A logarithmic lower bound for time-spread multiple-access (TSMA) protocols

Stefano Basagni [a,*] and Danilo Bruschi [b]

[a] Center for Advanced Telecommunications Systems and Services, Erik Jonsson School of Engineering and Computer Science, The University of Texas at Dallas, Richardson, TX 75083, USA

[b] Dipartimento di Scienze dell'Informazione, Università degli Studi di Milano, Italy

Time-Spread Multiple-Access (TSMA) protocols are scheduled access protocols for mobile multi-hop radio networks that guarantee deterministic access to the shared channel regardless of the possibility of radio interference. In scheduled access methods, time is considered to be slotted and time slots are cyclically organized into frames. In general, the shorter the frame, the more efficient the protocol. An $\Omega(\log \log n)$ lower bound is known on the minimum length of the frame of TSMA protocols in networks with n nodes. In this note we improve that lower bound by characterizing the multiple access to the radio channel as a combinatorial problem. The proposed characterization allows us to prove that no TSMA protocols can successfully schedule the transmissions of the nodes of a multi-hop radio network in frames with less than $\log n$ time slots.

1. Introduction

Guaranteed delivery and bounded delay are fundamental requirements for implementing real-time services and multimedia applications in any networked environment. In the case of *mobile multi-hop radio networks*, i.e., radio networks without a fixed infrastructure in which all the nodes can move, these requirements are also crucial for the implementation of reliable control channel mechanisms for managing the operation of the network itself.

From a channel access perspective, a multi-hop radio network is characterized by the fact that several of its nodes may share the same transmission channel. Thus, in these networks, selective transmission is impossible – whenever a node transmits, all its neighbors (nodes within its transmission range) will receive the message. *Collisions* of received messages may occur if transmissions overlap, preventing correct message reception.

Thus, each node of the network should be able to schedule the transmission of a message so that collisions are avoided. This poses several difficulties. For instance, each node is required to know its current neighbors, which is not always feasible in a rapidly mobile environment. Moreover, protocol operations should not be compromised by changes in the network topology due to failures or due to the mobility of the nodes.

Currently, the only approach in wireless systems that combines rapid mobility with guaranteed delivery, stability and bounded delay, while avoiding any overhead to reorganize the protocol in case of topology change, is the Time-Spread Multiple-Access (TSMA) protocol family introduced in [2] (see also [3,4] and [1]).

* Stefano Basagni was supported in part by the Army Research Office (DARPA) under contract No. DAAG55-97-1-0312.

TSMA protocols are *time division scheduled access* protocols. In these solutions, time is divided into *slots* of fixed length which are in turn organized in cycles (*frames*), each made up of L slots. In the case of TSMA protocols, each node is assigned a unique code that deterministically specifies its frame transmission *schedule* (i.e., the slots of the frame in which the node is allowed to transmit a message). Unlike Time Division Multiple-Access (TDMA) protocols, collisions in certain slots of the schedule are not excluded, but due to the special coding method, each user is *guaranteed* a successful transmission to every neighbor within each frame, no matter which nodes are the current neighbors. Since the position of the successful transmission among the attempts in the frame is not known in advance (but its existence is guaranteed), the success is "spread over time", hence the name. The transmission schedule (code) for every user is permanently computed once during the network setup and it is proven to induce a polylogarithmic frame length (further details can be found, for instance, in [2]).

Since the correct transmission of a message is guaranteed within a frame, the shorter the frame, the more efficient the protocol. Thus, the challenge is to find a slot assignment for each node that guarantees message delivery within the shortest possible frame length L.

In this note we are interested in evaluating an estimate on the efficiency of such protocols by determining a lower bound on L. For this purpose, we present a new combinatorial characterization of the slot assignment problem that allows us to prove that no TSMA protocols can guarantee the correct delivery of a message within frames of length $L < \log n$. (Throughout the note all the logarithms are intended to be base 2.)

The best lower bound known for this problem was $\Omega(\log \log n)$. This lower bound can be deduced by the results contained in [5] about a local graph coloring problem (i.e., the problem of assigning a color to a vertex of an undirected graph based on the colors of that vertex and of its neighbors).[1] Thus, the lower bound presented in this note represents an exponential improvement on the previous result.

2. Basic definitions

A multi-hop radio network is modeled by an undirected graph $G = (V, E)$ in which V, $|V| = n$, is the set of (radio) nodes and there is an edge $\{u, v\} \in E$ if and only if u and v can mutually receive each other's transmissions. In this case, we say that u and v are neighbors. It is important to stress here that, due to node mobility, the graph can change in time. Thus, in the most general case, each pair of nodes can be neighbors at a certain time. Each node v in the network is assigned a unique identifier (ID), denoted by a number in $\{1, \ldots, n\}$. For the sake of simplicity, a node is identified with its ID, and both are denoted by v.

With TSMA protocols, time is slotted. *Slots* are grouped into *frames*, of length L, and are numbered 1 through L. Given a node v, a *transmission schedule* (or *slot assignment*) for v is a set $S(v) \subseteq \{1, \ldots, L\}$ that contains the slots in which v is allowed to transmit. Given $S(v)$, for every $v \in V$, we define the *transmission set* $T_s \subseteq V$ as the set of all the nodes allowed to transmit in the sth slot, $1 \leqslant s \leqslant L$. Thus, the transmissions of the nodes in a frame proceed according to the list $T_1, \ldots, T_L$ of the transmission sets. Since for each v, $S(v)$ is computed once during the network setup, the transmission set list remains unchanged in each frame, regardless of changes in the network topology.

3. Combinatorial characterization

This section introduces the combinatorial characterization of the slot assignment problem which allows us to obtain our lower bound.

In multi-hop radio networks the overlapping of two or more transmissions results in a *collision*. Two types of collisions are defined:

- *Primary collisions.* A node transmitting in a given slot cannot receive a message in the same slot.
- *Secondary collisions.* A node cannot receive more than one message in the same slot. If more than one message is transmitted to a node in a given slot, all the transmissions are considered unsuccessful in this slot.

A transmission is said to be successful if no collision has occurred. In order to guarantee each node $v \in V$ a successful transmission of a message to each of its neighbors u at least once in every frame, there should be at least one slot $s \in S(v)$ such that the following holds:

- $s \notin S(u)$ (this takes care of the primary collisions), and
- for each neighbor $u' \neq v$ of u, it should be $s \notin S(u')$ (this does not allow secondary collisions to occur).

In combinatorial terms, this corresponds to the requirement that, for each pair of neighboring nodes v and u and for at least a slot s, $1 \leqslant s \leqslant L$,

- $\{u, v\} \cap T_s = \{v\}$, and
- for each neighbor $u' \neq v$ of u, $\{u', v\} \cap T_s = \{v\}$.

Given this characterization, in order to prove a lower bound τ on the frame length L, it is enough to show that given *any* list of transmission sets $T_1, \ldots, T_L$ (i.e., given *any* TSMA protocol), with $L < \tau$, we can *always* find a set $R \subseteq V$ with two nodes such that:

- $|R \cap T_s| \neq 1$, for each slot $s \in \{1, \ldots, L\}$.

This means that, in each slot s, the two nodes in R are either both scheduled to transmit a message ($|R \cap T_s| = 2$: in this case we have a primary collision), or are not scheduled to transmit at all ($|R \cap T_s| = 0$). Thus, by the end of the frame, none of the nodes in R has been able to exchange any information with the other node in R.[2]

4. The lower bound

In this section we show how, by using the previous characterization for the slot assignment problem, we can prove a *logarithmic* lower bound for the frame length L of TSMA protocols.

The proof is based on a "halving technique" that iteratively removes nodes from the node set V depending on the nodes of the transmission sets. At every iteration the set V is at most halved, and the set $V^{(L)}$ resulting after $L < \log n$ iterations is then proven to have at least two nodes. We show that any pair of nodes in $V^{(L)}$ are scheduled to transmit in the same slots. Formally, we have the following:

Theorem 1. Given the set of nodes V of a multi-hop radio network, $|V| = n \geqslant 2$, and the list of transmission sets $T_1, \ldots, T_L$ of an arbitrary TSMA protocol such that $L < \log n$, it is always possible to find a set $R \subseteq V$, $|R| = 2$, such that

$$|R \cap T_i| \neq 1, \quad 1 \leqslant i \leqslant L.$$

Proof. The set R can be determined by the following procedure, which takes as input the set of the nodes V and the transmission sets $T_1, \ldots, T_L$.

[1] The main contribution of [5] concerns upper and lower bounds for the mentioned coloring problem. Moreover, the characterization that they give of the coloring problem is based on a set theoretic model which is different from the combinatorial setting used in this note.

[2] This is a worst case situation: the nodes in R do not *have to* exchange information. However, in the case they are currently neighbors and are willing to transmit a message, no frame shorter than τ can guarantee them to successfully receive each other's transmissions.

PROCEDURE *Find_R* $(V, T_1, \ldots, T_L)$
begin
  for $i := 1$ to L **do**
    if $|V \cap T_i| > |V \cap \overline{T_i}|$
      then $V := V \cap T_i$
      else $V := V \cap \overline{T_i}$
  $R :=$ any two elements $\in V$
end

Let $V^{(0)}$ be set V, and let $V^{(i)}$, $1 \leqslant i \leqslant L$, be the set computed after the ith iteration of the procedure *Find_R*. We prove that each subset R, $|R| = 2$, contained in $V^{(L)}$, has the desired property. We start by showing that $|V^{(L)}| \geqslant 2$, i.e., that at least a two-element set R always exists $\subseteq V^{(L)}$. We proceed by induction, showing that $|V^{(i)}| \geqslant n/2^i$, $0 \leqslant i \leqslant L$. When $i = 0$, the property trivially holds. Suppose now that $|V^{(i)}| \geqslant n/2^i$ is satisfied for a generic i, $0 < i < L$. During step $(i+1)$, the set $V^{(i)}$ is partitioned into two sets, namely $V^{(i)} \cap T_{i+1}$ and $V^{(i)} \cap \overline{T_{i+1}}$. Let $V^{(i+1)}$ be the set of greatest cardinality. Thus,

$$\left|V^{(i+1)}\right| \geqslant \frac{|V^{(i)}|}{2},$$

and, by the induction hypothesis:

$$\left|V^{(i+1)}\right| \geqslant \frac{n/2^i}{2} = \frac{n}{2^{i+1}}.$$

Thus, after $L < \log n$ steps,

$$\left|V^{(L)}\right| \geqslant \frac{n}{2^L} \geqslant \frac{n}{2^{\log n - 1}} = 2.$$

Now we show that any $R \subseteq V^{(L)}$, $|R| = 2$, is such that $|R \cap T_i| \neq 1$, $1 \leqslant i \leqslant L$. We start by noticing that $R \subseteq V^{(L)} \subseteq V^{(L-1)} \subseteq \cdots \subseteq V^{(i)} \subseteq \cdots \subseteq V^{(0)}$. Let us now assume, for the sake of contradiction, that there exists an R and an ℓ such that $|R \cap T_\ell| = 1$. The ℓth iteration of the procedure *Find_R* sets $V^{(\ell)}$ either as $V^{(\ell-1)} \cap \overline{T_\ell}$ or as $V^{(\ell-1)} \cap T_\ell$. In the first case, when $\ell < L$, our assumption becomes true only if during a subsequent iteration $i > \ell$ of the procedure *Find_R* the element $R \cap T_\ell$ is added to $V^{(i)}$, but then we would have $V^{(i)} \not\subseteq V^{(\ell)}$, a contradiction. In the case $\ell = L$ the contradiction is straightforward, since $V^{(L)} = V^{(L-1)} \cap \overline{T_L}$ cannot contain an element from T_L. In the second case, since both the elements in R belong to $V^{(L)}$, the element $R \setminus T_\ell$, $\ell < L$, has to be added to $V^{(i)}$ during some iteration $i > \ell$ of the procedure *Find_R*. Thus, $V^{(i)} \not\subseteq V^{(\ell)}$, which is again a contradiction. In the case $\ell = L$ the contradiction is given by the fact that both the elements in $R \subseteq V^{(L)}$ have to be in T_L. □

The previous theorem guarantees that given the $L < \log n$ transmission sets of *any* TSMA protocol, we can always find two nodes that cannot exchange messages with each other correctly, i.e., such that they are scheduled to transmit in the same slots of every frame. Thus, when these two nodes become neighbors, the given TSMA protocol cannot guarantee their correct communication.

We notice that our proof has the advantage of being constructive: for each TSMA protocol with a frame length $< \log n$, the set V as modified by the procedure *Find_R* contains all the nodes that, in order to have guaranteed delivery within a frame, cannot be neighbors. While in the most general case no assumption can be made about the mobility of the nodes, this information can be useful for the design and the layout of multi-hop networks for which it is possible to know in advance which nodes will never be neighbors (e.g., networks of satellites or static multi-hop radio networks).

As a final note, we point out that the lower bound proven in this note holds for *every* channel access protocol whose transmission sets are computed in advance. In particular, even if *randomized* techniques are used to compute $T_1, \ldots, T_L$ (thus guaranteeing the correct access to the channel within a frame length L *with a given probability*), we need $L \geqslant \log n$.

References

[1] S. Boztas, A robust multi-priority topology-independent transmission schedule for packet radio networks, Information Processing Letters 55(5) (15 September 1997) 291–295.

[2] I. Chlamtac and A. Faragó, Making transmission schedule immune to topology changes in multi-hop packet radio networks, IEEE/ACM Transactions on Networking 2(1) (February 1994) 23–29.

[3] I. Chlamtac, A. Faragó and H.Y. Ahn, A topology transparent link attivation protocol for mobile CDMA radio networks, IEEE Journal on Selected Areas in Communications 12(8) (October 1994) 1426–1433.

[4] I. Chlamtac, A. Faragó and H. Zhang, Time-spread multiple-access (TSMA) protocols for multihop mobile radio networks, ACM/IEEE Transactions on Networking 5(6) (December 1997) 804–812.

[5] M. Szegedy and S. Vishwanathan, Locality based graph coloring, in: *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, San Diego, CA (16–18 May 1993) pp. 201–207.

Stefano Basagni is currently a full-time Ph.D. student in electrical engineering at the University of Texas at Dallas. He also holds a B.Sc. in computer science from the Università degli Studi di Pisa, Italy (1991) and a Ph.D. in computer science from the Università degli Studi di Milano, Italy (1998). His research concerns mainly algorithms for network systems with special emphasis on routing, multicast and broadcast protocols for wireless networks. Stefano Basagni has been awarded the Texas Public Education Grant (TPEG) for the academic years 1997/1998 and 1998/1999, and the 1998/1999 and 1999/2000 TxTEC Consortium Fellowship.
E-mail: basagni@utdallas.edu

Danilo Bruschi received a B.Sc. and a Ph.D. in computer science from the Università degli Studi di Milano, Italy, in 1985 and 1989, respectively. In 1990, he spent a postdoctoral year at the University of Wisconsin, Madison. In 1991 he joined the Computer Science Department at the Università degli Studi di Milano, where he is currently an Associate Professor. His current research interests include: distributed systems, mobile systems, computer and network security.
E-mail: bruschi@dsi.unimi.it