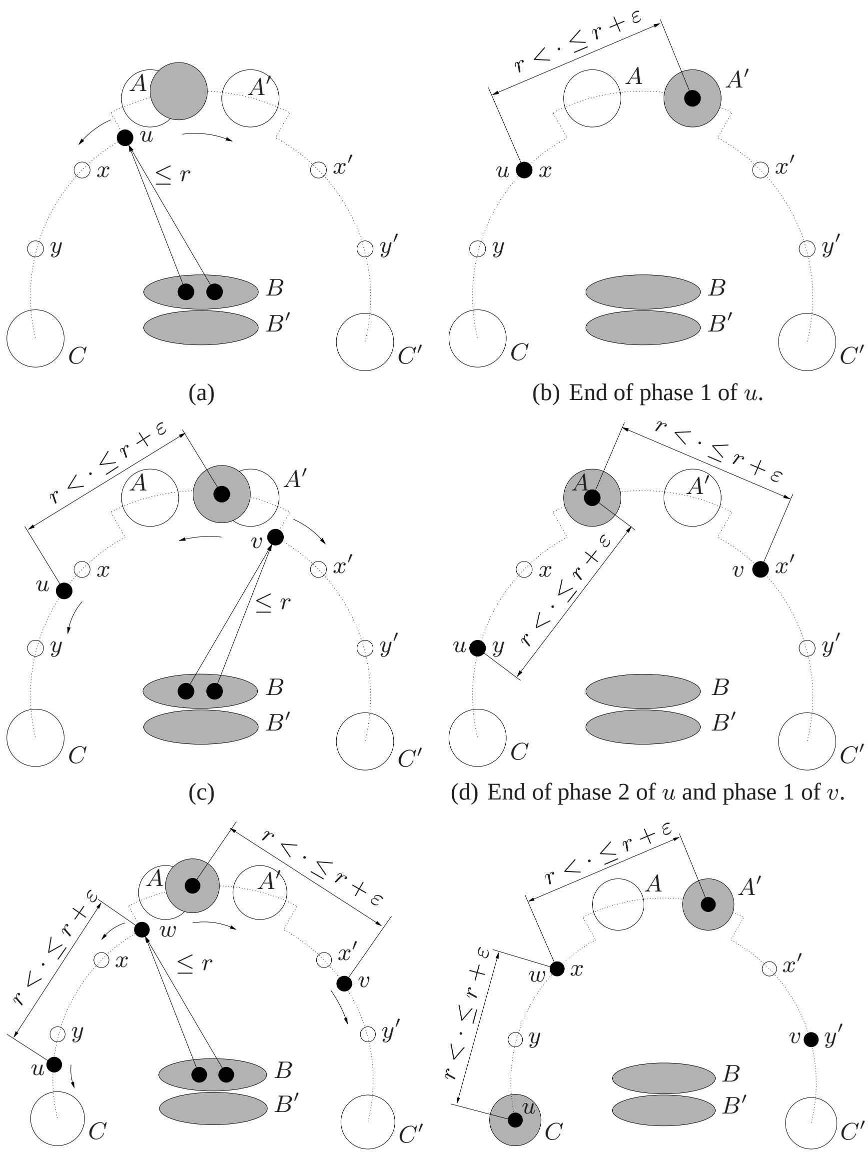

(a) (b) End of phase 1 of u.

(c) (d) End of phase 2 of u and phase 1 of v.

(e) (f) End of phase 3 of u, 2 of v, and 1 of w.

Fig. 4. Illustration of Theorem 5. A small empty circle depicts a point in the plane. A small black circle depicts a node. A big empty circle/ellipse depicts an empty set. A big shaded circle/ellipse depicts a non-empty set.

is placed in point x'. Thus, even if v becomes covered in the first step of its second phase, v cannot inform nodes in A because they are separated by a distance $> r$. (See Figure 4(d).)

Phase 3 During this phase, u is moved counterclockwise towards set C maintaining a distance at most r with respect to every node in B and B'. Simultaneously, nodes that were in A at the end of the second phase are moved clockwise to A' except for one node w that moves in its own first phase to x. Also simultaneously, v continues its movement towards set C' in its own second phase. (See Figure 4(e).)

Nodes moving from A to A' maintain a distance $> r$ with respect to v. Thus, even if v becomes covered the information cannot be passed to the former. Also, nodes u and w are moved maintaining a distance $> r$ between them. Thus, u cannot inform w. At the end of this phase u has reached set C, v is placed in point y', and w is placed in point x. Thus, even if w becomes covered in the first step of its second phase, w cannot inform nodes in A because they are separated by a distance $> r$. (See Figure 4(f).) Upon completing the third phase, u stays static in C forever so that (α, β)-connectivity is preserved.

The three-phase movement detailed above is produced for each node initially in A, overlapping the phases as described, until all nodes have became covered. Given that when a node u reaches the point x, its phase 1 is stretched until the time step before u becomes covered by a node v in B and all other nodes remain static, the next node w that will be moved from A' to x' does not become covered by v, because w stays in A' until u becomes covered. In each phase of at least $\alpha + n/2 - 2$ 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}. Given that $n/2$ nodes initially in A are covered one by one, each at least within $\alpha + n/2 - 2$ time slots after the previous one, the overall running time is lower bounded as claimed, even if $t = 1$. □

The quadratic lower bound shown in Theorem 5 holds for any deterministic protocol, even if it is adaptive. Building upon the argument used in that theorem, but additionally exploiting the adversarial node activation, the following theorem shows a roughly cubic lower bound for oblivious protocols, even relaxing the constraint on v_{max}.

Theorem 6. *For any $n \geq 9$, $d > r$, $\alpha > 0$, $\beta > 0$, and $v_{max} > \pi r/6(\alpha + \lfloor (n/3)(n/3 - 1)/\ln(n/3(n/3-1)) \rfloor - 2)$, there exists an (α, β)-connected MANET of n nodes, for which any oblivious deterministic protocol takes $\Omega((\alpha + n^2/\ln n)n)$ time slots to solve Geocast.*

Proof (Sketch). The following adversarial configuration and movement of nodes shows the claimed lower bound. Consider again the configuration described in Theorem 5 (illustrated in Figure 3(a)). The sets and distance invariant are the same as in that theorem. Regarding node location, the adversary places $2n/3$ nodes, including the source node s, in set B', and the remaining $n/3$ nodes in the set A. (For clarity, assume that n is a multiple of 3.) All the other sets are initialized empty. (See Figure 3(b).) Also, ε is again set appropriately so that a node can be moved as required without exceeding v_{max}, and all nodes in A that are active at the time that the source node transmits, must be covered for the particular value of $d > r$.

The adversarial node activation schedule follows. Let $m = \lfloor (n/3)(n/3-1)/\ln(n/3(n/3-1)) \rfloor$. Consider a partition $\{B'_o, B'_e\}$ of the nodes in B' such that $|B'_o| = |B'_e| = n/3$ and $s \in B'_o$. The adversary activates all nodes in the set A and s simultaneously. Let t_1 be the time slot when the source node is scheduled to transmit the information for the first time. Adversarially, let t_1 be the first time slot when the source is informed. Given that all nodes in A are active and within distance d of s at time t_1, all nodes in A must be covered to solve the problem.

The rest of the nodes in B'_o are activated for the first time during the m time slots before t_1 so that some node $o \in B'_o$ does not transmit uniquely among the nodes in B'_o before the time slot $t_1 + m - 1$. As proved in Lemma 2, such a node exists because the protocol is oblivious and $n \geq 9$. At the end of the time slot t_2 when node o transmits uniquely for the first time, all nodes in B'_o are deactivated.

Nodes in B'_e are activated for the first time during the m time slots before t_2 so that some node $e \in B'_e$ does not transmit uniquely among the nodes in B'_e before the time slot t_2+m-1. As proved in Lemma 2, such a node exists because the protocol is oblivious and $n \geq 9$. At the end of the time slot t_3 when node e transmits uniquely for the first time, all nodes in B'_e are deactivated.

These alternating rounds of activation are repeated until the problem is solved. In each of the rounds of activation described here, and concurrently with the movement of nodes in A in the second interval of the first phase described in Theorem 5, nodes are moved between sets B' and B as described in Theorem 5, but now using the last node to transmit uniquely in B'_o and B'_e alternately.

The movement of nodes from A and A' to C and C' is performed slower than in Theorem 5 since, for example, an uncovered node from A has now at least $\alpha + m - 2$ time steps to be moved from A to x before being informed. In each phase of at least $\alpha + m - 2$ time slots every node is moved a distance at most $\pi(r+\varepsilon)/6+\varepsilon$. Setting ε appropriately, the adversarial movement described does not violate v_{max}.

Given that $n/3$ nodes initially in A are covered one by one, each at least within $\alpha + m - 2$ time slots after the previous one, the overall running time is lower bounded as claimed.

The movement of nodes is the same as in Theorem 5, although now nodes fail. However, nodes witnessing (α, β)-connectivity do not fail while doing so. Thus, the network is (α, β)-connected. □

4 Upper Bounds

Solving the Dissemination problem under arbitrary node-activation schedule and node-movement is not a trivial task. To the best of our knowledge, deterministic protocols for such scenarios were not studied before, not even for potentially epidemic networks such as an (α, β)-connected MANET, and not even for specific instances of Dissemination. In this section, a quasi-oblivious protocol and an oblivious one that solve Dissemination, both based on known algorithms particularly suited for our setting, are described and their time efficiency proved. The first bound is asymptotically tight with respect to the more powerful class of adaptive protocols.

A Quasi-Oblivious Protocol. The idea behind the protocol is to augment the well-known Round-Robin protocol with the synchronization of the clock of each node with the time elapsed since the dissemination started, which we call the *global time*. (See Algorithm 1.) This is done by embedding a counter τ, corresponding to the global time, in the messages exchanged to disseminate the information I. Given that the schedule of transmissions of a node depends only in its ID and the global time, the protocol is quasi-oblivious.

In the following theorem, we prove that this quasi-oblivious algorithm solves Dissemination for arbitrary values of v_{max} in at most $n(\alpha + n)$ time steps. Formally,

Theorem 7. *Given an (α, β)-connected MANET where $\beta \geq n$, there exists a quasi-oblivious deterministic protocol that solves Dissemination for arbitrary values of v_{max} in at most $n(\alpha + n)$ time steps.*

Algorithm 1: Dissemination algorithm for node i. I is the information to be disseminated. The source node initially sets $\tau = 0$.

upon *receiving* $\langle I, \tau \rangle$ **do**
    for *each time slot* **do**
        increase τ by one
        if $\tau \equiv i \mod n$ **then** transmit $\langle I, \tau \rangle$
        if $\tau \geq n(\alpha + n)$ **then** stop

Proof. As a worst case, we assume that all nodes in the network must be covered, and v_{max} is arbitrarily big.

Using τ, each node achieves synchronization with respect to the time step when I was transmitted for the first time. Thus, since each node transmits in a unique time slot, collisions are avoided.

Let us denote by C the set of covered nodes, and $C(t)$ the set C at the end of time slot t. Let us denote by t_i the time slot at which the set C increases from size i to $i + 1$[5]. The time slot when the information is transmitted for the first time is then t_0.

We prove that $t_{i+1} - t_i \leq \alpha + n$ and thus in the worst case, after time $n(\alpha + n)$ all nodes are informed. At time t_i, the set of nodes $C(t_i)$ and the remaining nodes in the system $V - C(t_i)$ constitute a non trivial partition.

Given that the network is (α, β)-connected, for each time slot t, and in particular for time slot t_i, there exists a time slot $t' \geq t_i$ such that a pair of nodes $p \in C(t')$ and $q \in V - C(t')$ such that, p is informed and q is uncovered at time t', p and q are neighbors and active for all the interval $[t^{'}, t^{'} + \beta)$ and this latter intersects with the interval $[t_i, t_i + \alpha)$. In the worst case at time $t^{'} + \beta$ a new node is covered, since $\beta \geq n$.

Remember that every n time steps each informed node will re-transmit I unless the stop condition is verified. (I.e., $n(\alpha + n)$ time slots after the source node transmitted for the first time.) Since $\beta \geq n$, we know that at time $t_{i+1} \leq t^{'} + n$ node q is informed. Since $t^{'} + n \leq t_i + \alpha + n$, we have that $t_{i+1} \leq t_i + \alpha + n$. Thus, $t_{i+1} - t_i \leq \alpha + n$. □

Recall that $\beta \geq n - 1$ is required for the problem to be solvable as shown in Theorem 1. This upper bound is asymptotically tight with respect to the lower bound for general deterministic Geocast protocols when $v_{max} > \pi r/(3(2\alpha + n - 4))$ shown in Theorem 5. Thus, we can conclude that having extra information in this case does not help.

An Oblivious Protocol. We finally describe how to implement an oblivious protocol for Dissemination, based on *Primed Selection*, a protocol presented in [20] for the related problem of Recurrent Communication. Given that in this protocol the schedule of transmissions of a node depends only in its ID, the protocol is oblivious. This upper bound is only a poly-logarithmic factor away from the lower bound shown in Theorem 6.

In order to implement Primed Selection, one of n prime numbers is stored in advance in each node's memory, so that each node holds a different prime number. Let p_ℓ denote the ℓ-th prime number. We set the smallest prime number used to be p_n, which is at least n, because Primed Selection requires the smallest prime number to be at least the maximum number of neighbors of any node, which in our model is unknown. The algorithm is simple to

[5] It may happen that two nodes receive the message at the same time. But this does not affect the correctness of the proof, it just reduces the time necessary for the problem to be solved.

describe, upon receiving the information, each node with assigned prime number p_i transmits with period p_i.

It was shown in [20] that, for any Radio Network formed by a set V of nodes running Primed Selection, for any time slot t, and for any node i such that the number of nodes neighboring i is $k-1$, $1 < k < n$, i receives a transmission without collision from each of its neighbors within at most $k \max_{j \in V} p_j$ steps after t. Given that in our setting the biggest prime number used is p_{2n-1}, that $p_x < x(\ln x + \ln\ln x)$ for any $x \geq 6$ as shown in [40], and that due to mobility all nodes may get close to i in the worst case, we have that $k \max_{j \in V} p_j < n(2n-1)(\ln(2n-1) + \ln\ln(2n-1))$, for $n \geq 4$. Which is in turn less than $4n(n-1)\ln(2n)$ for $n \geq 3$. Hence, given that in the worst case all nodes must be covered at least one at a time and that the network is (α, β)-connected, the overall running time is less than $n(\alpha + 4n(n-1)\ln(2n))$. We formalize this bound in the following theorem. Recall that $\beta > \lfloor (n-1)(n-3)/4 \ln((n-1)(n-3)/4) \rfloor$ is required for the problem to be solvable when $n \geq 8$ as shown in Theorem 2.

Theorem 8. *Given an (α, β)-connected MANET, where $\beta \geq n(2n-1)(\ln(2n-1)+\ln\ln(2n-1))$ and $n \geq 4$, there exists an oblivious deterministic protocol that solves Dissemination for arbitrary values of v_{max} in at most $n(\alpha + 4n(n-1)\ln(2n))$ time steps.*

References

1. James Aspnes and Eric Ruppert. An introduction to population protocols. *Bulletin of the European Association for Theoretical Computer Science*, 93:98–117, 2007.
2. Baruch Awerbuch and Shimon Even. Efficient and reliable broadcast is achievable in an eventually connected network(extended abstract). In *Proc. 3rd Ann. ACM Symp. on Principles of Distributed Computing*, pages 278–281, 1984.
3. Baruch Awerbuch, Boaz Patt-Shamir, David Peleg, and Michael Saks. Adapting to asynchronous dynamic networks (extended abstract). In *Proc. of the 24th Ann. ACM Symp. on Theory of Computing*, pages 557–570, 1992.
4. R. Baldoni, K. Ioannidou, and A. Milani. Mobility versus the cost of geocasting in mobile ad-hoc networks. In *Proc. 21st Ann. EATCS International Symposium on Distributed Computing*, volume 4731 of *Lecture Notes in Computer Science*, pages 48–62. Springer-Verlag, Berlin, 2007.
5. 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.
6. H. Baumann, P. Crescenzi, and P. Fraigniaud. Parsimonious flooding in dynamic graphs. In *Proc. 28th Ann. ACM Symp. on Principles of Distributed Computing*, pages 260–269, 2009.
7. D. Bruschi and M. Del Pinto. Lower bounds for the broadcast problem in mobile radio networks. *Distributed Computing*, 10(3):129–135, 1997.
8. B. Chlebus, L. Gąsieniec, A. Lingas, and A. Pagourtzis. Oblivious gossiping in ad-hoc radio networks. In *Proc. of 5th Intl. Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, pages 44–51, 2001.
9. B. S. Chlebus, L. Gąsieniec, A. Gibbons, A. Pelc, and W. Rytter. Deterministic broadcasting in unknown radio networks. *Distributed Computing*, 15:27–38, 2002.
10. B. S. Chlebus, D. R. Kowalski, and T. Radzik. On many-to-many communication in packet radio networks. In *Proceedings of the 10th International Conference on Principles of Distributed Systems*, pages 258–272, 2006.
11. A.E.F. Clementi, C. Macci, F. Pasquale, A. Monti, and R. Silvestri. Flooding time in edge-markovian dynamic graphs. In *Proc. 27th Ann. ACM Symp. on Principles of Distributed Computing*, pages 213–222, 2008.

12. A.E.F. Clementi, A. Monti, and R. Silvestri. Selective families, superimposed codes, and broadcasting on unknown radio networks. In *Proc. of the 12th Ann. ACM-SIAM Symp. on Discrete Algorithms*, 2001.
13. A.E.F. Clementi, A. Monti, and R. Silvestri. Round robin is optimal for fault-tolerant broadcasting on wireless networks. *Journal of Parallel and Distributed Computing*, 64(1):89–96, 2004.
14. A.E.F. Clementi, F. Pasquale, A. Monti, and R. Silvestri. Communication in dynamic radio networks. In *Proc. 26th Ann. ACM Symp. on Principles of Distributed Computing*, pages 205–214, 2007.
15. A.E.F. Clementi, F. Pasquale, A. Monti, and R. Silvestri. Information spreading in stationary markovian evolving graphs. In *Proc. 23rd IEEE International Parallel and Distributed Processing Symposium*, 2009.
16. A. Dessmark and A. Pelc. Broadcasting in geometric radio networks. *Journal of Discrete Algorithms*, 5:187–201, 2007.
17. K. Fall. A delay-tolerant network architecture for challenged internets. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM)*, pages 27–34, 2003.
18. M. Farach-Colton and M. A. Mosteiro. Sensor network gossiping or how to break the broadcast lower bound. In *Proc. of the 18th Intl. Symp. on Algorithms and Computation*, volume 4835 of *Lecture Notes in Computer Science*, pages 232–243. Springer-Verlag, Berlin, 2007.
19. A. Fernández Anta and A. Milani. Bounds for deterministic reliable geocast in mobile ad-hoc networks. In *Proceedings of the 12th International Conference on Principles of Distributed Systems*, pages 164–183, 2008.
20. A. Fernández Anta, Miguel A. Mosteiro, and Christopher Thraves. Deterministic communication in the weak sensor model. In *Proceedings of the 11th International Conference on Principles of Distributed Systems*, pages 119–131, 2007.
21. M. J. Fischer, N. A. Lynch, and M. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, 1985.
22. R. G. Gallager, P. A. Humblet, and P. M. Spira. A distributed algorithm for minimum-weight spanning trees. *ACM Trans. Program. Lang. Syst.*, 5(1):66–77, 1983.
23. L. Gasieniec, E. Kranakis, A. Pelc, and Q. Xin. Deterministic m2m multicast in radio networks. In *Proc. of 31st Intl. Colloquium on Automata Languages and Programming*, pages 670–682, 2004.
24. L. Gąsieniec, A. Pelc, D. Peleg, Y. Emek, E. Kantor, and C. Su. Broadcasting in udg radio networks with unknown topology. In *Proc. 26th Ann. ACM Symp. on Principles of Distributed Computing*, 2007.
25. S.K.S. Gupta and P.K. Srimani. An adaptive protocol for reliable multicast in mobile multi-hop radio networks. In *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, 1999.
26. P. Humblet. Selecting a leader in a clique in o(n log n) messages. *Proc. 23rd Conf. on Decision and Control*, page 11391140, 1984.
27. X. Jinag and T. Camp. A review of geocasting protocols for a mobile ad hoc network. In *Proceedings of Grace Hopper Celebration*, 2002.
28. E. Korach, S. Moran, and S. Zaks. Optimal lower bounds for some distributed algorithms for a complete network of processors. *Theor. Comput. Sci.*, 64(1):125–132, 1989.
29. D. R. Kowalski. On selection problem in radio networks. In *Proc. 24th Ann. ACM Symp. on Principles of Distributed Computing*, pages 158–166, 2005.
30. 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.
31. 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.
32. D. S. Mitrinović. *Elementary Inequalities*. P. Noordhoff Ltd. - Groningen, 1964.
33. M. Mohsin, D. Cavin, Y. Sasson, R. Prakash, and A. Schiper. Reliable broadcast in wireless mobile ad hoc networks. In *Proceedings of the 39th Hawaii International Conference on System Sciences*, page 233, 2006.

34. Regina O'Dell and Roger Wattenhofer. Information dissemination in highly dynamic graphs. In *DIALM-POMC '05: Proceedings of the 2005 joint workshop on Foundations of mobile computing*, pages 104–110, New York, NY, USA, 2005. ACM.
35. E. Pagani and G.P. Rossi. Reliable broadcast in mobile multihop packet networks. In *Proc. of the 3rd ACM Ann. Intl. Conference on Mobile Computing and Networking*, pages 34–42, 1997.
36. D. Peleg and Tomasz Radzik. Time-efficient broadcast in radio networks. In A. M. C. A. Koster and Xavier Muñoz, editors, *Graphs and Algorithms in Communication Networks*, chapter 12, pages 315–339. Springer-Verlag, 2009.
37. 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.
38. R. Prakash, A. Schiper, M. Mohsin, D. Cavin, and Y. Sasson. A lower bound for broadcasting in mobile ad hoc networks. Technical report, Ecole Polytechnique Federale de Lausanne, 2004.
39. L. G. Roberts. Aloha packet system with and without slots and capture. *Computer Communication Review*, 5(2):28–42, 1975.
40. J. B. Rosser and L. Schoenfeld. Approximate formulas for some functions of prime numbers. *Illinois Journal of Mathematics*, 6(1):64–94, 1962.