Proving Lower Bounds for Distributed Ad Hoc Broadcast

Stefano Basagni
Department of Computer Science
The University of Texas at Dallas
E-mail: basagni@utdallas.edu

Abstract

Broadcast is the network operation that concerns disseminating a message from a given node to every other node in the network. When a network is comprised of mobile nodes, as in *ad hoc* networks where potentially all the nodes can be mobile, broadcast protocols need to be *distributed*, i.e., they should be executed at each node with the minimum possible knowledge of the dynamically changing network topology. This paper presents a general method for proving lower bounds for the problem of distributed broadcast in mobile ad hoc networks. As a consequence of the generality of the proposed method, previously proven lower bounds, both for randomized and deterministic broadcast, can be seen as specific applications of the proposed method. The effectiveness of our framework is also demonstrated by proving in a simpler and direct way the linear lower bound proven for deterministic broadcast in networks with constant diameter.

Keywords: Broadcast, Lower Bounds, Ad Hoc Networks, Distributed Computing.

1 Introduction

The term *mobile ad hoc* network is now commonly used to refer to a set of geographically dispersed nodes that can be stationary or mobile. In these networks, each node is willing to forward packets for other nodes that cannot communicate directly with each other (namely, *each* node is also a router).

One of the basic characteristics of ad hoc networks is the use of shared transmission (radio) channels, which makes selective transmission impossible: Whenever a node transmits, all its neighbors (nodes within transmission range) will receive the message, and a *collision* may occur if some transmissions overlap, preventing the correct reception of the message.

The mobility of all nodes also implies that communication protocols for ad hoc networks should be *distributed*, i.e., executed at each node with the minimum possible knowledge of the network topology. Since the topology changes in time, the updates needed to make each node aware of the current network configuration could detrimentally affect the performance of the network, especially in networks with a large number of nodes.

In this paper we are interested in investigating the *broadcast* problem for ad hoc network, i.e., the problem of disseminating a message from a given node, the *source* of the broadcast, to any other node in the network. In particular, the aim of this paper is to provide a general answer to the question of how much the mentioned two "restrictions" (possibility of collisions and dynamically changing topology) affect the design of efficient distributed broadcast protocols for ad hoc networks, especially in terms of execution time.

Previous work. The broadcast problem has been extensively studied for ad hoc networks and several broadcast protocols have been proposed. A quite thorough review of these protocols and and an exhaustive list of references can be obtained from papers such as [1] and [2]. In this paper we review only the results concerning lower bounds.

In order to create a uniform way of comparing these bounds we express them in terms of the number of *time slots*, or *rounds* required to complete the broadcast of a message m. The parameters used for expressing such a complexity measure are n, the number of nodes in the network, and D, the diameter of the network.

A lower bound for *deterministic* ad hoc broadcast has been discovered by Bruschi et al. in [3]. In that paper it is proven that, in order to complete the dissemination of m, any protocol needs $\Omega(D \log n)$ rounds. Remaining on the deterministic side, in [4] the authors show an ad hoc network with constant diameter in which any deterministic broadcast protocol needs at least a linear amount of rounds for the complete dissemination of m. The methods followed in the two papers to prove the lower bounds are quite different. In the first case, a "bad" network is produced such that any deterministic protocol needs at least $D \log n$ rounds to perform broadcast. In the second case the technique is more involved, resorting to the "conversion" of the broadcast problem to a combinatorial game for which, eventually, the lower bound is proven.

In [5], by using a technique similar to that used in [3], the authors prove that *randomized* broadcast needs $\Omega(D \log \frac{n}{D})$ rounds to be completed. When $D < n^{1-\epsilon}$, $\epsilon \in O(1)$, the obtained lower bound, combined with the results in [6] and [4], gives the optimal bounds $t \in \Theta(\log^2 n + D \log n)$ for randomized ad hoc broadcast protocols.

Our contribution. In this paper we propose a general framework for determining deterministic and randomized lower bounds for distributed broadcast.

The basic idea is to define a class of networks so that, when the nodes cannot have extensive topology knowledge (as it is the case for centralized protocols) *any* broadcast protocols needs at least a certain amount of rounds to complete the diffusion of the message.

Our approach is a generalization of the methods used in [3] and [5], in that we define a class of networks that forces any broadcast protocol to proceed from the source to the farthest nodes in a layer by layer fashion: The source transmits the message m to its immediate neighbors; some of these, in turn, transmit m to their neighbors, etc., till the nodes in the last layer have correctly received m. This allows us to reduce the problem of finding a lower bound for the broadcast problem in the entire network to that of finding a lower bound for the problem of forwarding the message from a layer to the consecutive one. The overall lower bound is then obtained by combining the results on the layer by layer forwarding.

As a result, our method achieves all the mentioned lower bounds as special cases, and permits us to prove the linear lower bound on constant diameter network very easily.

2 Preliminaries

An ad hoc network may be modeled by an undirected graph $G = (V, E)$ in which V is the set of (radio) *nodes* and there is an *edge* $\{v_i, v_j\} \in E$, $v_i, v_j \in V$, if and only if v_j is in the *hearing range* (namely, can hear the transmissions) of v_i and vice versa. In this case we say that v_i and v_j are neighbors. Due to mobility, the graph can change in time.

The set of the neighbors of a node v will be indicated by $\Gamma(v)$ and its cardinality, $\delta(v) = |\Gamma(v)|$, is called the *degree* of v. With $\Delta = \max\{\delta(v) : v \in V\}$ we indicate the maximum *degree* of the network G. The *distance* $d(v_i, v_j)$ between two nodes v_i and v_j is defined as the length of the shortest path (minimum number of hops) between v_i and v_j. The maximum distance between any pair of nodes is called the *diameter* D of the network. Given the source s of a message to be broadcasted, all the nodes v such that $d(s, v) = \ell \leq D$ are said to belong to the ℓth *layer* of the network, $0 \leq \ell \leq D$. Every node v in the network is assigned a unique identifier (ID). Here, for simplicity, the node is identified with its ID and both are denoted by v.

A *distributed broadcast protocol for ad hoc networks* Π is a protocol which is executed *at each* node in the network in the following way:

a. Time of execution is considered to be slotted and the time slots, or *rounds*, are numbered $0, 1, \ldots$ At round 0 a specific node s, called the *source*, transmits a message m.

b. In each round a node acts either as a transmitter or as a receiver. A node receives a message m in a specific round if and only if in that round it acts as a receiver and *exactly* one of its neighbors acts as a transmitter. In this case m is the same message transmitted by the neighbor (i.e., the message is received correctly).

c. The action of a node v in a specific round is determined by its initial input (namely, its own ID and possibly n the number of the network nodes and Δ, the network degree) and by the sequence of messages received by v in previous rounds.

*d.*The broadcast is *completed* at round t if all the nodes have received the message m by round t.

In each of the t rounds some of the nodes are allowed to transmit the message (if they have received it) according to a node's *local schedule* which it has been given to the nodes at network set up and that can possibly change in time. The entire broadcast is thus performed according to a *global schedule* $L_\Pi = \langle T_1, \ldots, T_t \rangle$, i.e., according to a list of *transmissions (sets)* which specifies for each round i the set of nodes which act as (potential) transmitters, $1 \leq i \leq t$.

A distributed broadcast protocol Π for radio network is *correct* for a specific class of networks C if for every network $G \in C$ there exists an integer $t \geq 0$ such that Π is completed at round t when executed in the network G. Thus, *our strategy for proving a lower bound on such protocols is that of exhibiting a particular class of networks for which an arbitrary protocol has that lower bound.* In the reminder of this section we define such network class, that we denote with C_D^n.

The network class C_D^n. All networks in C_D^n have exactly $n + 1$ nodes and diameter $D + 1$. We assume that the protocol gets both n and D as inputs. Of course, in a mobile environment, or, in general, when the topology is unstable, both the number of the nodes in the network and its diameter are unknown to the nodes. However, this assumption makes our lower bound stronger. Without loss of generality, in the remaining part of this paper, we consider $D \leq \frac{n}{4}$.

A generic network in C_D^n will be denoted as G_{CP}. Such a network has $n+1$ nodes distributed among $D+2$ layers. Layer 0 contains node s, which is the initiator of the broadcast. The next D layers (layers 1 through D) consist of $h = \lfloor \frac{n}{D} \rfloor \geq 4$ nodes, and layer $D+1$ contains all the remaining nodes. All the nodes in layer ℓ, $1 \leq \ell \leq D+1$, are connected with the nodes in $CP_{\ell-1}$, which are a subset of the nodes of layer $\ell - 1$ called the *connection nodes.* Only the node(s) in CP_ℓ, $0 \leq \ell \leq D$, are connected to *all* the nodes in layer $\ell+1$. There are no connections among nodes in the same layer. A generic network $G \in C_D^n$ is depicted in Figure 1 (the connection nodes are indicated by a $\bullet$).

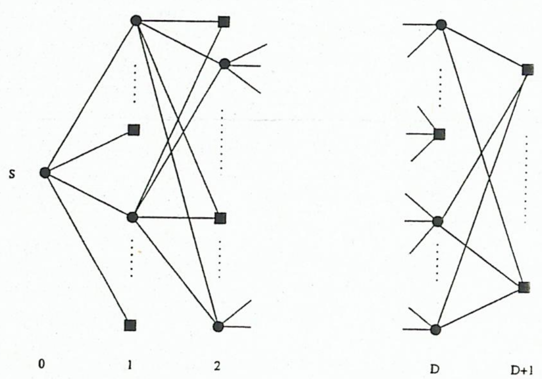

Figure 1: A generic network $G_{CP} \in C_D^n$: $n+1$ nodes are distributed into $D+2$ layers.

3 Proving Lower Bounds

In this section we present a unified framework for proving lower bounds for distributed broadcast by using the class C_D^n of ad hoc networks introduced above.

Due to the specific topology of the networks in C_D^n, and by the fact that for all layers and all the runs of any broadcast protocol, all the nodes in the same layer have the same *view* (every message received at one node of a layer is received by all the other nodes of that layer simultaneously), it can be easily verified that on any network $G_{CP} \in C_D^n$ the broadcast progresses in a layer by layer fashion. *Thus the total amount of rounds required to complete the broadcast on a network in C_D^n is given by the sum of the number of rounds required for forwarding the message from one layer to the following one.* For this reason, the first step in proving a lower bound for distributed broadcast will be that of proving a lower bound on the number of rounds needed for executing such an elementary step. If τ_ℓ is the minimum number of rounds required to correctly forward a message m from layer ℓ to layer $\ell+1$, $1 \leq \ell \leq D$, then $\Omega(\sum_{\ell=1}^{D} \tau_\ell)$ is the lower bound obtained for the broadcast of m on the entire network.

In order to prove a lower bound for τ_ℓ we have to show that *no* schedule $L_\Pi^\ell = \langle T_1^\ell, \ldots, T_\tau^\ell \rangle$ with $\tau < \tau_\ell$ transmission sets for the nodes of layer ℓ can guarantee the correct reception of m at the nodes of layer $\ell+1$. This is proven by the following proposition whose proof has a "hole," that will be filled according to the particular lower bound that has to be proven (either deterministic or randomized).

Proposition 1 *Given an arbitrary distributed broadcast protocol Π and two integers > 0, n and D, there exists a network $G_{CP} \in C_D^n$ for which it is always possible to prove the existence of a set CP_ℓ of connection nodes of layer ℓ such that Π needs at least τ_ℓ rounds for forwarding m from layer ℓ to layer $\ell+1$, $1 \leq \ell \leq D$.*

Proof We prove the lower bound assuming that when a node at layer ℓ receives m, together with m it also receives the whole history of the network and the identities of all the nodes at layers $< \ell$. This information is not necessarily available to the nodes at layer ℓ: in fact, such nodes only know the list of the previous successful transmitters and the successful rounds (in practice, only these information can be piggybacked to the message m). As noticed already, the availability of additional information makes the lower bound stronger.

In order to forward m from layer ℓ to layer $\ell+1$, the protocol Π has to generate a schedule for the nodes of layer ℓ which, independently of the $|CP_\ell|$ connection nodes of layer ℓ guarantees the correct reception of m at the nodes of layer $\ell+1$, $1 \leq \ell \leq D$. Not knowing the identities of the nodes in CP_ℓ, Π will draw them among a set of possible candidates known to Π, which we assume to be $S = V \setminus S_{\ell-1}$, where $S_{\ell-1}$ is the set containing all the nodes in G_{CP} belonging to layers $< \ell$.

Thus, to prove that at least τ_ℓ rounds are needed for the correct forwarding of m, we have to show that, regardless of the schedule $L_\Pi^\ell = \langle T_1^\ell, \ldots, T_\tau^\ell \rangle$ of length $\tau < \tau_\ell$ we consider, it is always possible to find $|CP_\ell|$ connection nodes whose transmissions, according to L_Π^ℓ, always collide, thus preventing the correct forwarding of m. In other words, given the $|S| = n - |S_{\ell-1}| = n - (\ell-1)\lfloor \frac{n}{D} \rfloor$ nodes of the network that have not transmitted m yet, we have always to be able to prove the existence of a set

CP_ℓ of "bad" connection nodes of layer ℓ, that forces the layer to layer schedule to have at least τ_ℓ transmission sets to forward m correctly. The problem of showing that such a bad set always exist can be characterized in combinatorial terms as follows: Since the schedule L^ℓ_Π is a sequence of transmitting sets $T^\ell_1, \ldots, T^\ell_\tau$, we have to show that $|CP_\ell|$ nodes in $S \subseteq \{v_1, \ldots, v_n\}$ always exist for which $|T^\ell_i \cap CP_\ell| \neq 1$, $1 \leq i \leq \tau < \tau_\ell$. This means that if $\tau < \tau_\ell$ there is no round i, $1 \leq i \leq \tau$, in which only one of the connection nodes is the only transmitter, and therefore a collision always occurs. Thus, at least τ_ℓ rounds are needed in order to forward m. We assume that the existence of such a bad set of connection nodes is guaranteed by a combinatorial lemma, the discussion on which is the main contribution of this note, and it is given below. •

We stress the fact that the construction in the previous proof can be applied to an arbitrary schedule of length $< \tau_\ell$, $1 \leq \ell \leq D$, and it is solely based on the assumption that such a schedule has been "built" ignoring the identities of the connection nodes. *It is worth noticing that our proof also holds in the case when the nodes have neighbors' knowledge.* In this case, each connection node v of a given layer could know the identities of *all* the nodes in the same layer (such identities could have been sent to it by their predecessors, piggybacked to m) but the lack of network topology knowledge, still prevents v from knowing the identities of the other connection nodes. Thus, the schedule L^ℓ_Π can be determined by Π using any kind of information which can be contained on the view of a node (except, of course, the identities of the connection nodes) without affecting the proof.

Now, we extend the result of Proposition 1 from two consecutive layers of a network G_{CP} to the entire network.

Theorem 1 *Given an arbitrary distributed broadcast protocol Π and two integers > 0, n and D, there always exists a network $G_{CP} \in C^n_D$ on which Π needs $\Omega(\sum_{\ell=1}^D \tau_\ell)$ rounds of transmission for broadcasting a message m.*

Proof Let Π be an arbitrary distributed broadcast protocol for radio networks. We build the "bad" network $G_{CP} \in C^n_D$ that forces Π to execute at least $1+\sum_{\ell=1}^D \tau_\ell$ rounds to complete the broadcast of m layer by layer in the following way. Let us consider the set $V = \{s, v_1, \ldots, v_n\}$. The only node of layer 0 will be the source of the broadcast, s. Let $S = V \setminus \{s\}$. For each layer ℓ, $1 \leq \ell \leq D$, we execute the following steps:

1. We choose from S a set CP_ℓ of connection nodes such that Π needs to execute at least τ_ℓ rounds to forward m to the nodes of the following layer correctly;
2. $S := S \setminus CP_\ell$;
3. We complete the ℓth layer by choosing $\lfloor \frac{n}{D} \rfloor - |CP_\ell|$ from S;
4. We remove from S the nodes determined in the previous step 3.

The remaining nodes (if any) form layer $D+1$.

Correctness. In step 1. the existence of the set of the bad connection nodes of layer ℓ is guaranteed by Proposition 1, $1 \leq \ell \leq D$. The second and forth steps, ensure us that the choices of the nodes determined for each layer are consistent, because once a node has been assigned to a layer it is no longer reconsidered in further iterations since it is permanently removed from the set S.

Complexity. Since the existence of a bad network is always guaranteed given the consistency of the assignments of the nodes to the layers, and since by Proposition 1 a layer to layer schedule has to have at least τ_ℓ rounds, a network G_{CP} obtained as indicated above forces the protocol Π to execute a total number of rounds t for broadcasting m which is at least $1 + \sum_{\ell=1}^D \tau_\ell$. •

As mentioned in the proof of Proposition 1, the previous result depends on a "combinatorial lemma" that guarantees the existence of the bad set of connection nodes. A general form for the needed lemma can be stated in the following way:

Lemma 1 *Given a set $S \subseteq \{v_1, \ldots, v_n\}$, $|S| \geq 2$, and an arbitrary family of subsets of S, $\mathcal{T} = \{T_1, \ldots, T_\tau\}$, $\tau < \tau_\ell$, a set $\Gamma \subseteq S$, $2 \leq |\Gamma| \leq h = \lfloor \frac{n}{D} \rfloor$, always exists such that $|T_i \cap \Gamma| \neq 1$, $1 \leq i \leq \tau$.*

In broadcast terms, the previous lemma states that regardless of the layer to layer schedule $\mathcal{T}$ with less that τ_ℓ transmission sets, it is always possible to show that a subset Γ (the connection nodes) of the set S (nodes that have to transmit m yet) always exists such that $2 \leq |\Gamma| \leq h$ ($|\Gamma|$ does not exceed the number of nodes in a layer of a network in C^n_D) and no node in Γ is ever allowed to be the only transmitter. Thus, choosing Γ as the set of the connection nodes for a given layer, m cannot be correctly forwarded to the following layer.

4 Lower Bounds for Distributed Ad Hoc Broadcast

In this section we show how any previously proven lower bound for the problem of distributed broadcast in ad hoc radio networks can be obtained following the guidelines stated in Proposition 1 and in Theorem 1 once a suitable instance of Lemma 1 is proven.

A logarithmic lower bound for deterministic broadcast

The $\Omega(D \log n)$ lower bound presented in [3] is obtained by proving the following instance of Lemma 1, where $\tau_\ell = \log(n - (\ell - 1)h)$, $h = \lfloor \frac{n}{D} \rfloor$.

Lemma 2 (Lemma 3.1,[3]) *Given a set $S \subseteq \{v_1, \ldots, v_n\}$, $|S| \geq 2$, and an arbitrary family of subsets of S, $\mathcal{T} = \{T_1, \ldots, T_\tau\}$, $\tau < \log |S|$, a set $\Gamma \subseteq S$, $|\Gamma| = 2$, always exists such that $|T_i \cap \Gamma| \neq 1$, $1 \leq i \leq \tau$.*

In the original proof of the previous lemma the two-elements set Γ is actually computed, i.e., given a family $\mathcal{T} = \{T_1, \ldots, T_\tau\}$ of $\tau \leq \log |S|$ subsets of S, a set Γ with two elements is always found that has either no intersection with any of the sets in $\mathcal{T}$ or the intersection is Γ itself. Thus, for each layer ℓ, $1 \leq \ell \leq D$, CP_ℓ is always a two-elements set that can be actually constructed. It is proven that $t \geq 1 + \sum_{\ell=1}^{D} \log(\frac{n-(\ell-1)h}{2}) \geq (D-1)\log n + k \log(D-1) - 3D \in \Omega(D \log n)$, where $h = \lfloor \frac{n}{D} \rfloor$ and k is a constant. In [1] a broadcast algorithm is presented for ad hoc networks that completes the broadcast in $O(D 2^d \log^d n)$ rounds, where d is such that $\Delta < 2^{d+1}$. Thus, for networks with maximum degree $\Delta < 4$ the time complexity of that protocol matches the $\Omega(D \log n)$ lower bound.

A logarithmic lower bound for randomized broadcast

The probabilistic lower bound presented in [5] is obtained by proving the following lemma.

Lemma 3 (Lemma 1,[5]) *Given a set $S \subseteq \{v_1, \ldots, v_n\}$, $|S| \geq 2$, and an arbitrary family of subsets of S, $\mathcal{T} = \{T_1, \ldots, T_\tau\}$, $\tau < \tau_\ell$, a set $\Gamma \subseteq S$, $|\Gamma| = 2^j$, always exists such that $|T_i \cap \Gamma| \neq 1$, $1 \leq i \leq \tau$, where $E_j[E[\tau_\ell]] \in \Omega(\log \frac{n}{D})$ (the expectation is taken over the probabilistic choice of 2^j nodes from S, and j is uniformly chosen in $[1, 2, \ldots, \lfloor \log \frac{n}{D} \rfloor]$).*

For each layer ℓ, $1 \leq \ell \leq D$, the set CP_ℓ of the connection nodes is chosen randomly (from S) in such a way that $|CP_\ell| = 2^{j_\ell}$, where j_ℓ is uniformly chosen in $[1, 2, \ldots, \lfloor \log \frac{n}{D} \rfloor]$.

The expected number t of rounds necessary to any (probabilistic) protocol Π to complete the broadcast of a message is then computed by showing that, for some choice of $j_1, \ldots, j_D$, $E_\Pi[\sum_{\ell=1}^{D} \tau_\ell] \in \Omega(D \log \frac{n}{D})$, where the expectation is taken over the random choices of Π.

As mentioned in the Introduction, when $D < n^{1-\epsilon}$, $\epsilon \in O(1)$, the obtained lower bound becomes $\Omega(D \log n)$, and this, combined with the results in [6] and [4] gives the optimal bounds $t \in \Theta(\log^2 n + D \log n)$.

A linear lower bound for deterministic broadcast in ad hoc networks with constant diameter

In [4] it is shown that it is always possible to find a network with $D \in O(1)$ in which any distributed broadcast protocol needs $\Omega(n)$ rounds to complete the broadcast. In that paper, the broadcast is reduced to a combinatorial game. The reduction is performed in three steps: in the first two steps the problem is "simplified" by restricting and strengthening broadcast protocols in such a way that the lower bound it is not affected, and in the third step, the obtained "abstract" protocols are reformulated as a combinatorial game, for which a linear lower bound is finally shown. Lots of new notations and long proofs are needed in order to prove the lower bound. Here we show how that lower bound can be easily proved by following the steps indicated above.

Let us consider a generic network $G_{CP} \in C_1^{2n}$ with just three layers (Figure 2). According to the definition of C_D^n given earlier, layer 0 only contains the source node s. Each one of the other two layers has n node (so that each network of this particular subclass of C_D^n has $2n+1$ nodes and diameter 2). Without loss of generality, we suppose that the nodes in the first layer are in the set $S = \{v_1, \ldots, v_n\}$. The source is connected to all the nodes in the first layer (all these nodes will be the receivers of the initial transmission), and only nodes from a set CP_1 (connection nodes), $|CP_1| \geq 2$, are connected to the nodes of the last layer. The broadcast problem in these networks reduces to reaching the nodes in layer 2.

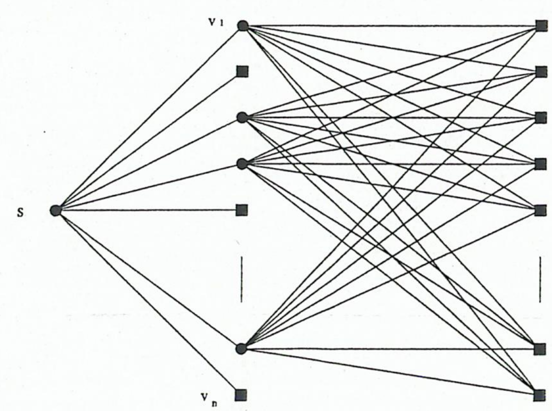

Figure 2: A generic network $\in C_1^{2n}$

Let us consider the following lemma.

Lemma 4 *Given the set $S = \{v_1, \ldots, v_n\}$, $|S| \geq 2$, and an arbitrary family of subsets of S, $\mathcal{T} = \{T_1, \ldots, T_\tau\}$, $\tau < |S| - 1$, a set $\Gamma \subseteq S$, $|\Gamma| \geq 2$, always exists such that $|T_i \cap \Gamma| \neq 1$, $1 \leq i \leq \tau$.*

Proof. Consider the set Γ returned by the following procedure.

PROCEDURE $Find_\Gamma$ (**input:** S; $T_1, \ldots, T_\tau$);
begin
$J := \emptyset$;
for $i := 1$ **to** τ **do**
  if $T_i \cap S = \{v\}$ **then**
    begin
      $S := S \setminus \{v\}$;
      for all $j \in J$ **do**
        if $T_j \cap S = \{u\}$ **then**
          begin
            $S := S \setminus \{u\}$;
            $J := J \setminus \{j\}$
          end
    end
    else $J := J \cup \{i\}$;
$\Gamma := S$;
output Γ
end;

The set Γ verifies the desired properties, namely:
a. $|\Gamma| \geq 2$. The procedure removes from the set S at most one element for any component of $\mathcal{T}$. Thus, at most $\tau \leq n-2$ elements are removed from S by the end of the procedure.
b. $|T_i \cap \Gamma| \neq 1$, $1 \leq i \leq \tau$. Let $S^{(0)}$ be the set S and let $S^{(i)}$, $1 \leq i \leq \tau$, be the set computed by the procedure $Find_\Gamma$ after the ith iteration of the main loop. It can be easily seen that $\Gamma = S^{(\tau)} \subseteq S^{(\tau-1)} \subseteq \ldots \subseteq S^{(i)} \subseteq \ldots \subseteq S^{(0)} = S$. Now, assume that there exists an l, $1 \leq l \leq \tau$, such that $|T_l \cap \Gamma| = 1$.

The following two cases have to be considered:

1. During the lth iteration of the main loop of the procedure *Find_*Γ, $T_l \cap S^{(l)}$ is a singleton $\{v\}$. In this case, v is removed from $S^{(l)}$. Thus, our assumption becomes true only if during a subsequent iteration $i > l$ of the procedure *Find_*Γ the element v is added to $S^{(i)}$. But in such a case we would have $S^{(i)} \not\subseteq S^{(l)}$, with $i > l$, thus getting a contradiction.

2. During the lth iteration of the main loop of the procedure *Find_*Γ $|T_l \cap S^{(l)}| \neq 1$. We distinguish two further cases: The case for which $|T_l \cap S^{(l)}| = \emptyset$ is similar to the previous case *1.* When $|T_l \cap S^{(l)}| > 1$, the index l is added to the set J, and our assumption becomes true if, during subsequent iterations of the procedure $Find_\Gamma$, elements belonging to T_l are continuously removed until a certain value i, $l < i \leq \tau$, such that $|T_l \cap S^{(i)}| = 1$. Such a situation, however, would be recognized by the instructions in the inner loop of the procedure and the element belonging to $T_l \cap S^{(l)}$ would be removed from S. Then, the considerations made for the case *1.* hold.

It is clear that Γ can be considered as the set of the connection nodes of layer 1, and thus, any scheduling has to have at least $\tau_\ell = n-1$ transmission sets in order to forward m to the nodes of layer 2.

5 Conclusions

This paper presented a unified framework for proving lower bounds for the problem of distributed one-to-all communication (broadcast) in mobile ad hoc networks. As a consequence of the generality of the proposed method, previously proven lower bounds, both for randomized and deterministic broadcast, are special cases of the proposed method. The effectiveness of our framework is also demonstrated by proving in a simple and direct way the linear lower bound proven in [4] for deterministic broadcast in networks with constant diameter.

References

[1] S. Basagni, D. Bruschi, and I. Chlamtac. A mobility transparent deterministic broadcast mechanism for ad hoc networks. *ACM/IEEE Transactions on Networking*, 7(6):799–807, December 1999.

[2] B. S. Chlebus, L. Gąsieniec, A. Gibbons, A. Pelc, and W. Rytter. Deterministic broadcasting in unknown radio networks. In *Proceedings of 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'00)*, San Francisco, CA, January 2000.

[3] D. Bruschi and M. Del Pinto. Lower bounds for the broadcast problem in mobile radio networks. *Distributed Computing*, 10(3):129–135, April 1997.

[4] 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(1):104–126, August 1992.

[5] 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, June 1998.

[6] N. Alon, A. Bar-Noy, N. Linial, and D. Peleg. A lower bound for radio broadcast. *Journal of Computer and System Sciences*, 43(2):290–298, October 1991.