Available at
www.ComputerScienceWeb.com
POWERED BY SCIENCE @ DIRECT®

Theoretical Computer Science 302 (2003) 337–364

Theoretical
Computer Science

www.elsevier.com/locate/tcs

Distributed broadcast in radio networks of unknown topology☆

Andrea E.F. Clementi[a,*], Angelo Monti[b], Riccardo Silvestri[b]

[a] *Dipartimento di Matematica, Università di Roma, Via della Ricerca Scientifica, "Tor Vergata", Rome I-00133, Italy*

[b] *Dipartimento di Informatica, Università "La Sapienza" di Roma, Italy*

Received 5 July 2002; received in revised form 31 October 2002; accepted 7 November 2002
Communicated by D. Peleg

Abstract

A multi-hop synchronous radio network is said to be *unknown* if the nodes have no knowledge of the topology. A basic task in radio network is that of *broadcasting* a message (created by a fixed source node) to all nodes of the network. Typical operations in real-life radio networks is the *multi-broadcast* that consists in performing a set of r independent broadcasts. The study of broadcast operations on unknown radio network is started by the seminal paper of Bar-Yehuda et al. [J. Comput. System Sci. 45 (1992) 104] and has been the subject of several recent works.

In this paper, we study the completion and the termination time of distributed protocols for both the (single) broadcast and the multi-broadcast operations on unknown networks as functions of the number of nodes n, the maximum eccentricity D, the maximum in-degree Δ, and the congestion c of the networks. We establish new connections between these operations and some combinatorial concepts, such as *selective families*, *strongly selective families* (also known as *superimposed codes*), and *pairwise r-different families*. Such connections, combined with a set of new lower and upper bounds on the size of the above families, allow us to derive new lower bounds and new distributed protocols for the broadcast and multi-broadcast operations. In particular, our upper bounds are almost tight and strongly improve over the previous bounds for a large class of networks.

☆ Research partially supported by the European RTN Project *ARACNE* and by the Italian MURST Project "Resource Allocation in Communication Networks". The results of this paper have been presented at *ACM-SIAM* SODA'01 [13] and at *ACM* PODC'01 [14].

* Corresponding author.

E-mail addresses: clementi@mat.uniroma2.it (A.E.F. Clementi), monti@dsi.uniroma1.it (A. Monti), silver@dsi.uniroma1.it (R. Silvestri).

0304-3975/03/$ - see front matter
doi:10.1016/S0304-3975(02)00851-4

1. Introduction

1.1. Radio networks

Radio networks have been the subject of several works in recent years due to their potential applications in scenarios such as battlefields, emergency disaster relief, and in any situation in which it is very difficult (or impossible) to provide the necessary infrastructure [37,39]. As in other network models, a challenging task is to enable fast communication.

A *radio network* is a set of radio stations that are able to communicate by transmitting and receiving radio signals. A transmission range is assigned to each station s and any other station within this range can directly (i.e. by one *hop*) receive messages from s. Communication between two stations that are not within their respective ranges can be achieved by *multi-hop* transmissions. This implies that radio networks in the ideal environment can be seen as a particular class of geometric graphs. However, the presence of natural and artificial environment hurdles (such as mountains, buildings, etc.) potentially yields all possible network topologies. A useful (and sometimes unavoidable) paradigm of radio communication is the structuring of communication into synchronous *time-slots*. This paradigm is commonly adopted in the practical design of protocols and hence the use of the paradigm in theoretical analysis is well motivated [3,24,38].

A radio network can be modeled as a directed graph where an edge (u, v) exists if and only if u can communicate with v in one hop. The nodes of a radio network are processing units, each of which is able to perform local computations. It is also assumed that every node can perform *any* local computation required for deciding the next send/receive operation during the current time-slot. In every time-slot, each node can be *active* or *non-active*. When active, it can decide to be either *transmitter* or *receiver*: in the former case the node transmits a message along all of its outgoing edges while, in the latter case, it tries to recover messages from all its incoming edges. When it is not active, it does not perform operations of any kind. The fundamental feature here is that a node v can recover a message from one of its incoming edges if and only if this edge is the only one bringing in a message. If two or more neighbors of a node are transmitting at the same time-slot then a *collision* occurs. Nodes do not distinguish between the background noise and the noise due to a collision (i.e., we are in the case of absence of *collision detection* [36,3]). A radio network is said to be *unknown* when every node knows nothing about the network but its own label [9–11]. Informally speaking, unknown radio networks, with absence of collision detection, model communication networks in which the assumptions on the nodes' knowledge are minimal. An important motivation in studying unknown networks with no collision detection is that, in several applications, the network topology is unstable or dynamic and it is difficult to distinguish the presence of a collision from the background noise of the channel. Another motivation comes from its strong connection with the *fault-tolerance* issue [35,30,32,13]. Unknown networks can also be seen as "known" networks (i.e. networks in which nodes have the knowledge of the entire initial topology) with *unknown permanent faults*. A node (or an edge) is said to suffer a permanent fault if it is never active during the entire execution of the protocol [30,15].

It is easy to show that a broadcast protocol for unknown fault-free networks is also an (unknown permanent-)fault tolerant protocol for general (known) networks and vice versa. So, the results obtained in the unknown model immediately apply to the permanent-fault tolerance issue.

One of the fundamental tasks in network communication is the *broadcast* operation. It consists in transmitting a message from one source node to all the nodes.

According to the network model described above, the communication protocol operates in *time-slots*: at every time-slot, each active node decides to either transmit or receive, or turn into the non-active state. Two kinds of broadcast protocols have been considered in the literature [3,6,9,10]: *spontaneous protocols*, in which the starting time-slot is known to all the nodes and, thus, every node can transmit even if it has not received any message in previous time-slots; *non-spontaneous protocols* in which a node (which is not the source) may act as a transmitter in a time-slot only if it has received a message in some previous time-slots (while the source starts at time-slot 0). A deterministic (randomized) broadcast protocol is said to have *completed broadcasting* when all nodes, reachable from the source, have received (with high probability [1]) the source message. Notice that when this happens, the nodes not necessarily stop to run the protocol since they might not know that the operation is completed. We also say that a broadcast protocol *terminates* in time t if, after the time-slot t, all the nodes are in the non-active state (i.e. when all nodes stop to run the protocol).

The completion and termination time of *Deterministic* (*Randomized*) *Broadcast protocols*, in short DB (RB) protocols, will be analyzed as functions of the following parameters of the network: the number n of nodes, the maximum in-degree Δ, and the maximum eccentricity D over all possible source nodes. Given a source node s, the eccentricity of s is the largest distance between s and any node of the network. Observe that the maximum eccentricity equals the diameter in the case of symmetric networks.

A typical task in real-life radio networks is that of performing a set of simultaneous and independent broadcast operations: a *multi-broadcast* operation is to perform $r \geqslant 1$ broadcasts (from an arbitrary multiset of source nodes). The completion time of a *Deterministic* (*Randomized*) *multi-Broadcast* protocol, in short multi-DB (multi-RB) protocols, is defined as follows. A multi-DB (multi-RB) protocol on a radio network has *completion time* t if, (with high probability) every broadcast message is received by all the nodes reachable from the source of the message within the first t time-slots. The termination time of a multi-DB (multi-RB) protocol is defined as for DB protocols.

As for the channel bandwidth, we distinguish two models. In the *Unbounded-Bandwidth* (in short UB) model [11], a node can send/receive messages of unbounded size (so, a node can send an arbitrary large subset of the r messages in one time-slot). In the *Bounded-Bandwidth* (in short BB) model [4], every node can send messages of size at most $O(\log n + \log r)$ in one time-slot. In this model, the completion time of the protocols also depends on the *congestion* (denoted as c) which is defined as

[1] A formal definition of completion time for randomized protocols will be given later.

the maximum number of broadcast messages that a node has to receive (where the maximum is computed over all possible nodes).

1.2. Previous results

1.2.1. Broadcast

The first result on broadcasting in unknown radio networks can be found in the seminal paper [3] by Bar-Yehuda et al. They provide a non-spontaneous RB protocol having expected completion time $O((D+\log n)\log n)$. Observe that their randomized protocol does not terminate when no upper bound on n is known. In [1], a lower bound $\Omega(\log^2 n)$ is shown for RB protocols that holds even for graphs of constant eccentricity (and diameter). The best known general lower bound for RB protocols is $\Omega(D\log(n/D))$, obtained in [31]. As for non-spontaneous DB protocols, Bruschi and Del Pinto [6] obtained a lower bound $\Omega(D\log n)$ for symmetric networks of diameter D. Moreover, an equivalent lower bound for spontaneous DB protocols on directed networks has been proved by Chlebus et al. [9]. More recently, Kowalski and Pelc [29] proved an $\Omega(n^{1/4})$ lower bound for spontaneous DP protocols on symmetric network of diameter 4. The first DB protocol for unknown radio networks has been presented in [7,8]. The protocol is based on the construction of *superimposed codes* [28] (see Section 2), even though this is not explicitly stated. The protocol has $O(D\Delta^2\log^3 n/\log^2\Delta))$ completion time. In [5], another combinatorial tool, recently called *selective families*,[2] has been introduced in order to derive a DB protocol that works in $O(D\Delta(\log n)^{\log\Delta})$ completion time. Observe that this performance is worse than that of [7,8] since the protocol relies on an inefficient, explicit construction of selective families. The same combinatorial tool has been used in [9] to design an $O(n^{11/6})$ DB protocol. By means of a better use of selective families, more efficient DB protocols have been obtained in [10]. Among others, they present a DB protocol having $O(n^{3/2})$ completion time. The best presently known deterministic upper bound for general unknown networks is $O(n\log^2 n)$ and is obtained by means of a DB protocol introduced by Chrobak et al. in [11]. The upper bound is not constructive, however, Indyk [27] shows that this upper bound can be made constructive by paying a polylogarithmic factor. We also notice that a DB protocol for symmetric unknown networks is presented in [9] that has $O(n)$ completion time.

It thus turns out that all previous deterministic upper bounds on general unknown networks are either *superlinear* in n (independently of the parameters D and Δ) or they contain a Δ^2 factor.

1.2.2. Multi-broadcast

As for the UB model, Chrobak et al. [11] provide a multi-DB protocol for the *gossiping* operation (i.e., the special case of n simultaneous broadcast operations, each starting from one different node) that has $O(n^{3/2})$ completion time. Multi-broadcast in the BB model has been studied in [4], where a randomized distributed protocol is presented that performs the broadcast of r messages in $O((D+r)\log\Delta\log n)$ completion

[2] A formal definition will be given later.

Table 1

	Previous results		Our results
	Deterministic	Randomized	Deterministic
Lower bound	$\Omega(D\log n)$	$\Omega(D\log(n/D))$	$\Omega(D\Delta\log(n/\Delta))$ for $\Delta \leqslant n/D$
Upper bound	$O(n\log^2 n)$, $O\left(\frac{D\Delta^2\log^3 n}{\log^2\Delta}\right)$	$O((D+\log n)\log n)$	$O(D\Delta\log(n/\Delta)\log^{1+\alpha} n)$

time. However, this protocol *does not work on unknown networks*: it indeed assumes that every node knows its respective neighborhood and that the network is symmetric. Moreover, it requires a set-up phase in which a Breadth First Search tree is computed in $O((n+D\log n)\log\Delta)$ time-slots.

1.3. Our results

1.3.1. Broadcast

Our first contribution is the construction of a family of directed layered graphs that yields an $\Omega(n\log D)$ lower bound on the completion time of non spontaneous DB protocols on unknown networks. The lower bound given in [9,3] (i.e. $\Omega(D\log n)$) is superlinear only when D is "almost" linear, i.e., $D=\omega(n/\log n)$. Instead, our lower bound implies a superlinear number of time-slots for *any* $D=\omega(1)$ and, moreover, it implies that the $O(n\log^2 n)$ deterministic upper bound given in [11] is almost optimal when $D=\Omega(n^\alpha)$, for any constant $\alpha>0$. We emphasize that our lower bound also holds when every node knows n and Δ. A simple variant of our family of graphs allows us to get the first lower bound that also depends on Δ: we indeed provide an $\Omega(D\Delta\log(n/\Delta))$ lower bound that holds for any $\Delta\leqslant n/D$. This lower bound implies that the bound $O(n\log^2 n)$ given in [11] is almost optimal whenever $D\Delta=\Omega(n)$.

On the other hand, we provide a new broadcast technique that yields the first DB protocols having a completion-time that does not contain the factor n and contains D and Δ as linear factors. More precisely, we obtain an $O(D\Delta\log(n/\Delta)\log^{1+\alpha} n)$ upper bound, where α is any fixed real positive constant.

Our protocols are thus not efficient when $D\Delta=\omega(n\,\mathrm{poly}\log n)$. However, by comparing them with our $\Omega(D\Delta\ \log(n/\Delta))$ lower bound, we can see that these upper bounds are almost optimal when $\Delta=O(n/D)$.

Table 1 summarizes previous and our results for the broadcast operation on unknown networks.

1.3.2. Multi-broadcast

Let us first consider the BB model. By combining the trivial lower bound $\Omega(D)$ with the fact that a node cannot receive more than one message per time-slot, it is easy to derive an $\Omega(D+c)$ lower bound for the multi-broadcast operation for both randomized and deterministic protocols (observe that, in the UB model, we can only get $\Omega(D)$

since the congestion c is always 1). On the other hand, we are not aware of any lower bound of the form $\Omega(f(c)\cdot g(n))$ where both f and g are some unbounded functions. Such a kind of lower bounds is important since it implies that a "perfect pipeline" protocol (i.e., a protocol yielding an $O(SB(D,n)+c)$ upper bound, where $SB(D,n)$ is the best upper bound available for the broadcast operation) is not achievable. We provide the first lower bound of the kind defined above, even under very restrictive topology conditions. We indeed derive a family of graphs with $\Delta = 2$ that forces any multi-DB protocol to perform at least $\Omega(c+(c/\log c+D)\log n)$ time-slots to complete multi-broadcast operations. Then, we derive an $\Omega(c+(c/\log c)\log n+D\log n/D)$ lower bound to multi-RD protocols. Hence, perfect pipelining is not achievable even with the help of distributed random choices.

We observe that the above lower bounds also hold in presence of collision detection and when the nodes know n and/or Δ.

On the other hand, we combine a variant of our (single) broadcast technique with a suitable "local" scheduling (that solves the congestions arising inside every node) in order to get a multi-DB protocol for the BB model. This protocol has $O((D+c)\Delta^2\log^{2+\alpha}n)$ completion time (where α is any fixed real positive constant) and it can be converted into an efficiently constructible one having $O((D+c)\Delta^2\log^{3+\alpha}n)$ completion time.

By comparing the above upper bounds with our deterministic lower bound, we have that our multi-DB protocols turn out to be "almost" optimal (i.e., only a polylogarithmic factor away from the lower bound) when $\Delta = O(\text{poly}\log n)$. We also emphasize that, for $\Delta = O(1)$, our deterministic upper bound is almost equivalent to the $O((D+r)\log\Delta\log n)$ randomized upper bound [4] in which it is even assumed that the network is symmetric and the nodes know their respective neighborhood.

As for the UB model, since arbitrary large concatenation of the messages inside a node can be sent along the outgoing edges, we can use a simpler version of our multi-DB protocols. In this version, the local scheduling is not required and, thus, we get an $O(D\Delta^2\log^{2+\alpha}n)$ upper bound.

1.4. Organization of the paper

Section 2 provides an overview of the connections between the issue of radio broadcasting and some combinatorial concepts and results. In Section 3, the proofs of such combinatorial results are given. Section 4 describes the results on the broadcast operation. In Section 5, the results on the multi-broadcast operation are presented. Finally, Section 6 discusses the obtained results and proposes some open problems.

2. Techniques and combinatorial results: an overview

The proofs of our upper and lower bounds exploit some new combinatorial results that we believe to have a per se interest. In this section, we provide a description of such results and we outline their connection with broadcast and multi-broadcast operations.

In [9], Chlebus et al. make an explicit use of *selective families* in designing DB protocols in unknown networks. In what follows, the set $\{1,\dots,n\}$ is denoted as $[n]$.

Definition 2.1. Let n and k be any integers with $k \leqslant n$. A family $\mathscr{F}$ of subsets of $[n]$ is (n,k)-*selective* if, for every non empty subset Z of $[n]$ such that $|Z| \leqslant k$, there is a set F in $\mathscr{F}$ such that $|Z \cap F| = 1$.

Let $\mathscr{F} = \{F_1,\dots,F_m\}$ be an (n,Δ)-selective family. Then, it is easy to define DB protocol on a network of n nodes and maximum in-degree Δ (this protocol has been introduced in [9,10]).

A node u transmits at time-slot $t \leqslant m$ iff it has received the source message (i.e. it is informed) and $u \in F_t$.

During the execution of these m time-slots, thanks to the selective property of $\mathscr{F}$, for each set of $d \leqslant \Delta$ nodes, there is at least one time-slot in which only one of the nodes of the set can transmit. This guarantees that at least one non-informed node gets informed. By iterating this process n times, the broadcast is completed. Thus the completion time of such DB protocol is $n|\mathscr{F}|$.

We prove that there exist (n,k)-selective families of size $\mathrm{O}(k\log(n/k))$. We also prove that this upper bound is optimal, that is, any (n,k)-selective family has size $\Omega(k\log(n/k))$. On one hand, such small selective families are combined with a new broadcast technique in order to obtain our DB protocols. On the other hand, the lower bound on the size of selective families allows us to obtain lower bounds on the completion time of DB protocols. Very recently (after the conference versions of our paper), Indyk [22] has provided an efficient construction of (n,k)-selective families of size $\min\{n, k\,\mathrm{poly}\log n\}$.

In designing multi-DB protocols, we introduce families of sets having a stronger selective property.

Definition 2.2. Let $k \leqslant n$. A family $\mathscr{F}$ of subsets of $[n]$ is (n,k)-*strongly-selective* if for every subset Z of $[n]$ such that $|Z| \leqslant k$ and for every element $z \in Z$ there is a set F in $\mathscr{F}$ such that $Z \cap F = \{z\}$.

If we have at hand an $(n,\Delta+1)$-strongly selective family $\mathscr{F} = \{F_1,\dots,F_m\}$, then it is easy to define an multi-DB protocol for the UB model on a network of n nodes and maximum in-degree Δ. A node u transmits (all the messages it knows) at time-slot $t \leqslant m$ iff $u \in F_t$. By the strongly selective property of $\mathscr{F}$, for each set X of $d \leqslant \Delta + 1$ nodes and for each node $u \in X$, there is at least one time-slot in which only u transmits among the nodes in X. This guarantees that every message reaches at least a new node during the m time-slots. By iterating this process n times the multi-broadcast is completed in $n|\mathscr{F}|$ time-slots.

Similarly to our single DB protocol, our multibroadcast protocol exploits the existence and the construction of strongly selective families of small size. Actually, strongly selective families are a new appearance of a well-known notion. Let $\mathscr{F} = \{F_1,\dots,F_m\}$ be a family of subsets of $[n]$ and consider the matrix $M^{\mathscr{F}}$ of m rows and n columns where $M^{\mathscr{F}}_{t,i}$ is set to 1 iff $i \in F_t$. It turns out that $\mathscr{F}$ is an (n,k) strongly selective family iff the OR of any set of at most k columns of $M^{\mathscr{F}}$ covers only the columns of the set. Hence the columns of $M^{\mathscr{F}}$ form a *superimposed code* [28]. Superimposed codes are also known in combinatorics as *cover free families* [19]. Our protocol uses the (n,k) strongly selective families of size $O(\min\{n, k^2 \log n\})$ whose existence is proved in [19]. This existence proof does not provide an efficient construction. However, efficient construction of such families can be found in [28]: in this case the size is $O(\min\{n, k^2 \log^2 n\})$.

Since a lower bound on the size of strongly selective families determines a lower bound on the completion time of *our* multi-broadcast technique, we have also investigated this combinatorial aspect. In [12], Chaudhuri and Radhakrishnan obtain a lower bound $\Omega((k^2/\log k)\log n)$ for sufficiently large k such $k \leqslant n^{1/3}$. Our contribution here is the extension to every k of that bound, that is, $\Omega(\min\{n, (k^2/\log k)\log n\})$. This implies that there are no significantly smaller strongly selective families than those adopted by our protocols.

As in the case of selective families, we tried to exploit strongly selective families to obtain good lower bounds on the completion time of multi-DB protocols. But we could not obtain anything better than the lower bound for single broadcast.

In the BB model, besides the interference problems, the bound on the channel bandwidth yields further delays due to the congestion inside the nodes of the network. We exploit this aspect to obtain a non-trivial lower bound. Consider the following situation that can happen during the execution of a multi-broadcast protocol. There are n nodes each of them having r messages. They are not connected each other, and exactly two of them are the only in-neighbors of another node u. Thus, all the r messages can be received by u only if, for every pair of nodes and for each message, there is a time-slot in which exactly one of the nodes of the pair transmits that message and the other does not transmit anything. This suggests to introduce the following notion.

Definition 2.3. Two sequences $\vec{x}$ and $\vec{y}$ of equal length over the alphabet $\{0\} \cup [r]$ are r-different if for any $z \in [r]$ there is a coordinate i for which $\{x_i, y_i\} = \{z, 0\}$.

From the situation described above, we will show that a protocol which performs any multi-broadcast operation on unknown networks within t time-slots must yield a set of n pairwise r-different sequences of length not greater than t. By combining this connection with a new lower bound on the length of such sequences, i.e. $\Omega((r/\log r)\log n)$, we derive the lower bounds for the multi-broadcast operation on the BB model.

3. Combinatorial results

3.1. Selective families

We now show the existence of (n,k)-selective families of small size by a suitable application of the probabilistic method [2].

Theorem 3.1. *For any $n>2$ and $k\geqslant 2$, there exists an (n,k)-selective family of size* $O(k\log(n/k))$.

Proof. If $k\geqslant n/4$, then the family of all singletons from $[n]$ is an (n,k)-selective family of size n. In the sequel we assume $k<n/4$ and we say that a family $\mathscr{F}$ is selective for a family $\mathscr{S}$ if, for each $S\in\mathscr{S}$, there is a set $F\in\mathscr{F}$ such that $|F\cap S|=1$. Let $\mathscr{S}_i$, $1\leqslant i\leqslant\lceil\log k\rceil$, be the family of all the subsets of $[n]$ having size in the range $(2^{i-1},2^i]$. Consider now a family $\mathscr{F}_i$ of l_i sets (the value of l_i is specified later) in which each set is defined by randomly picking every element of $[n]$ independently, with probability $1/2^i$.

Fix a set $S\in\mathscr{S}_i$ and consider a set $F\in\mathscr{F}_i$; then it holds that

$$\Pr[|F\cap S|=1]=\frac{|S|}{2^i}\left(1-\frac{1}{2^i}\right)^{|S|-1} > \frac{|S|}{2^i}\left(1-\frac{1}{2^i}\right)^{2^i} \geqslant \frac{|S|}{4\cdot 2^i}\geqslant\frac{1}{8},$$

where the second inequality is due to the fact that $(1-1/t)^t\geqslant\frac{1}{4}$ for $t\geqslant 2$.

The sets in $\mathscr{F}_i$ have been constructed independently, so, from the above inequality, the probability that $\mathscr{F}_i$ does not select S is at most

$$\left(1-\frac{1}{8}\right)^{l_i}\leqslant e^{-l_i/8}.$$

Hence we have that

$$\begin{aligned}\Pr[\mathscr{F}_i \text{ is not selective for } \mathscr{S}_i] &\leqslant \sum_{S\in\mathscr{S}_i}\Pr[\mathscr{F}_i \text{ does not select } S]\\ &\leqslant \sum_{d=2^{i-1}+1}^{2^i}\binom{n}{d}e^{-l_i/8}.\end{aligned}$$

By choosing $l_i>8\ln(\binom{n}{2^i}2^i)$, we get

$$\sum_{d=2^{i-1}+1}^{2^i}\binom{n}{d}e^{-l_i/8}\leqslant\sum_{d=2^{i-1}+1}^{2^i}\frac{\binom{n}{d}}{\binom{n}{2^i}2^i}\leqslant\sum_{d=2^{i-1}+1}^{2^i}\frac{1}{2^i}\leqslant\frac{2^{i-1}}{2^i}=\frac{1}{2},$$

where the second inequality follows from $d\leqslant 2^i\leqslant n/2$. Since $\log\binom{n}{t}=O(t\log(n/t))$, it holds that $l_i\leqslant c2^i\log(n/2^i))$ for some constant $c>0$, thus there exists a family $\mathscr{F}_i$ selective for $\mathscr{S}_i$ and having size at most $c2^i\log(n/2^i)$. Finally, we consider the (n,k)-selective family

$$\mathscr{F}=\bigcup_{i=1}^{\lceil\log k\rceil}\mathscr{F}_i$$

whose size is

$$\sum_{i=1}^{\lceil\log k\rceil}c2^i\log(n/2^i)=O(k\log(n/k)). \qquad\square$$

In what follows, we provide a lower bound on the size of selective families. To this aim, we make use of the notion of *intersection free* family

Definition 3.1. Let $l \leqslant k \leqslant n$. A family $\mathscr{F}$ of k-subsets of $[n]$ is (n,k,l)-*intersection free* if $|F_1 \cap F_2| \neq l$ for every F_1 and F_2 from $\mathscr{F}$.

Roughly speaking, the intersection free property is somewhat "complementary" to the selectivity property we are using in this paper. So, even though an explicit mathematical connection between the two properties will be determined later, the reader can already imagine our interest in introducing the following result obtained by Frankl and Füredi.

Theorem 3.2 (Frankl and Füredi [22]). *Let $\mathscr{F}$ be an (n,k,l)-intersection free family where $2l+1 \geqslant k$ and $k-l$ is a prime power. Then it holds that*

$$|\mathscr{F}| \leqslant \binom{n}{l}\binom{2k-l-1}{k}\bigg/\binom{2k-l-1}{l}.$$

In particular, we first prove the following consequence of the above theorem

Corollary 3.1. *Let $\mathscr{F}$ be an $(n,k,k/2)$-intersection free family where k is a power of 2 and $k \leqslant n/64$. Then it holds that*

$$\log|\mathscr{F}| \leqslant \frac{11k}{12}\log\left(\frac{n}{k}\right).$$

Proof. By using the following inequalities involving binomial coefficients

$$\left(\frac{a}{b}\right)^b \leqslant \binom{a}{b} \leqslant \left(\frac{ea}{b}\right)^b, \quad \binom{a-1}{b} = \frac{a-b}{a}\binom{a}{b}$$

we obtain

$$\begin{aligned}
\log|\mathscr{F}| &\leqslant \log\left(\binom{n}{k/2}\frac{\binom{3k/2-1}{k}}{\binom{3k/2-1}{k/2}}\right) = \log\left(\frac{1}{2}\binom{n}{k/2}\frac{\binom{3k/2}{k}}{\binom{3k/2}{k/2}}\right)\\
&\leqslant \log\left(\frac{1}{2}\left(\frac{2en}{k}\right)^{k/2}\left(\frac{3e}{2}\right)^k 3^{-k/2}\right)\\
&= \frac{k}{2}\log\frac{n}{k} + \frac{k}{2}\log 3 + \frac{3k}{2}\log e - \frac{k}{2} - 1 < \frac{k}{2}\log\frac{n}{k} + \frac{5}{2}k\\
&\leqslant \frac{11k}{12}\log\frac{n}{k}. \qquad \square
\end{aligned}$$

We are now ready to prove the lower bound.

Theorem 3.3. *For any* $n>2$, *let* $\mathscr{F}$ *be an* (n,k)*-selective family with* $2\leqslant k\leqslant n/64$. *Then it holds that*

$$|\mathscr{F}| \geqslant \frac{k}{24}\log\frac{n}{k}.$$

Proof. Let k', $k/2<k'\leqslant k$, be a power of 2. Let $\chi(G)$ be the chromatic number of the graph G whose vertices are all the k'-subsets of $[n]$ and whose edges connect vertices having exactly $k'/2$ common elements. The theorem is an immediate consequence of the following inequalities:

$$\log\chi(G) \geqslant \frac{k}{24}\log\frac{n}{k}, \tag{1}$$

$$|\mathscr{F}| \geqslant \log\chi(G). \tag{2}$$

We first prove Eq. (1). For any graph $G(V,E)$ with stability number $\alpha(G)$ it holds that

$$\chi(G) \geqslant \frac{|V|}{\alpha(G)}. \tag{3}$$

Clearly a stable set of vertices in G forms an $(n,k',k'/2)$-intersection free family satisfying the conditions of Corollary 3.1. Hence, from Eq. (3) and Corollary 3.1, we have that

$$\begin{aligned}\log\chi(G) &\geqslant \log|V| - \log\alpha(G) \geqslant \log\binom{n}{k'} - \frac{11k'}{12}\log\frac{n}{k'}\\ &\geqslant k'\log\frac{n}{k'} - \frac{11k'}{12}\log\frac{n}{k'} = \frac{k'}{12}\log\frac{n}{k'} \geqslant \frac{k}{24}\log\frac{n}{k}.\end{aligned}$$

We now prove Eq. (2). Here we use the straightforward inequality

$$\chi\left(\bigcup_{i=1}^{t} G_i\right) \leqslant \prod_{i=1}^{t}\chi(G_i) \tag{4}$$

that holds for any set of graphs having the same set of vertices.

Let be $\mathscr{F}=\{F_1,\ldots,F_{|\mathscr{F}|}\}$. We define the graph G_i, $1\leqslant i\leqslant|\mathscr{F}|$, by setting $V(G_i)=V(G)$ and by drawing an edge between two vertices of G_i if they are adjacent in G and furthermore $|F_i\cap X|=1$, where X is the symmetric difference of the sets corresponding to the two vertices. Since $\mathscr{F}$ is a (n,k)-selective family and the symmetric difference of these sets has cardinality k', for any edge of G, there will be at least a graph G_i having this edge. Hence we have $G=\bigcup_{i=1}^{|\mathscr{F}|}G_i$. It thus follows that

$$\log\chi(G) = \log\chi\left(\bigcup_{i=1}^{|\mathscr{F}|}G_i\right) \leqslant \log\prod_{i=1}^{|\mathscr{F}|}\chi(G_i) = \sum_{i=1}^{|\mathscr{F}|}\log\chi(G_i) \leqslant |\mathscr{F}|,$$

where the first inequality follows from Eq. (4) and the last inequality follows by noting that the graphs G_i are bipartite graphs (i.e. $\chi(G_i)\leqslant 2$): indeed, for any two adjacent vertices in G_i one has odd intersection with the elements of F_i and the other has even intersection. □

3.2. Strongly selective families

In [8], Dyachkov and Rykov proved a lower bound $\Omega(c_k \log n)$ on the size of (n,k)-strongly selective families, where

$$\frac{c_k}{k^2/\log k} = \Theta(k^2/\log k) \quad \text{for } k \to \infty.$$

Observe that this *does not* imply the standard two-variable lower bound $\Omega((k^2/\log k) \log n)$: for instance when $k = \Omega(n)$, this would imply a lower bound $\Omega(n^2)$. The latter is clearly false. Indeed, the family consisting of all the singletons from $[n]$ is (n,k)-strongly selective for any $k = 1,\dots,n$, and it has size n. In [12], Chaudhuri and Radhakrishnan obtain a lower bound $(k^2 \log n)/(100 \log k)$ for sufficiently large k such that [3] $k \leqslant n^{1/3}$. Our contribution here is the generalization (and an improvement) of the Chaudhuri and Radhakrishnan's result.

We prove a lower bound that is only an $O(\log k)$ factor away from the $O(\min\{n, k^2 \log n\})$ bound in [19].

Theorem 3.4. *Let $\mathscr{F}$ be an (n,k)-strongly selective family.*
(i) *If $3 \leqslant k \leqslant \sqrt{2n} - 1$ then it holds that $|\mathscr{F}| \geqslant \frac{k^2}{48 \log k} \log n$.*
(ii) *If $k \geqslant \sqrt{2n}$ then it holds that $|\mathscr{F}| \geqslant n$.*

Proof. (i). The proof relies on a result by Füredi [23] and a result by Bassalygo [18] on superimposed codes. For the sake of convenience, we state such results in terms of strongly selective families.

Let $\mathscr{F}$ be an (n,k)-strongly selective family then Bassalygo proved that

$$|\mathscr{F}| \geqslant \min\left\{ \binom{k+1}{2}, n \right\}, \tag{5}$$

and Füredi proved that, for $k \geqslant 3$,

$$n \leqslant k - 1 + \binom{|\mathscr{F}|}{\left\lceil \frac{|\mathscr{F}| - k + 1}{\binom{k}{2}} \right\rceil}. \tag{6}$$

Let $3 \leqslant k \leqslant \sqrt{2n} - 1$. From Eq. (6) and the inequality $\binom{a}{b} \leqslant (ea/b)^b$, we get

$$\frac{k(k-1)}{\log \frac{e|\mathscr{F}|k(k-1)}{2(|\mathscr{F}|-k+1)}} \log(n-k+1) \leqslant 2|\mathscr{F}| - 3k + 2 + k^2.$$

[3] Notice that the conditions on k are stated immediately before Lemma 5.1 of [12].

Since $k^2/2 \leqslant k(k-1) < k^2$ and $\sqrt{n} < n - k + 1$, it follows that

$$\frac{k^2}{4\log\frac{e|\mathscr{F}|k^2}{2(|\mathscr{F}|-k+1)}}\log n \leqslant 2|\mathscr{F}| - 3k + 2 + k^2.$$

Moreover, since $k \leqslant \sqrt{2n} - 1$, Eq. (5) implies that

$$|\mathscr{F}| \geqslant \frac{k^2+k}{2}$$

and then

$$\frac{|\mathscr{F}|}{2(|\mathscr{F}|-k+1)} \leqslant 1 \quad \text{and} \quad -3k+2+k^2 \leqslant 2|\mathscr{F}|.$$

We thus obtain

$$\frac{k^2}{4\log(ek^2)}\log n \leqslant 4|\mathscr{F}|.$$

Finally, since $k \geqslant e$, $|\mathscr{F}| \geqslant (k^2/48\log k)\log n$.

(ii). When $k \geqslant \sqrt{2n}$, the thesis follows immediately from Eq. (5). □

3.3. Sets of pairwise r-different sequences

Our goal in this section is to prove a lower bound on the length of n sequences which are pairwise r-different. In the sequel, we will make use of the binary entropy function $h(t) = -t\log t - (1-t)\log(1-t)$.

The proof of our lower bound relies on the following theorem proved in [20].

Theorem 3.5 (Fachini and Körner [20]). *Let S be a subset of $\binom{[r]}{2}$ and C be a set of sequences of length m over the alphabet $[r]$ with the property that for each $\{x,y\} \in \binom{C}{2}$ and $\{a,b\} \in S$ there exists an $i \in [m]$ such that $\{x_i, y_i\} = \{a,b\}$. Then it holds that*

$$\log|C| \leqslant m \max_P \min_{\{a,b\}\in S}\left\{(p_a+p_b)h\left(\frac{p_b}{p_a+p_b}\right)\right\},$$

where, in the maximum, P is running over all the probability distributions on $[r]$.

Theorem 3.6. *Let $M(n,r)$ denote the minimum length of n sequences which are pairwise r-different. Then*

$$M(n,r) = \Omega\left(\frac{r}{\log r}\log n\right).$$

Proof. Let C be a set of n sequences which are pairwise r-different and define $S = \{\{0,i\} \mid i \in [r]\}$. From Theorem 3.5 we have that

$$\log n \leqslant M(n,r)\max_P \min_{i\in[r]}\left\{(p_0+p_i)h\left(\frac{p_i}{p_0+p_i}\right)\right\}.$$

Let $f(x)=(p_0+x)h(x/(p_0+x))$, $0\leqslant x\leqslant 1$. We have that $f'(x)=\log(p_0+x)/x\geqslant 0$ for any $0\leqslant x\leqslant 1$. Thus f is not decreasing and we can restrict the search of the maximum value, in the right hand of the above inequality, to those probability distributions in which p_i have the same value for all $i\in[r]$, i.e., $p_i=(1-p_0)/r$.

We thus consider for any $x\in[0,1]$, the probability distribution $p_i=x/r$ for $i\in[r]$ and $p_0=1-x$. Then the inequality can be written as

$$\log n \leqslant M(n,r)\max_{x\in[0,1]}\left\{\left(1-x+\frac{x}{r}\right)h\left(\frac{x/r}{1-x+x/r}\right)\right\}.$$

In order to prove the theorem, we show that

$$\max_{x\in[0,1]}\left\{\left(1-x+\frac{x}{r}\right)h\left(\frac{x/r}{1-x+x/r}\right)\right\}=\mathrm{O}\left(\frac{\log r}{r}\right).$$

Indeed, the function

$$f(x)=\left(1-x+\frac{x}{r}\right)h\left(\frac{x/r}{1-x+x/r}\right)$$

can be written as

$$f(x)=(1-x)\log\left(1+\frac{x}{r(1-x)}\right)+\frac{x}{r}\log\frac{r-rx+x}{x}.$$

Then, by using the well-known inequality $1+t\leqslant \mathrm{e}^t$ (that holds for any real t), we get

$$\begin{aligned}(1-x)\log\left(1+\frac{x}{r(1-x)}\right)&\leqslant(1-x)\log \mathrm{e}^{x/(r(1-x))}\leqslant\frac{x}{r}\log \mathrm{e}\\&\leqslant\frac{1}{r}\log \mathrm{e}=\mathrm{O}\left(\frac{1}{r}\right).\end{aligned}$$

It thus suffices to prove that

$$g(x)=\frac{x}{r}\log\frac{r-rx+x}{x}=\mathrm{O}\left(\frac{\log r}{r}\right).$$

Since $(r-rx+x)/x$ is a decreasing function in the interval set $[r/(3r-1),1]$ then, for $x\in[r/(3r-1),1]$,

$$g(x)\leqslant\frac{x}{r}\log 2r\leqslant\frac{1}{r}\log 2r\in\mathrm{O}\left(\frac{\log r}{r}\right).$$

Furthermore,

$$g'(x)=\frac{1}{r}\log\frac{r-rx+x}{x}-\frac{1}{(r-rx+x)\ln 2}$$

is strictly positive in the interval set $(0, r/(3r-1)]$. Thus, in $(0, r/(3r-1)]$, the function $g(x)$ is increasing. Hence, for $x \in [0, r/(3r-1)]$,

$$g(x) \leqslant g\left(\frac{r}{3r-1}\right) = \frac{1}{3r-1}\log 2r = \mathrm{O}\left(\frac{\log r}{r}\right).$$

This completes the proof. □

4. Broadcast operations

4.1. The lower bounds

In this section, we show the existence of an infinite family of directed graphs that force any DB protocol to perform, in the worst-case, $\Omega(n \log D)$ time-slots before completing a broadcast. Then, we provide a simple variant of this family of graphs yielding a lower bound that also depends on Δ. Our lower bound holds for the UB model (and, thus, for the BB model too). We first formalize the notion of DB protocol according to Bar-Yehuda et al. [3].

Definition 4.1. A Deterministic distributed Broadcast DB protocol P is a protocol that works in time-slots (numbered $0, 1, \ldots$) according to the following rules:

1. In the initial time-slot a specified node (i.e. the source) transmits a message (called the source message).
2. In each time-slot, each node either acts as transmitter or as receiver or is non-active.
3. A node receives a message in a time-slot if and only if it acts as receiver and exactly one of its in-neighbors acts as transmitter in that time-slot.
4. The action of a node in a specific time-slot is a function of its own label, the number of the current time-slot t, and the messages received during the previous time-slots.

Theorem 4.1. *For any DB protocol P, for any n and for any $D \leqslant n/6$, there exists an n-node directed graph G^P of maximum eccentricity D such that P completes broadcasting on G^P in $\Omega(n \log D)$ time-slots. The lower bound holds even when every node knows n.*

Proof. The graph G^P is a layered n-node graph with $D+1$ levels $L_0, L_1, \ldots, L_D$; Level L_0 contains only the source s, level L_j has no more than $\lfloor n/(2D) \rfloor$ nodes for $j = 1, \ldots, D-1$ and, finally, the level L_D contains all the remaining nodes. All nodes of L_{j-1} have outgoing edges to all nodes in L_j. As we will see later, the actions specified by P determine the node assignment in the levels $j \geqslant 1$ in such a way that the protocol is forced to execute $\Omega((n/D)\log D)$ time-slots in order to successfully transmit the source message between two consecutive levels. This assignment will be performed by induction on the levels.

From Theorem 3.3, there exists a constant $c > 0$ such that, if $2 \leqslant D \leqslant n/6$, any $(\lceil n/2 \rceil, \lfloor n/(2D) \rfloor)$-selective family must have size at least T, where $T = \lfloor c(n/D)\log D \rfloor$.

The theorem is then an easy consequence of the following:

Claim 1. *For any $j=0,\dots,D-1$, it is possible to assign nodes in $L_0, L_1, \dots, L_j$ in such a way that P does not broadcast the source message to level L_j before the time-slot $j \cdot T$.*

Proof. The proof is by induction on j. For $j=0$, the claim is trivial. We thus assume the thesis be true for any j and we prove it for $j+1$. Let us define

$$R = \{\text{nodes not already assigned to levels } L_0, \dots, L_j\}.$$

Notice that $|R| \geqslant \lceil n/2 \rceil$. In fact

$$|R| = n - \sum_{h=0}^{j} |L_h| \geqslant n - \left(\left\lfloor \frac{n}{2D} \right\rfloor \right)(D-2) - 1 \geqslant \left\lceil \frac{n}{2} \right\rceil.$$

Let L be an arbitrary subset of R. Consider the following two cases: (i) L_{j+1} is chosen as L, and (ii) L_{j+1} is chosen as R (i.e. all the remaining nodes are assigned to L_{j+1}). In both cases, the predecessor[4] subgraph G_u^P of any node $u \in L$ is that induced by $L_0 \cup L_1 \cup \cdots \cup L_{j+1} \cup \{u\}$ in G^P. It follows that the behavior of node u, according to protocol P, is the same in both cases. We can thus consider the behavior of P when $L_{j+1} = R$. Then, we define

$$F_t = \{u \in R | u \text{ acts as transmitter at time-slot } j \cdot T + t\}$$

and the family $\mathscr{F} = \{F_1, \dots, F_{T-1}\}$ of subsets from R. Since $|\mathscr{F}| < T$, $\mathscr{F}$ is not $(\lceil n/2 \rceil, \lfloor n/(2D) \rfloor)$-selective; so, a subset $L \subset R$ exists such that $|L| \leqslant \lfloor n/(2D) \rfloor$ and L is not selected by $\mathscr{F}$ (and thus by P) in any time-slot t such that $jT + 1 \leqslant t \leqslant (j+1)T - 1$. The proof is completed by choosing L_{j+1} as L. $\square$

Theorem 4.2. *Let P be a DB protocol. Then, for any n, for any $D \leqslant n/6$, and for any $\Delta \leqslant n/D$, there exists an n-node directed graph G^P of maximum eccentricity D and in-degree bounded by Δ such that P completes broadcasting on G^P in $\Omega(D\Delta \log(n/\Delta))$ time-slots. The lower bound holds even when every node knows n and Δ.*

Proof. The proof is based on the same construction of the proof of Theorem 4.1. The only difference is that, for every $j = 1, 2, \dots, D-1$, level L_j of G^P consists of at most Δ nodes and L_D (consisting of all the remaining nodes) is connected to the previous level in such a way that the maximum in-degree is kept not larger than Δ. $\square$

4.2. The upper bounds

This section provides a DB protocol for unknown networks. For case of exposition, we first describe the protocol that assume the knowledge of n and Δ. Then, we show

[4] Given a graph G, the predecessor subgraph G_u of a node u is the subgraph of G induced by all nodes v for which there exists a directed path from v to u.

how to extend the same technique to the cases in which Δ and n are not known by the nodes.

The following protocol assumes the knowledge of Δ and n.

4.2.1. Description of protocol BROAD-A(n, Δ)

The protocol uses an (n, Δ)-selective family $\mathscr{F}$. It starts by setting all the nodes to the active state, and by let s transmit the source message. After the first time-slot, it turns s to the non-active state. Then, it performs a sequence of consecutive identical *phases*. Let us fix an arbitrary ordering for the sets of $\mathscr{F}$; at time-slot j of *phase* i, each node v acts according to the following rule: v transmits the source message along its outgoing edges if and only if

(1) the label of v belongs to the jth set of $\mathscr{F}$, and
(2) v has received the source message *for the first time* during the phase $i-1$.

After the phase in which a node v acts as a transmitter, it turns to the non-active state (so, this is a first change w.r.t. the straightforward protocol described in Section 2). The active nodes that, at time-slot j of any phase, have a state not satisfying Conditions (1) and (2) act as receivers. Observe that Condition (2) is the *key* difference between our technique and the straightforward one. As we will see in the analysis of the protocol, this difference will play a crucial role in order to achieve an upper bound not containing the linear factor n.

Theorem 4.3. *Protocol* BROAD-A(n, Δ) *completes broadcasting and* terminates *in* $O(D\Delta \log(n/\Delta))$ *time-slots on any n-node graph of maximum eccentricity D and maximum in-degree* Δ.

Proof. Since, Theorem 3.1 implies that $|\mathscr{F}| = O(D\Delta \log(n/\Delta))$, the thesis is an easy consequence of the following claim.

Claim. *A node v receives (for the first time) the source message at phase i of protocol* BROAD-A(n, Δ) *if and only if v is at distance* $i+1$ *from the source s.*

Proof. The proof is by induction on i. ($\Leftarrow$). For $i=0$ the claim is obvious. We thus assume that all nodes at distance i have received the source message during phase $i-1$. Let us consider a node v at distance $i+1$ during phase i. This node has at least one informed in-neighbor at distance i. According to the protocol, only the neighbors of v informed in phase $i-1$ will act as transmitters in phase i. Then, from the (n, Δ)-selectivity of $\mathscr{F}$, there will be a step of phase i, in which only one of these informed in-neighbors will transmit to v. ($\Rightarrow$). If v is not at distance $i+1$ from s, then two cases may arise. If v is at distance less than $i+1$ then, by the inductive hypothesis, v has been informed before phase i. Otherwise, v is at distance greater than $i+1$, so none of its in-neighbors has been informed before phase i. $\square$

The next protocol assumes the knowledge of n.

4.2.2. Description of protocol BROAD-B(n)

Each node runs a sequence of *phases*, each of them consisting of $\lceil \log n \rceil$ time-slots. In time-slot l ($1 \leqslant l \leqslant \lceil \log n \rceil$) of phase h, each node runs time-slot h of BROAD-A($n, 2^l$). Furthermore, if a node v is set to the non-active state in a time-slot of BROAD-A($n, 2^l$) for some $1 \leqslant l \leqslant \lceil \log n \rceil$, then it will stay inactive for all the rest of BROAD-B(n).

Theorem 4.4. *Protocol* BROAD-B(n) *completes broadcasting and terminates in* $O(D\Delta \log n \log(n/\Delta))$ *time-slots on any n-node graph of maximum eccentricity D and maximum in-degree* Δ.

Proof. Since G has maximum in-degree d, the execution of BROAD-A($n, 2^{l_d}$) where l_d is the minimum integer such that $d \leqslant 2^{l_d}$ satisfies claim in the proof of Theorem 4.3. Observe that during this execution, a node v, that satisfies the two conditions for transmitting, could be already in the non-active state because of the execution of some BROAD-A($n, 2^l$) with $l \leqslant l_d$. However, if this is the case, v has already successfully transmitted the source message to all its out-neighbors.

In the third protocol, the nodes only know their respective labels. □

4.2.3. Description of protocol BROAD$^\alpha$($\alpha > 1$)

Informally speaking, this protocol consists in running BROAD-B(n) with $n = 2^\ell$, for $\ell = 1, 2, \ldots$. One of these executions will be the "good" one. However, applying a direct "dovetail" scheduling would result into a completion time of $O((D\Delta \log n \log(n/\Delta))^2)$ (recall that nodes do not know n). So, in order to bound the extra-time by a factor of $O(\log^\alpha n)$, Protocol BROAD$^\alpha$ executes different applications of BROAD-B($\cdot$) according to a more sophisticated *dovetail* technique. Consider the following family of functions:

$$f_0^\alpha(z) = 0, \quad f_k^\alpha(z) = 2^{\lceil k^{2/\alpha} \rceil}(k - z), \quad k = 1, 2, 3, \ldots \ .$$

Protocol BROAD$^{(\alpha)}$ consists of a sequence of *phases*, denoted as PHASE(k), $k = 1, 2, 3, \ldots$. The PHASE(k) is in turn formed by k *stages*: in STAGE(k, ℓ) (with $\ell = 0, 1, \ldots, k-1$), the nodes execute the time-slots

$$f_{k-1}^\alpha(\ell) + 1, \ f_{k-1}^\alpha(\ell) + 2, \ldots, f_k^\alpha(\ell)$$

of BROAD-B(2^ℓ). If a node v is not active in a time-slot of BROAD-B(2^ℓ), for some ℓ, then it will remain non-active for all the rest of BROAD$^{(\alpha)}$. This new dovetail technique is shown in Fig. 1. Observe that a node v during the execution of a time-slot of BROAD-B(2^ℓ) could have been informed for the first time during a time-slot of the execution of BROAD-B($2^{\ell'}$) for some $\ell' \neq \ell'$. In this case, by definition of Protocol BROAD-A($\cdot, \cdot$), the node v acts as an informed node.

Theorem 4.5. *For any positive constant* $\alpha > 0$, BROAD$^{(\alpha)}$ *completes broadcasting and terminates in* $O(D\Delta \log(n/\Delta) \log^{1+\alpha} n)$ *time-slots on any n-node graph of maximum eccentricity D and maximum in-degree* Δ.

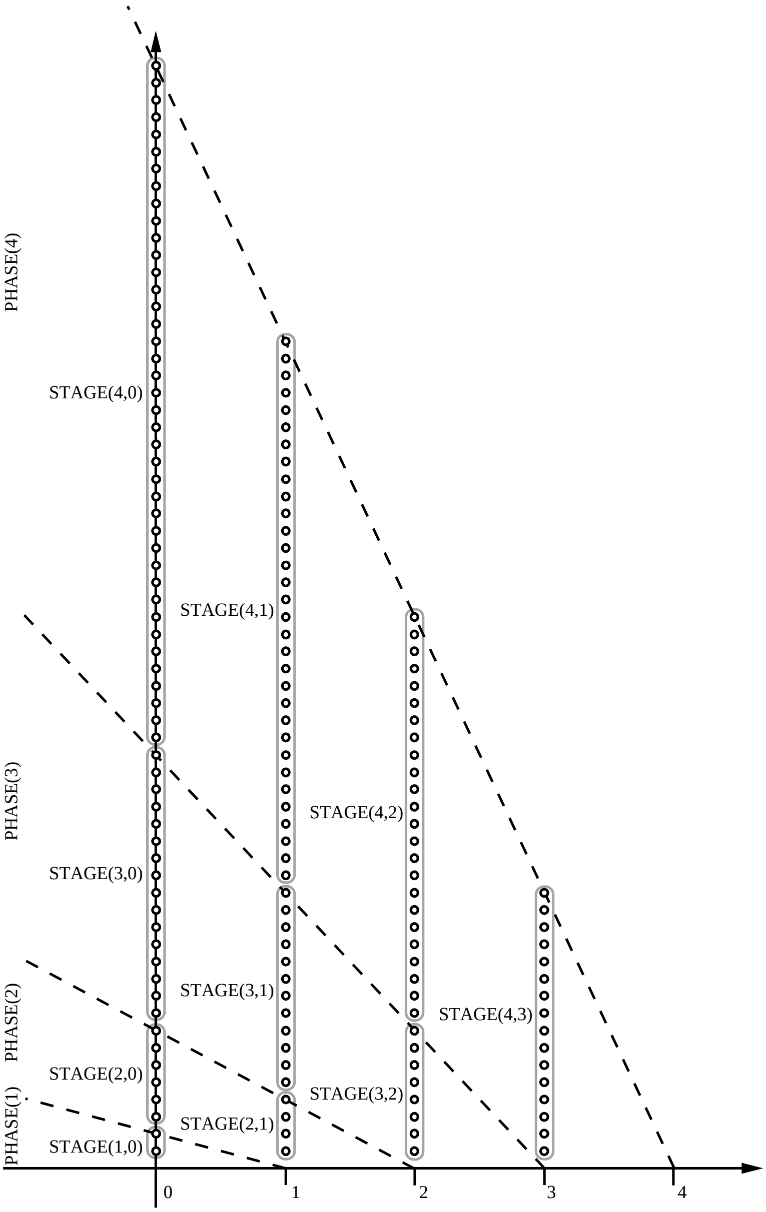

Fig. 1. The figure refers to the case $\alpha = 2$: the abscissa represents the executions of BROAD-B(2^{ℓ}), while the ordinate represents the time-slots of BROAD-B(2^{ℓ}).

Proof. The execution of BROAD-B(2^{ℓ_n}), for $\ell_n = \lceil \log n \rceil$, will be the good one and it has completion time $O(D\Delta \log(n/\Delta) \log n)$. By definition of BROAD$^{(\alpha)}$, it follows that all the nodes turn to the non active state within the last time-slot $t_{\text{end}} = O(D\Delta \log(n/\Delta) \log n)$ of BROAD-B(2^{ℓ_n}). We thus need to upper bound the time in which this happens, i.e., when PHASE(k_{end}) is completed, where k_{end} is the smallest integer such that $k_{\text{end}} > \ell_n$ and $f^{\alpha}_{k_{\text{end}}}(\ell_n) \geqslant t_{\text{end}}$. From the definition of f^{α}_{k}, it follows that

$$k_{\text{end}} \leqslant \log^{\alpha/2} t_{\text{end}}. \tag{7}$$

Let T be the number of time-slots required to complete PHASE(k_{end}). From the definition of phase and stage of the protocol, it holds that

$$T = \sum_{k=1}^{k_{\text{end}}} \sum_{\ell=0}^{k-1} (f^{\alpha}_{k}(\ell) - f^{\alpha}_{k-1}(\ell)) = \sum_{\ell=0}^{k_{\text{end}}-1} f^{\alpha}_{k_{\text{end}}}(\ell).$$

It thus follows that

$$T = \sum_{\ell=0}^{k_{\text{end}}-1} f^{\alpha}_{k_{\text{end}}}(\ell) \leqslant 2^{\lceil k_{\text{end}}^{2/\alpha} \rceil} \sum_{\ell=0}^{k_{\text{end}}-1} (k_{\text{end}} - \ell) \leqslant 2^{\lceil k_{\text{end}}^{2/\alpha} \rceil} k^2_{\text{end}}. \tag{8}$$

Finally, by combining Eq. (7) with Eq. (8), we get $T \leqslant t_{\text{end}} \log^{\alpha} t_{\text{end}}$ and so

$$T = \mathrm{O}(D\Delta \log(n/\Delta) \log^{1+\alpha} n). \qquad \square$$

5. Multi-broadcast: the BB model

We first need to extend Definition 4.1 to multi-broadcast operations. We assume here that each message has an header containing a unique ID number so that two messages have different ID numbers.

Definition 5.1. A multi-DB (*multi-RB*) protocol P is a protocol that, given a graph G and a set of r broadcast operations on it (in short, an r-broadcast operation, $r \geqslant 1$), works in time-slots (numbered $0, 1, \ldots$) according to the following rules.

1. In every time-slot, each node either acts as transmitter or as receiver. When transmitting, the node sends *one* message.
2. All the nodes share the same *message-recovering* function $\mathscr{R}$ that takes any message m as input and, if any, returns the (unique) broadcast message contained in m.
3. A node receives a message in a time-slot if and only if it acts as receiver and *exactly* one of its *in*-neighbors acts as transmitter in that time-slot.
4. The actions of a node in a specific time-slot are function of its own label, the number of current time-slot t, and the messages received during the previous time-slots (for multi-RB protocols, the actions also depend on the output of a random bit generators).

5.1. The lower bound for deterministic protocols

The next theorem provides a lower bound which is a function of the congestion c and n.

Theorem 5.1. *Let P be any multi-DB protocol. Then, for any $n \geqslant 4$ and $c \geqslant 2$, there exist an n-node directed graph G^P, with $D = 3$, $\Delta = 2$, and an r-broadcast operation* (*with* $r \geqslant c$) *on* G^P (*yielding a congestion* c), *such that* P *completes this operation in* $\Omega((c/\log c)\log n)$ *time*.

Proof. Without loss of generality, we consider the case in which c is arbitrary fixed and $r = c$ (i.e., maximum congestion). The graph G^P will be one of the family $\mathscr{G}_n$ described below. Any graph in $\mathscr{G}_n$ is an n-nodes directed graph of 3 levels. The first level consists of the (unique) source node (with label 1) containing all the r messages. The source is then connected to $n-2$ nodes that form the second level. Finally, the third level has only one node (with label n), i.e., the *sink*. The sink has exactly two in-neighbors among the nodes in the second level. We denote by $G_{u,v}$ the graph in which the sink has u and v as its in-neighbors. So,

$$\mathscr{G}_n = \{G_{u,v} \mid u, v \in \{2, 3, \ldots, n-1\}\}.$$

Since the sink cannot send any information to any other node, the execution of P, with respect to any non sink node, is the same for every graph of the family. Let T be defined as follows:

$$T = \max_{G_{u,v} \in \mathscr{G}_n} \{t \mid P \text{ has completion time } t \text{ on } G_{u,v}\}.$$

We represent the execution of the first T time-slots of P,with respect to a node v in the second level, as a sequence $\vec{x}_v$ over the alphabet $\{0, 1, \ldots, r\}$ with the following meaning: $\vec{x}_v(t) = z$ ($z \geqslant 1$) if, at time-slot t, v sends a message m such that $\mathscr{R}(m)$ is the zth source message, where $\mathscr{R}$ is the message-recovering function. Furthermore, $\vec{x}_v(t) = 0$ if, at time-slot t, either v acts as receiver or it sends a message m such that $\mathscr{R}(m)$ is not defined. We thus obtain a set $\mathscr{D}^P$ of $n-2$ sequences of length T over the alphabet $\{0, 1, \ldots, r\}$.

We claim that a necessary condition to complete the r-broadcast on every graph in $\mathscr{G}$ is that *any* two sequences $\vec{x}_u$ and $\vec{x}_v$, with $v \neq u$, must be *r-different*, i.e., for any element $z \in [r]$ there is a coordinate t for which the set $\{\vec{x}_u(t), x_v(t)\}$ is equal to $\{z, 0\}$.

Indeed, assume by contradiction that this is not true. So, there are two sequences $\vec{x}_u$, $\vec{x}_v$ with $v \neq u$ and $z \in [r]$, such that $\{\vec{x}_u(t), x_v(t)\}$ is not equal to $\{z, 0\}$ for every $t \leqslant T$. We then consider protocol P on the graph $G_{u,v}$. It is easy to verify that, in $G_{u,v}$, the sink is reachable from the source, but the sink does not receive the zth message during the first T time-slots.

From the above discussion, the $n-2$ sequences of length T in $\mathscr{D}^P$ must be pairwise r-different. From Theorem 3.6, in order to have a set of $n-2$ pairwise r-difference sequences of length T, it must hold that

$$T = \Omega\left(\frac{r}{\log r} \log n\right).$$

Since T is a lower bound on the worst-case completion time of P over the graph family $\mathscr{G}_n$, the theorem follows. □

Theorem 5.2. *Let* P *be any multi-DB protocol. Then*, *for any* $n \geqslant 4$, $3 \leqslant D \leqslant n/6$ and $2 \leqslant c \leqslant r$, *it is possible to define an n-node directed graph* G^P *with maximum* *eccentricity D and maximum in-degree $\Delta = 2$, and a set of r independent broadcast operations on G^P (yielding a congestion c) such that P completes these operations on G^P in $\Omega(c + (c/\log c + D)\log n)$ time.*

Proof. Since a node can receive at most one message per time-slot, it is easy to obtain the lower bound $\Omega(c)$. Then, by combining the family of graphs of Theorem 5.2 (with $\Delta = 2$) with that yielding Theorem 4.2, we easily get the thesis. □

Notice that the proof of Theorem 5.1 does not rely on the fact that the nodes do not know n. Thus, Theorem 5.2 also holds under this condition.

Finally, we observe that the above construction can be easily modified in order to let each of the r broadcast messages start from a different source (i.e. we have r messages in r different sources). It suffices to replace the source node in $\mathscr{G}_n$ (see the proof of Theorem 5.1) with the root of a binary directed tree in which the r sources are the leaves of the tree. The only difference is that D is now a logarithmic function of n.

5.2. Lower bound for randomized protocols

Any worst-case time bound of a randomized protocol can be considered reliable if it happens within a high probability on every possible instance. This concept is widely adopted in the field of randomized algorithms [34], and it can be easily adapted to the case of multi-broadcast operations on unknown networks.

Definition 5.2. A multi-RB protocol P has completion time T (where, clearly, T depends on n) if, for any $n \geqslant 1$ and for any n-node graph G, P completes, with probability at least $1 - 1/n$ any r-broadcast operation on G within T time-slots.

Lemma 5.1. *Let P be a multi-RB protocol for unknown networks. If P has completion time T, then it must holds that $T = \Omega((c/\log c)\log n)$.*

Proof. The proof makes use of the families $\mathscr{G}_n$ ($n \geqslant 4$) of directed graphs, and the corresponding multi-broadcast operations, which have been introduced in the proof of Theorem 5.1. In particular, we will show that, for any $n \geqslant 4$, a graph $G^P \in \mathscr{G}_n$ exists on which P has $\Omega((c/\log c)\log n)$ completion time with probability larger than $1/n$. As in the proof of Theorem 5.1, an execution of T time-slots of P, with respect to the nodes of the second level of $G_{u,v}$, can be represented as a set $\mathscr{D}$ of $n-2$ sequences of length T over the alphabet $\{0, \ldots, r\}$ (with the same meaning of that given in the proof of Theorem 5.1). A multi-RB protocol (restricted to the nodes in the second level of $G_{u,v}$) can thus be seen as a probability distribution $\mathscr{P}$ over the set $\mathscr{A}$ of all possible sequence sets $\mathscr{D}$. Consider the following function

$$\chi^{\mathscr{D}}(u,v) = \begin{cases} 1 & \text{if sequences } u \text{ and } v \text{ in } \mathscr{D} \text{ are } r\text{-different,} \\ 0 & \text{otherwise.} \end{cases}$$

From the hypothesis of the theorem we have that, for all $u \neq v$,

$$\Pr\{P \text{ has completion time } T \text{ on } G_{u,v}\} \geqslant 1 - \frac{1}{n},$$

it follows that

$$\forall u \neq v \quad \sum_{\mathscr{D}\in\mathscr{A}} \mathscr{P}(\mathscr{D})\chi^{\mathscr{D}}(u,v) \geqslant 1 - \frac{1}{n}. \tag{9}$$

Consider now the sum

$$M^{\mathscr{D}} = \sum_{u,v\in L_2} \chi^{\mathscr{D}}(u,v),$$

where L_2 denotes the nodes of the second level of any graph in $\mathscr{G}_n$; this equals the number of r-different pairs yielded by the protocol P. We now prove that there has to exist a $\bar{\mathscr{D}}$ such that

$$M^{\bar{\mathscr{D}}} \geqslant \left(1 - \frac{1}{n}\right)\frac{(n-2)(n-3)}{2}. \tag{10}$$

In fact, from Eq. (9), we have that

$$\sum_{u,v\in L_2} \sum_{\mathscr{D}\in\mathscr{A}} \mathscr{P}(\mathscr{D})\chi^{\mathscr{D}}(u,v) \geqslant \left(1 - \frac{1}{n}\right)\frac{(n-2)(n-3)}{2}$$

and, hence

$$\sum_{\mathscr{D}\in\mathscr{A}} \mathscr{P}(\mathscr{D}) \sum_{u,v\in L_2} \chi^{\mathscr{D}}(u,v) = \sum_{\mathscr{D}\in\mathscr{A}} \mathscr{P}(\mathscr{D})M^{\mathscr{D}} \geqslant \left(1 - \frac{1}{n}\right)\frac{(n-2)(n-3)}{2}.$$

It follows that a $\bar{\mathscr{D}}$ exists that verifies Eq. (10). We now prove that $\bar{\mathscr{D}}$ must contain a large subset of sequences which are pairwise r-different, thus the same property derived in the proof of Theorem 5.1.

Claim. *There exists a subset $\bar{\mathscr{S}} \subseteq \bar{\mathscr{D}}$ such that $|\bar{\mathscr{S}}| = n/2$ and, for any $u,v \in \bar{\mathscr{S}}$ (with $u \neq v$), the pair (u,v) is r-different.*

Proof. The following simple algorithm finds the desired subset $\bar{\mathscr{S}}$ (we assume here that n is an even number).

begin
Choose an arbitrary $\bar{\mathscr{S}} \subseteq \bar{\mathscr{D}}$ s.t. $|\bar{\mathscr{S}}| = n/2$
$\bar{\mathscr{S}}^c := \emptyset$:
while ($\bar{\mathscr{S}}$ does not satisfy the claim) **do**
  begin
  Choose (arbitrarily) $u,v \in \bar{\mathscr{S}}$ that are not r-different
  Choose (arbitrarily) $k \in \bar{\mathscr{D}} \backslash (\bar{\mathscr{S}} \cup \bar{\mathscr{S}}^c)$
  $\bar{\mathscr{S}} := (\bar{\mathscr{S}} - \{v\}) \cup \{k\}$;
  $\bar{\mathscr{S}}^c := \bar{\mathscr{S}}^c \cup \{v\}$
  end
return $\bar{\mathscr{S}}$.
end

We first notice that, from Eq. (10), there are at most $(n-3)/2$ different pairs in $\bar{\mathscr{D}}$ that are not r-different. Since at every iteration of the while loop the algorithm "discards" a new not r-different pair in $\bar{\mathscr{D}}$, the algorithm always returns a set $\bar{\mathscr{S}}$ satisfying the claim in $O(n)$ steps.

The claim implies that the multi-RB protocol P on the family $\mathscr{G}_n$ must yield a set of $n/2$ sequences that are pairwise r-different. So, by applying the lower bound of Theorem 3.6, we can state that

$$T = \Omega\left(\frac{r}{\log r}\log\frac{n}{2}\right) = \Omega\left(\frac{r}{\log r}\log n\right).$$

Since T is a lower bound on the completion time of P, the lemma follows. □

The proof of the following theorem is an easy consequence of the $\Omega(D\log(n/D))$ lower bound for randomized protocols given in [31], the trivial lower bound $\Omega(c)$, and Lemma 5.1.

Theorem 5.3. *Let P be any multi-RB protocol. Then, for any $n \geqslant 4$, $D \geqslant 3$ and $2 \leqslant c \leqslant r$, there exist* (i) *an n-node directed graph G^P with maximum eccentricity D and maximum in-degree $\Delta = 2$, and* (ii) *a set of r independent broadcast operations on G^P* (*yielding a congestion c*), *such that the completion time of P is $\Omega(c + (c/\log c)\log n + D\log(n/D))$.*

The above theorem in fact holds for any probability lower bound of the form $1 - 1/n^a$ (with any fixed constant $a > 0$). However, we choose the form $1 - 1/n$ in order to simplify the proof of Lemma 5.1.

5.3. The multi-broadcast protocols

As mentioned in the Introduction, our multi-DB protocols make use of superimposed codes [19,25,26,16]. In particular, we will use the following upper bounds [19,28].

Theorem 5.4. *For any $n \geqslant 3$ and for $k \geqslant 2$:*

there exist (n,k)-strongly selective families of size $O(\min\{n, k^2\log n\})$ [19];

it is possible to construct, in time polynomial in n and k, an (n,k)-strongly selective family (*based on q-ary* error-correcting codes) *of size* $O(\min\{n, k^2\log^2 n\})$ [28, p. 370].

We recall that an $\Omega((k^2/\log k)\log n)$ lower bound is proved in Theorem 3.4.

In what follows, we describe the protocol MULTI-BB-BROAD-A(n,Δ): It assumes that nodes know Δ and n. However, when Δ and/or n are not known, we can adopt the same dovetail technique described in Section 4.2. The cost of this further task is $O(\log^{1+\alpha} n)$ additional time-slots (for any fixed $\alpha > 0$) per each time-slot of MULTI-BB-BROAD-A(n,Δ).

5.3.1. Description of protocol MULTI-BB-BROAD-A(n, Δ)

With each message is associated a priority so that the priorities induce a total ordering on the set of broadcast messages.[5] The priorities will be used by the nodes to schedule the messages to send. In fact, every node stores the messages by means of a priority queue. The protocol MULTI-BB-BROAD-A(n, Δ) uses an $(n, \Delta+1)$-strongly selective family $\mathscr{F} = \{F_1, F_2, \ldots, F_{|\mathscr{F}|}\}$. It consists of a sequence of consecutive *phases*. Each phase consists of $|\mathscr{F}|$ time-slots. At the very beginning, the priority queues of the source nodes contain their broadcast messages, the other priority queues are empty. At the beginning of each phase, every node v with a non-empty priority extracts from the queue the message m_v of highest priority. At the jth time-slot of the phase, node v acts according to the following rules:

- If the label of v belongs to F_j and m_v exists then v transmits m_v.
- In all the other cases, v acts as a receiver. If v receives a message that v has never received before, the message is enqueued, otherwise, the message is discarded.

Theorem 5.5. *Protocol* MULTI-BB-BROAD-A(n, Δ) *completes any r-broadcast operation within* $\mathrm{O}((D + c)\min\{n, \Delta^2 \log n\})$ *time-slots on any n-node graph of maximum eccentricity D, maximum in-degree Δ, and congestion c.*

Proof. Firstly, we prove the following:

Claim 1. *If a node v transmits a message during a phase t then all the out-neighbors of v receive the message by the end of phase t.*

Proof. In fact, let u be any out-neighbor of v. By virtue of the $(n, \Delta+1)$-strong selectivity of $\mathscr{F}$, there is a time-slot of phase t in which u acts as a receiver and v is the only node, among the in-neighbors of u (notice that these are at most Δ), that transmits. Hence, in that time-slot, u receives the message of v. □

Now, we can show that all messages reach their destinations.

Claim 2. *After a finite number of phases, any message m is received by all the nodes that are reachable from the source of m.*

Proof. The proof is by induction on the distance ℓ from the source. For $\ell = 0$ the claim is obvious. Let u be a node at distance $\ell + 1$ from the source. Consider an in-neighbor v of u at distance ℓ. From the inductive hypothesis, node v receives the message m. The first time this happens, m is enqueued in the priority queue of v. According to the protocol, at the beginning of each phase, node v extracts the message of highest priority and transmits it. Since there are less than r messages of priority higher than that of m and any message is enqueued at most once, it follows that v

[5] A possible choice for the priorities is the following. If m is the hth message of a source of label l then the priority of m is given by the pair (l, h). Thus, for any two messages m and m' of priorities (l, h) and (l', h'), m has priority higher than m' if either $l < l'$ or $l = l'$ and $h < h'$.

extracts and transmits m within at most r phases. Thus, from Claim 1 node u receives the message m. □

The last step consists of showing that the messages cannot be delayed too much.

Claim 3. *Assume that node u, at distance ℓ from the source of a message m, receives m for the first time during phase $\ell+t_u$. Then, during the first $\ell+t_u$ phases, u transmits at least t_u-1 messages with higher priority than that of m.*

Proof. The proof is by induction on the distance ℓ. For $\ell=0$ the claim is obvious. Let u be a node at distance $\ell+1$ and let v be the node at distance ℓ from which u receives m for the first time. Let $\ell+t_v$ be the phase in which node v receives m for the first time. By inductive hypothesis, in the first $\ell+t_v$ phases, v transmits at least t_v-1 messages of higher priority (than that of m). This fact and Claim 1 imply that, in the first $\ell+t_v$ phases, u receives from v a set M_1 of at least t_v-1 messages of higher priority. Now, let $k \geqslant 0$ be the number of phases during which the message m remains in the priority queue of v. This implies that v transmits a set M_2 of k further messages of higher priority before transmitting m. As a consequence, u receives all the messages in M_2 before receiving m.

Let $t_u=t_v+k$; then node u receives m for the first time in phase $\ell+1+t_u$. When this happens u has received from v at least $t_v-1+k=t_u-1$ messages of higher priority (i.e., all the messages in $M_1 \cup M_2$). Since all these messages have been received in distinct phases among the first $\ell+t_u$, the node u transmits at least t_u-1 messages of higher priority in the first $\ell+1+t_u$ phases. □

Notice that any message m has to reach nodes at distance at most D from its source and there are less than c messages of priority higher than m that collide with m. Hence, Claims 2 and 3 imply that any message reaches its destinations in less than $D+c$ phases. Since each phase requires $O(\min\{n, \Delta^2 \log n\})$ time-slots (see Theorem 5.4) the thesis follows. □

6. Conclusions and open problems

The main contribution of this paper is that of providing explicit and strong connections between a set of old and new combinatorial tools and the issue of broadcast operations in radio networks of unknown topology. Due to such connections, we have obtained new lower and upper bounds on the completion time of this important operation. We believe that the concept of selectivity and that of r-different sequences can have further applications to other distributed models in which the local knowledge is extremely low. Some evidence of our opinion is also given by some previous results [33] in which strongly selective families have been used for the distributed coloring problem in unknown graphs.

As for specific research directions, the following ones appear to be the most relevant.

The lower bounds for multi-DB (and multi-RB) protocols are consequences of the combinatorial lower bound $\Omega(r/\log r)\log n)$ on the length of n pairwise r-different sequences given in Theorem 3.6. We do not know whether the latter is tight. As far as we know, the best upper bound is $O(r\log n)$. The deterministic upper bounds in Theorem 5.5 almost match the lower bound in Theorem 5.2, when $\Delta = O(\text{poly}\log n)$. An interesting future research goal is that of reducing the gap between upper and lower bounds for larger Δ. To this aim, we believe that a generalization of the $\Omega((r/\log r)\log n)$ lower bound to d-wise r-different sequences could give a stronger lower bound that also depends on Δ.

A further research direction is that of using the same combinatorial tools to investigate the issue of *dynamical-fault tolerance* in radio networks. Dynamical edge and node faults may happen at any instant, *even during the execution of a protocol*. Some results on this direction have been recently obtained in [15].

Acknowledgements

We wish to thank Jànos Körner for helpful discussions. We are also very grateful to Andrzej Pelc for pointing us an incorrect statement in the conference version of our paper.

References

[1] N. Alon, A. Bar-Noy, N. Linial, D. Peleg, A lower bound for radio broadcast, J. Comput. System Sci. 43 (1991) 290–298 (preliminary version in 21st ACM STOC 1989).
[2] N. Alon, J. Spencer, The Probabilistic Method, Wiley, New York, 1992.
[3] R. Bar-Yehuda, O. Goldreich, A. Itai, On the time-complexity of broadcast operations in multi-hop radio networks: an exponential gap between determinism and randomization, J. Comput. System Sci. 45 (1992) 104–126 (preliminary version in 6th ACM PODC, 1987).
[4] R. Bar-Yehuda, A. Israeli, A. Itai, Multiple communication in multi-hop radio networks, SICOMP 22 (4) (1993) 875–887.
[5] S. Basagni, D. Buschi, I. Chlamtac, A mobility-transparent deterministic broadcast for ad hoc networks, IEEE/ACM Trans. Networking 7 (1999) 799–807.
[6] D. Bruschi, M. Del Pinto, Lower bounds for the broadcast problem in mobile radio networks, Distributed Comput. 10 (1997) 129–135.
[7] I. Chlamtac, A. Faragó, Making transmission schedule immune to topology changes in multi-hop packet radio networks, IEEE/ACM Trans. Networking 2 (1994) 23–29.
[8] I. Chlamtac, A. Faragó, H. Zhang, Time-spread multiple access (TSMA) protocols for multihop mobile radio networks, IEEE/ACM Trans. Networking 5 (1997) 804–812.
[9] B.S. Chlebus, L. Gąsieniec, A.M. Gibbons, A. Pelc, W. Rytter, Deterministic broadcasting in unknown radio networks, Distributed Comput. 15 (2002) 27–38 (preliminary version in 11th ACM-SIAM SODA 200).
[10] B.S. Chlebus, L. Gąsieniec, A. Ostlin, J.M. Robson, Deterministic radio broadcasting, 27th ICALP, Lecture Notes in Computer Science, Vol. 1853, 2000, pp. 717–728.
[11] M. Chrobak, L. Gąsieniec, W. Rytter, Fast broadcasting and gossiping in radio networks, 41st IEEE FOCS, 2000, pp. 575–581.
[12] S. Chaudhuri, J. Radhakrishnan, Deterministic restriction on circuit complexity, 28th ACM STOC, 1986, pp. 30–36.
[13] A.E.F. Clementi, A. Monti, R. Silvestri, Selective families, superimposed codes, and broadcasting in unknown radio networks, 12th ACM-SIAM SODA, 2001, pp. 709–718.

[14] A.E.F. Clementi, A. Monti, R. Silvestri, Distributed multi-broadcast in unknown radio networks, 20th ACM PODC, 2001, pp. 255–263.
[15] A. Clementi, A. Monti, R. Silvestri, Round robin is optimal for fault-tolerant broadcasting on wireless networks, 9th ESA, Lecture Notes in Computer Science, Vol. 2161, 2001, pp. 452–463.
[16] R. Cole, R. Hariharan, P. Indyk, Tree pattern matching and subset matching in deterministic $O(n \log^3 n)$-time, 10th ACM-SIAM SODA, 1999, pp. 245–254.
[17] K. Diks, E. Kranakis, D. Krizanc, A. Pelc, The impact knowledge on broadcasting time in radio networks, 7th ESA, Lecture Notes in Computer Science, Vol. 1643, 1999, pp. 41–52.
[18] A.G. Dyachkov, V.V Rykov, A survey of superimposed code theory, Problems Control Inform. Theory 12 (1983) 1–13.
[19] P. Erdös, P. Frankl, Z. Füredi, Families of finite sets in which no set is covered by the union of r others, Israel J. Math. 51 (1985) 79–89.
[20] E. Fachini, J. Körner, E. Fachini, J. Körner, A note on counting very different sequences, Combin. Probab. Comput. 10 (6) (2001) 501–504.
[21] C. Faloutsos, Signature files, in: W.B. Frakes, R. Baeza-Yates (Eds.), Information Retrieval—Data Structures and Algorithms, Prentice-Hall, Englewood Cliffs, NJ, 1992.
[22] P. Frankl, Z. Füredi, Forbidding just one intersection, J. Combin. Theory Ser. A 39 (1985) 160–176.
[23] Z. Füredi, On r-cover-free families, J. Combin. Theory Ser. A 73 (1996) 172–173.
[24] R. Gallager, A perspective on multiaccess channels, IEEE Trans. Inform. Theory (special issue on random access communications) 31 (1985) 124–142.
[25] F.K. Hwang, V.T. Sös, Non-adaptive hypergeometric group testing, Studia Sci. Math. Hungarica 22 (1987) 257–263.
[26] P. Indyk, Deterministic superimposed coding with application to pattern matching, 38th IEEE FOCS, 1997, pp. 127–136.
[27] P. Indyk, Explicit construction of selectors and related combinatorial structures, with applications, ACM-SIAM SODA, 2002, pp. 697–704.
[28] W.H. Kautz, R.C. Singleton, Nonrandom binary superimposed codes, IEEE Trans. Inform. Theory 10 (1964) 363–377.
[29] D.R. Kowalski, A. Pelc, Deterministic broadcasting time in radio networks of unknown topology, 43rd IEEE FOCS, 2002, pp. 63–72.
[30] E. Kranakis, D. Krizanc, A. Pelc, Fault-tolerant broadcasting in radio networks, 6th ESA, Lecture Notes in Computer Science, Vol. 1461, Springer, Berlin, 1998, pp. 283–294.
[31] E. Kushilevitz, Y. Mansour, An $\Omega(D \log(n/D))$ lower bound for broadcast in radio networks, SICOMP 27 (1998) 702–712 (preliminary version in 12th ACM PODC 1993).
[32] E. Kushilevitz, Y. Mansour, Computation in noisy radio networks, Ninth ACM-SIAM SODA, 1998, pp. 236–243.
[33] N. Linial, Locality in distributed graph coloring, SICOMP 21 (1992) 193–201.
[34] R. Motwani, P. Raghavan, Randomized algorithms, Cambridge University Press, Cambridge, 1995.
[35] E. Pagani, G. Rossi, Reliable broadcast in mobile multihop packet networks, Third ACM-IEEE MOBICOM, 1997, pp. 34–42.
[36] K. Pahlavan, A. Levesque, Wireless Information Networks, Wiley-Interscience, New York, 1995.
[37] T.S. Rappaport, Wireless Communications: Principles and Practice, Prentice-Hall, Englewood Cliffs, NJ, 1996.
[38] L.G. Roberts, Aloha packet system with and without slots and capture, ASS Notes 8, Advanced Research Projects Agency, Network Information Center, Stanford Research Institute, 1972.
[39] J.E. Wieselthier, G.D. Nguyen, A. Ephremides, On the construction of energy-efficient broadcast and multicast trees in wireless networks, IEEE INFOCOM, 2000, pp. 585–594.