Contention Resolution in a Non-Synchronized Multiple Access Channel ‡

Gianluca De Marco* Dariusz R. Kowalski†

Abstract

Multiple access channel is a well-known communication model that deploys properties of many network systems, such as Aloha multi-access systems, local area Ethernet networks, satellite communication systems, packet radio networks. The fundamental aspect of this model is to provide efficient communication and computation in the presence of restricted access to the communication resource: at most one station can successfully transmit at a time, and a wasted round occurs when more than one station attempts to transmit at the same time. In this work we consider the problem of contention resolution in a multiple access channel in a realistic scenario when up to k stations out of n join the channel at *different times*. The goal is to let at least one station to transmit alone, which results in successful delivery of the message through the channel.

We present three algorithms: two of them working under some constrained scenarios, and achieving optimal time complexity $\Theta(k \log(n/k) + 1)$, while the third general algorithm accomplishes the goal in time $O(k \log n \log \log n)$.

Keywords — Multiple access channel; Contention resolution; Deterministic algorithms; Distributed algorithms.

1 Introduction

Multiple access channels are well-known communication media that form the basis to many extensively studied network systems such as Aloha multi-access systems, local area Ethernet networks, satellite communication systems, packet radio networks [1, 2, 4].

Preliminaries. The model that is at the basis of theoretical studies of the multiple-access channel can be defined in the following terms. We are given a set of stations each of them having a unique integer *ID* from the set $\{1, 2, \ldots, n\} = [n]$, for some integer n.

These stations communicate by sharing one communication channel. At each time slot any station can either transmit a packet of data to the channel or listen to the channel. Notice that parameter n also establishes an upper bound on the number of stations that can be attached to the channel.

*Dipartimento di Informatica, Università di Salerno, 84084 Fisciano (SA), Italy.
Email: `demarco@dia.unisa.it`

†Department of Computer Science, University of Liverpool, Liverpool L69 3BX, UK.
Email: `d.kowalski@liverpool.ac.uk`

‡A preliminary version of this paper has been presented at the 27th Annual IEEE International Symposium on Parallel and Distributed Processing, (IPDPS 2013), Cambridge, MA, USA, May 2013.
This work was supported by the Polish National Science Centre grant DEC- 2012/07/B/ST6/01534.

There is no central unit controlling the stations. A transmission is successful at a given time slot if and only if at that time slot there is only one transmitting station; in such case all stations get the message (including the one which transmitted the data, as it possesses it by default). If two or more stations transmit simultaneously in a given time slot, the messages collide and the transmission is not successful — *i.e.*, no station receives any of the transmitted messages.

One of the fundamental problems in this context is a *contention resolution* problem, also known as a *wakeup problem*, where at least one station among those who joined the computation on the channel has to transmit successfully (*i.e.*, alone) on the channel. Of course, the possibility of having collisions among the transmissions makes this task particularly difficult. An *algorithm* for the wake-up problem is a collection of n transmission schedules, one for each station, which eventually allows exactly one of the active stations to transmit on the channel, therefore waking up every other station. Once one of the active stations manages to send its message successfully on the channel, the message is heard by all other stations. The efficiency of the algorithm is measured by the time complexity, *i.e.* the number of time slots necessary to find the first time slot at which the transmission schedules allow exactly one station to send a message, counted from the first slot with at least one active station.

In the literature, there are many important assumptions that can be made on the model described above, each of them may have an impact on the time complexity. The first important assumption is about the amount of feedback received from the channel in the case of collision. Substantially, two different scenarios are studied in the literature. In the *collision detection* model, any station is able to hear an *interference noise* in the case of collision, allowing it to deduce the information that two or more stations tried to transmit in a given time slot. A weaker scenario, used in the present paper, assumes that no feedback signal is supplied by the channel in the case of collision, making it consequently impossible to distinguish between an occurred collision and the case where no station transmits. Another crucial assumption concerns whether all the active stations wake-up simultaneously or as in the more general case considered in this paper, the stations wake up, spontaneously and independently, in different time slots. Finally, the third central issue is the measurement of the elapsed time. Essentially the possibilities range between two extreme situations: the *globally* synchronous and the *locally* synchronous model. In the first model all the stations have access to a global clock. When a processor wakes up, it can see the current round number ticked by the clock. The other model is weaker. Each station has its own local clock, therefore, although the communication is synchronous, (*i.e.* all the clocks tick with the same rate) there is no global round number visible by every station. Each station can start counting the time from the time slot it wakes up, without knowing anything about the other round numbers. In this paper, the globally synchronous model is considered. The contention resolution problem considered in this paper can be formally defined as follows.

The contention resolution problem. We are given a multiple access channel where each station knows only its own *ID* from the known range $[n]$. Some number k of stations, with $1 \leq k \leq n$, can *spontaneously and independently* wake up, *i.e.*, each of them can start its activity at any moment. Let $s \geq 0$ be the first time slot such that some station is woken up. The problem is to assign transmission schedules to the stations, one per each station, such that there exists a time slot $t \geq s$ at which exactly one station (among the conflicting awaken stations at time t) transmits. We consider the worst-case scenario over all possible patterns of spontaneous wake up times of stations and measure the efficiency by the number of time slots between the first spontaneous wakeup and the first successful transmission, *i.e.*, $t - s$.

It must be stressed that most of the collision resolution research on multi access communication has its main motivation from the fact that very often most transmitters are inactive most of the time, while only a few are busy. If all n stations connected to the channel were active, one could apply one of the simplest schedules to resolve conflicts: the *time division multiplexing protocol*. This means that when there are n stations, n time slots will be needed. Of course, this becomes very inefficient when the maximum number k of possible awaken stations is very small compared to n. Moreover, given the fully distributed nature of the system, it is often unrealistic to assume that the stations can rely on the knowledge of the bound k or the starting time s, as both these parameters depend on actions taken independently by the participating stations without any sort of coordination. This paper focuses on the impact that the knowledge of the parameters k and s can have on the time complexity of the wake-up problem in the realistic scenario when the stations wake up spontaneously and independently in different time slots. Namely, we consider the following three scenarios.

Scenario A (s is known). Each station knows its own *ID* and the parameter n. In addition, every station knows the starting time s, *i.e.* the first time at which some station has woken up.

Scenario B (k is known). Each station knows its own *ID* and the parameter n. In addition, every station knows the maximum number k of possible awaken stations, but doesn't have any a priori knowledge about the wakeup times of other stations (included the starting time s).

Scenario C (neither s nor k is known). Every station knows only its own *ID* and parameter n.

Previous and related work. The collision resolution research for the multi access communication began in 1970 with Abramson's ALOHA network [1].

Komlós and Greenberg [25] were the first to consider the typical situation when a subset of k among n stations are awakened and have messages, and *all of them* need to be sent (successfully) to the multiple access channel as soon as possible. Contrasted with our wake-up problem, their situation is more general in the sense that they seek for algorithms that will eventually allow *every* station to transmit its own message to the channel. (Of course, their algorithm, stopped at the first successful message sent, is actually a wake-up algorithm.) On the other hand, our model is more general in the sense that we allow any station (among a subset of at most k stations) to wake up and become active at any moment spontaneously and independently; while in their setting all k stations become active simultaneously at the beginning of the computation. They showed how to solve the problem in time $O(k + k\log(n/k))$, where either n or k is known. A lower bound of $\Omega(k(\log n)/(\log k))$ was then proved by Greenberg and Winograd [23], even if collision detection is available at stations. In the setting without collision detection, the best known lower bound on contention resolution, understood as a process of selecting *one* among the contending stations, was given by Clementi et al. [14]. They showed that for any $2 \leq k \leq n/64$, $\Omega(k\log(n/k))$ rounds are needed. All these results were under assumption that all participating stations wake up at the same time.

The first formalization of the wake-up problem on a *non-synchronized* multiple access channel is due to Gąsieniec, Pelc and Peleg [22]. They introduced many variations and assumptions concerning synchrony and knowledge. They also considered the randomized counterpart of the problem. For deterministic solutions in the globally synchronous model, which is the setting considered in the present paper, they showed an optimal algorithm that in time n solves the wake-up problem in the case the stations know n. In case of unknown n, they propose a wake-up algorithm working in time $4n$ in the worst case. The authors of [22] also introduced the most extreme model of synchronization, the locally synchronous model, and gave upper and lower bound both for n known and n unknown. The currently best upper bound for the locally synchronous model with n unknown has been set in [20]. It must be noted that, contrasted with our setting, in these works the number of possible awaken processors is not upper bounded as in our case (parameter k), but can eventually become as large as n. From this perspective, our situation is more general. On the other hand, in the model with known n, Chlebus et al. [9] presented a contention resolution protocol in the locally synchronous model, working in $O(k \log^2 n)$ rounds. As for the randomized solutions (not studied here), important improvements have been provided by Jurdziński and Stachowiak [24].

The contention resolution problem has been also studied in the more general framework of multi-hop radio networks, particularly in the context of problems such as (multi-)broadcast, gossip and others in the so-called blindfold model, *i.e.* in total absence of knowledge about topology and network parameters [8, 13, 15, 16].

Developments where similar issues on selecting stations (included broadcasing in multi-hop radio networks) under many assumptions, mainly regarding knowledge and synchrony, can be found in [5, 6, 7, 9, 10, 12, 17, 18, 19, 20].

Our results. To the best of our knowledge, this paper is the first to study the time complexity of contention resolution in the three scenarios defined above, among which Scenario C is the most general and realistic one. Our main result is an algorithm resolving contention, for the most general Scenario C, in time $\Theta(k \log n \log\log n)$ rounds. Our result is existential. More precisely, we introduce a combinatorial tool, a *waking matrix* (used by the stations running the algorithm), for which we prove the existence (see Subsection 5.3) by the probabilistic method. This represents the main technical challenge of the present paper. Note that the complexity of the algorithm differs by factor $O(\log\log n)$ from the lower bound, and is substantially better than the best known contention resolution protocol in the locally synchronous model given by Chlebus et al. [9]. For Scenarios A and B, we present two simple algorithms which resolve contention in $\Theta(k \log(n/k) + 1)$ rounds, which, as we will see, is optimal.

2 A lower bound

In this section we prove the following lower bound on the wake-up problem, which, though relatively simple, allows to match the upper bound for the problem in case of large values of k.

Theorem 2.1 *The wake-up problem requires* $\min\{k, n-k+1\}$ *rounds, even if the stations start simultaneously and parameters k and n are known.*

Proof: We know from the definition of the problem that $k \leq n$. If $k = n$, the proof is complete, since at least one round is needed to solve the wake-up problem. Therefore, assume from now on that $k < n$; we will count the number of additional rounds (apart from the initial round r) needed by the algorithm to accomplish its task.

An algorithm for the wake-up problem has to guarantee that for any set $X \subseteq [n]$ of k elements (representing stations), there will be a round such that one, and exactly one, station, among those in X, is selected, *i.e.*, transmits at that round. Let $X = \{x_1, x_2, \ldots, x_k\}$ be any set of k elements. In order to be correct, the algorithm must reserve a round r at which one, and exactly one, element of X transmits. That is, if we denote by T the set of transmitting stations at round r, the following property has to hold: there exists an element x such that $X \cap T = \{x\}$.

Let $Y = [n] \setminus X$. Let us pick any element y from Y and let $X' = (X \setminus \{x\}) \cup \{y\}$. Now the algorithm, in order to be correct, must also include a round r', $r' \neq r$, such that, if we denote by T' the set of transmitting stations at time r', the following property has to hold: there exists x' such that $T' \cap X' = \{x'\}$. Continuing this process, we can now build a new set X'' by substituting the selected element $x' \in X'$ with a *new* (*i.e.*, not considered before) element of its complement $[n] \setminus X'$. We can iterate as far as we are able to build distinct sets of k elements, *i.e.* $\min\{k, n-k\}$ times, where each iteration corresponds to a new round that the algorithm has to spend, and each such round involves one newly considered element in X and one newly considered element in Y. Therefore, if $k \leq n-k$, then the algorithm spends k rounds. If, on the other hand, $k > n-k$ then the algorithm spends a total of $(n-k)+1$ rounds (counting also the initial round r). This concludes the proof. □

As a straightforward consequence, and using the fact that $\log x = \Theta(x-1)$ for $x \in (1, const.]$, we have the following result.

Corollary 2.1 *For any constant c, if $k > n/c$ then $\Omega(n-k+1)$ rounds are needed. Moreover, for $k > n/c$, $n-k+1 = \Theta(k\log(n/k)+1)$.*

3 Contention resolution when the starting time s is known

In this section we show that when the stations know the first time slot s at which some stations become active, a simple *optimal* algorithm can be used to achieve the wake-up in time $\Theta(k\log(n/k)+1)$. Before giving the algorithm, we need to introduce the notion of *selective family*, a well-known combinatorial tool used in deterministic radio communication (see [8, 25]). Given an integer n, a (n,k)-*selective family*, for $2 \leq k \leq n$, is a family $\mathcal{F}$ of subsets of $[n] = \{1, 2, \ldots, n\}$ such that for any $X \subseteq [n]$ with $k/2 \leq |X| \leq k$, the following *selectivity property* holds: there exists a set $F \in \mathcal{F}$ such that $|X \cap F| = 1$. As a consequence of their already mentioned result (see the "Previous and related work" at the end of the Introduction), the authors of [25] showed that for all n and i, $1 \leq i \leq \log n$, there are $(n, 2^i)$-selective families of length $|\mathcal{F}| = O(k + k\log(n/k))$, where $k = 2^i$. The standard use of a selective family is as follows. A station $x \in X$ *transmitting according to a selective family* $\mathcal{F} = \{F_1, F_2, \ldots, F_{|\mathcal{F}|}\}$ of length $|\mathcal{F}|$ will transmit at time j if and only if $x \in F_j$. The F_j's are referred to as *transmission sets* in that they contain the nodes that are allowed to transmit at time j. This way, if the set X of transmitting stations has size falling in an interval $(2^{i-1}, 2^i)$ for some integer i, the selectivity property of a $(n, 2^i)$-selective family $\mathcal{F}$ guarantees that there exists a time t, $1 \leq t \leq |\mathcal{F}|$, at which only one station in X transmits. The length $|\mathcal{F}|$ of the selective family determines the time complexity of the algorithm.

The algorithm that we propose, called `wakeup_with_s`, is an interleaving between the well-known `round-robin`, exploited for particularly large values of k, and a new simple algorithm, called `select_among_the_first`, which is better for the large majority of cases, namely when k is smaller than a constant fraction of n. Notice that the interleaving is a very easy operation in a scenario with global clock (*e.g.*, one can execute `round-robin` in odd rounds and the other algorithm in even rounds) and does not require any knowledge about k.

It is straightforward to observe that the `round-robin` algorithm guarantees the completion of wake-up within $n-k+1$ rounds. Indeed, for any set X of k stations, at most $n-k$ transmissions can be wasted by `round-robin` (*i.e.* they generate silence) as there are no more than $n-k$ stations in the complement of X. Therefore, in view of Corollary 2.1, `round-robin` is asymptotically optimal for $k > n/c$, where c is any constant.

Algorithm `wakeup_with_s`. Interleave `round-robin` with `wakeup_with_s` described below.

In what follows, we first describe algorithm `select_among_the_first`, then we prove its correctness and finally we derive the complexity of the resulting algorithm obtained by interleaving `round-robin` and `select_among_the_first`.

Algorithm `select_among_the_first`. Only stations awakened in round s will be allowed to participate to the transmissions, while the others remain silent for the whole algorithm's execution (of course, since s is known, each station can locally determine whether to participate or remain silent by simply comparing its own waking time with s). Let X be the set of participating stations. Any station in X transmits according to a sequential composition of schedules defined by the concatenation of $(n, 2^j)$-selective families, for $j = 1, 2, \ldots$, until a successful transmission occurs.

Correctness of algorithm `select_among_the_first`**.** Let i be such that $2^{i-1} \leq |X| \leq 2^i$. To see the correctness of algorithm `select_among_the_first`, it is sufficient to observe that the selectivity property applied on a $(n, 2^i)$-selective family $\mathcal{F}$ guarantees the existence of a time slot t, $1 \leq t \leq |\mathcal{F}| = O(2^i + 2^i \log(n/2^i))$, at which only one station in X will transmit. Since $|X| \leq k$, by the end of the execution of the $(n, 2^{\lceil \log k \rceil})$-selective family a successful transmission must occur.

Time complexity of algorithm `wakeup_with_s`**.** The asymptotic time complexity of the final algorithm `wakeup_with_s` is the minimum between the asymptotic complexities of its two interleaved components. We have already observed that `round-robin` solves the wake-up problem within $n-k+1$ rounds, which is optimal for $k > n/c$, where c is any constant. Hence, to prove the optimality of `wakeup_with_s` it will suffice to show that `select_among_the_first` is optimal for $k \leq n/c$ for some constant $c > 0$.

To this aim, first observe that in [14] a lower bound $\Omega(k \log(n/k))$ is proved for $2 \leq k \leq n/64$, which holds even if k is known and all the stations (which can be at most k) start at the same time. The time complexity of `select_among_the_first` can be calculated as follows. The number of time slots up to the first successful transmission is upper bounded by $O(2+2\log(n/2)+4+4\log(n/4)+\ldots+k+k\log(n/k)) \subseteq O(k+k\log(n/k))$. Since $O(k+k\log(n/k)) \subseteq O(k\log(n/k))$ for $2 \leq k \leq n/64$, `select_among_the_first` is asymptotically optimal for $2 \leq k \leq n/64$.

We can conclude that our interleaved algorithm is asymptotically optimal and has performance $\Theta(\min\{n-k+1, k\log(n/k)+k\} = \Theta(k\log(n/k)+1)$.

4 Contention resolution when k is known

In this section we show that if the stations know the upper bound k on the number of stations that can become active, again a simple algorithm similar to that of the previous section achieves the wake-up in time $\Theta(k \log(n/k) + 1)$.

As in the previous section, we propose an algorithm, that we call `wakeup_with_k`, which is an interleaving between `round-robin` (exploited for large values of k) and a new simple algorithm, called `wait_and_go`.

Algorithm `wakeup_with_k`. Interleave `round-robin` with `wait_and_go` described below.

We start with a description of algorithm `wait_and_go`, then we prove its correctness and finally we derive the complexity of the interleaved algorithm `wakeup_with_k`. Let us start with the definition of the schedule used by `wait_and_go`. Given a set of selective families $\{\mathcal{F}_1, \mathcal{F}_2, \ldots, \mathcal{F}_l\}$, we will denote by $\langle \mathcal{F}_1, \mathcal{F}_2, \ldots, \mathcal{F}_l \rangle$ the schedule defined by the ordered sequence of selective families $\mathcal{F}_1, \mathcal{F}_2, \ldots, \mathcal{F}_l$. Let z_i, for $1 \le i \le \lceil \log k \rceil$, be the length of a $(n, 2^i)$-selective family and let $z = z_1 + z_2 + \ldots + z_{\lceil \log k \rceil}$. Algorithm `wait_and_go` will use the schedule

$$\mathcal{F} = \langle \mathcal{F}_1, \mathcal{F}_2, \ldots, \mathcal{F}_{\lceil \log k \rceil} \rangle = \langle F_1, F_2, \ldots, F_{z_1}, F_{z_1+1}, \ldots, F_{z_1+z_2}, F_{z_1+z_2+1}, \ldots, F_z \rangle$$

defined by the ordered sequence of $(n, 2^i)$-selective families for $i = 1, 2, \ldots, \lceil \log k \rceil$. Each round j ticked by the global clock corresponds to set F_j of the sequence $\mathcal{F}$. A crucial feature of the algorithm (which justifies its name) is that a station that has been activated at an arbitrary round j waits (remains silent) until round $\sigma \ge j$ such that the corresponding set F_σ is the first set of a selective family (in sequence $\mathcal{F}$). Formally, the algorithm can be described as follows.

Algorithm `wait_and_go`. Any station x that becomes active at some time j executes the following protocol. Let $\sigma \ge j$ be the smallest time slot such that set $F_{\sigma \bmod z}$ of $\mathcal{F}$ is the first transmission set of a $(n, 2^i)$-selective family for some $1 \le i \le \lceil \log k \rceil$. Station x remains silent from time j to time $\sigma - 1$. At every time $t \ge \sigma$, station x transmits according to transmission set $F_{t \bmod z}$, *i.e.*, if and only if $x \in F_{t \bmod z}$.

Correctness of algorithm `wait_and_go`. The fact that newly activated stations have to wait until the beginning of the next selective family before they are allowed to transmit, guarantees that the set of stations involved in any selective family of $\mathcal{F}$ remains unchanged during the execution of that selective family. Consequently, if we denote by X_i the set of transmitting stations involved in the execution of the selective family $\mathcal{F}_i = \{F_{z_{i-1}+1}, \ldots, F_{z_{i-1}+z_i}\}$, X_i will be formed by all those stations that became active within the time the algorithm "has reached" the first transmission set of $\mathcal{F}_i$, *i.e.*, not later than round $z_{i-1} + 1$. All the stations that are activated between times $z_{i-1} + 2$ and time z remain silent.

Since $|X_i| \le k$ for every $i = 1, 2, \ldots, \lceil \log k \rceil$, there will be a selective family $\mathcal{F}_i$ in $\mathcal{F}$, for some $1 \le i \le \lceil \log k \rceil$, such that $2^{i-1} \le |X_i| \le 2^i$. The selectivity property on $\mathcal{F}_i$ guarantees that there will be a time slot during the execution of $\mathcal{F}_i$ such that only one station is allowed to transmit.

Time complexity of algorithm `wakeup_with_k`. The time complexity of our algorithm `wait_and_go` can be derived as follows. The number of time slots from s up to the first successful transmission is upper bounded by the length of the ordered sequence of $(n, 2^i)$-selective families for $i = 1, 2, \ldots, \lceil \log k \rceil$. That is: $O(2 + 2\log(n/2) + 4 + 4\log(n/4) + \ldots + k + k\log(n/k)) \subseteq O(k + k\log(n/k))$.

Therefore, the time complexity of the interleaved algorithm `wakeup_with_k` will be, as for algorithm `wakeup_with_s`, $\Theta(\min\{n - k + 1, k + k\log(n/k)\}) = \Theta(k\log(n/k) + 1)$. Again, the optimality of `round-robin` (for $k > n/c$, where c is any positive constant) together with the lower bound in [14] proves the asymptotic optimality of `wakeup_with_k`.

5 Contention resolution without knowledge of s and k

Here we deal with the case when both the first time s at which there is some set of active stations and the upper bound k on the possible active stations are unknown (Scenario C). In this scenario, as in the other two, there is a global time accessed by every station. But, unlike Scenarios A and B, in Scenario C every awakened station does not know the wake-up times of other possible awake stations; nor does it know whether or not there are already awakened stations. Throughout the paper we omit all the floor and ceiling signs as they are not crucial for determining our asymptotic bound. In the following subsection we will present an algorithm achieving the wake-up in $O(k \log n \log\log n)$ time.

High-level description of the algorithm. The algorithm uses a combinatorial tool, the transmission matrix, which contains all the transmission sets used by the stations to decide whether to transmit at a given round or not. Roughly speaking, every row i includes a sequence of $(n, 2^i)$-selective families (as we will see they have some additional property, with respect to the classical selective families, that allows us to use them simultaneously to isolate a station).

The columns of the transmission matrix correspond to rounds, in the sense that during the execution of the algorithm all the stations that are active at time t will transmit according to transmission sets on the same column t of the matrix. Depending on the time at which they have been woken up, they can be involved in different rows of the matrix.

The behaviour of any station can be informally described as follows. Once a station has been woken up, it begins to scan the matrix starting from the column corresponding to the current round. It begins the transmission according to the transmission sets in the first row (corresponding to a $(n, 2)$-selective family). Once all the sets of a $(n, 2)$-selective family are executed, it goes down to the second row, following the transmission sets of a $(n, 4)$-selective family and so on.

5.1 Description of the protocol

Let $m_0 = 0$ and $m_i = c2^i \log n \log\log n$, for $i = 1, 2, \ldots, \log n$ and some sufficiently large constant $c > 0$ to be specified in the analysis. For any $j > 0$, let $\mu(j) = \min\{l \geq j : l \equiv 0 \mod \log\log n\}$. Any station u waking up at a time σ executes the following protocol `wakeup`(u, σ).

The protocol is provided with a $(\log n \times \ell)$ *transmission matrix* $\mathcal{M}$, for $\ell = 2cn \log n \log\log n$, whose entries $M_{i,j}$ are subsets of stations. The parameter ℓ will be called the *length of the transmission matrix*. Each station knows whether it belongs to set $M_{i,j}$, for any pair of admissible parameters i, j; this defines a binary transmission schedule for each of them. Sets $M_{i,j}$ are called *transmission sets*. Later we specify and analyze properties of transmission matrix $\mathcal{M}$ sufficient to guarantee efficiency and correctness of protocol `wakeup`.

Protocol `wakeup`(u, σ)

$t' \leftarrow \mu(\sigma)$
for $i = 1$ to $\log n$ **do**
  for $t = t'$ to $t' + m_i - 1$ **do**
    $j \leftarrow t \mod \ell$ $\{\mathcal{M}$ is scanned in a circular way$\}$
    if $u \in M_{i,j}$ **then**
      send a message at time t
  $t' \leftarrow t' + m_i$

Following the algorithm, we can observe that a station u, woken up at time σ_u, waits for rounds $t \in [\sigma_u, \mu(\sigma_u))$ before becoming operative at time $\mu(\sigma_u)$. Notice that interval $[\sigma_u, \mu(\sigma_u))$ may be empty in case $\sigma_u \equiv 0 \mod \log\log n$.

This implies another crucial property that will be exploited in the proof of Lemma 5.4: all stations woken up in the interval (j_1, j_2), where j_1 and j_2 are two consecutive rounds congruent

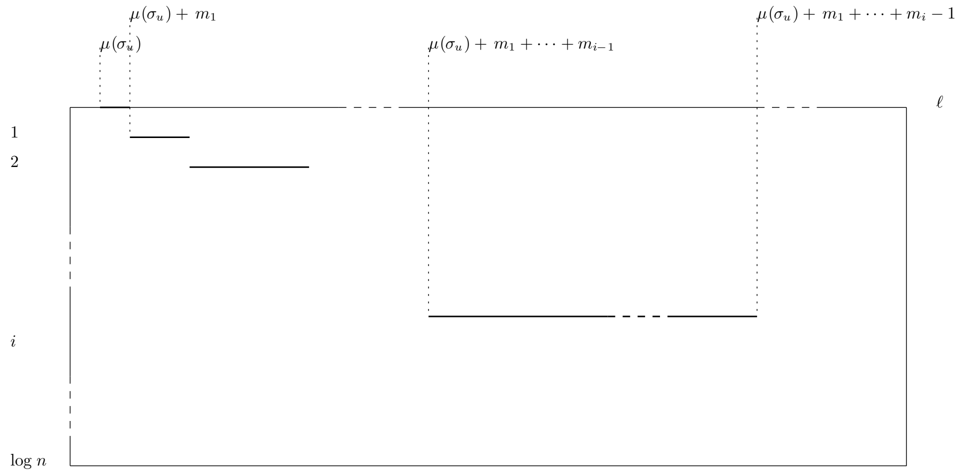

Figure 1: A graphical representation of the transmission sets of a $(\log n \times \ell)$ transmission matrix conditionally to which a station u, waking up at some time σ_u, transmits between time slots $\mu(\sigma_u)$ and $\mu(\sigma_u) + m_1 + \ldots + m_i - 1$, for $1 \le i \le \log n$.

modulo $\log\log n$ (*i.e.* such that $j_1 \equiv j_2 \mod \log\log n$) are not allowed to transmit. We will give a more formal version of such a property in Section 5.2, using the definition of *windows*.

For $t \ge \mu(\sigma_u)$, we say throughout this section, that it transmits *conditionally to a transmission set* $M_{i,j}$, for $j = t \mod \ell$, if the following conditions hold:

- if $u \in M_{i,j}$, it sends its message at round t;
- if $u \notin M_{i,j}$, it stays silent at round t.

Following protocol `wakeup`, any station u waking up at some time σ_u, has the following transmitting behaviour. It waits until time slot $\mu(\sigma_u)$ and then starts transmitting conditionally to transmission sets $M_{1,\mu(\sigma_u) \bmod \ell}, \ldots, M_{1,(\mu(\sigma_u)+m_1-1) \bmod \ell}$ (row 1 of matrix $\mathcal{M}$) for time slots $\mu(\sigma_u), \ldots, \mu(\sigma_u)+m_1-1$, respectively. Then it transmits conditionally to transmission sets $M_{2,(\mu(\sigma_u)+m_1) \bmod \ell}, \ldots, M_{2,(\mu(\sigma_u)+m_1+m_2-1) \bmod \ell}$ (row 2 of matrix $\mathcal{M}$) in time slots $\mu(\sigma_u) + m_1, \ldots, \mu(\sigma_u) + m_1 + m_2 - 1$, respectively. And so on (see Figure 1).

In order to simplify the notation, in the rest of the paper we will avoid specifying the modulo ℓ on the columns of the matrix: it is understood that the matrix is scanned circularly. This is equivalent to consider our matrix as a concatenation of a sufficiently large number of copies of the ℓ-column matrix defined above. This, in particular, allows us to have a direct correspondence between a time slot t and a column t on $\mathcal{M}$. Notice that the number of copies of the matrix, sufficient for the protocol to work, is bounded by the total execution time of the algorithm, that is $O(n \log n \log\log n)$ as we will show in the final theorem of this section.

5.2 Waking matrices

A station u will be hereafter viewed as a couple (u, σ_u), where σ_u is the time at which it is woken up (note, indeed, that, for any fixed execution, this wake up time is uniquely defined for each station's ID). Let $S(j) = \{(u, \sigma_u) : u \in [n] \text{ is woken up at time } \sigma_u \text{ such that } \mu(\sigma_u) \le j\}$ be the set of stations that are operational at time j, *i.e.* that transmit conditionally to transmission sets of $\mathcal{M}$.

Note that each station in $S(j)$, when following protocol `wakeup`, may transmit conditionally to transmission sets in different rows of $\mathcal{M}$, depending on the time at which it was waken up. Given a parameter $1 \le i \le \log n$ and a time slot j, we denote by $S_{i,j}$ the set of stations that at time j transmit conditionally to transmission set $M_{i,j}$. Set $S(j)$ can be partitioned using sets $S_{i,j}$ in the following way:

$$\begin{cases} S(j) = \bigcup_{i=1}^{\log n} S_{i,j} \\ \\ S_{i,j} \cap S_{l,j} = \emptyset \qquad \text{for every } i \neq l \quad . \end{cases}$$

In other words, at every time slot j all stations in $S(j)$ transmit conditionally to transmission sets that may be in different rows of $\mathcal{M}$, but that are *vertically aligned* on $\mathcal{M}$, *i.e.*, they are in the same column j of matrix $\mathcal{M}$ (see Figure 2).

Definition 5.1 *A* window *is a set of* $\log\log n$ *consecutive time slots* $W = \{p\log\log n, p\log\log n + 1, \dots, (p+1)\log\log n - 1\}$ *for some integer* $p \ge 0$.

As we have already observed in the previous subsection, stations that wake up inside a window are not allowed to transmit. A consequence of this fact is the following property:

P1: Let $W = \{p\log\log n, p\log\log n + 1, \dots, (p+1)\log\log n - 1\}$ be a window for some integer $p \ge 0$. Let $A = S_{i,p\log\log n}$. We have:

$$S_{i,j} = A, \text{ for all } j \in W.$$

In other terms, for any window W and for any $1 \le i \le \log n$ we have that $S_{i,j}$ is the same set for all $j \in W$.

In the sequel, we denote by s the first time slot such that $|S(s)| > 0$. A set of stations $S(t)$ is said to be *well-balanced* at time t if there exist $c \cdot |S(t)| \log n \log\log n$ time slots j, where $s \le j \le t$ and c is a sufficiently large integer constant, such that

S1: $\sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^i} \le \log n$.

S2: $|S_{i,j}| \ge 2^{i-3}$, for some $1 \le i \le \log n$.

We call t a *well-balanced round*. As a consequence of Property P1, we can observe that for any round j that satisfies conditions S1 and S2, all rounds j' that belong to the same window of j also satisfy the conditions. In other words, we can state the following property:

P2: for any window W, either conditions S1 and S2 are satisfied for all $j \in W$ or there is no round $j \in W$ that satisfies both of them.

This property allows us to partition the execution of the algorithm into windows that satisfy conditions S1 and S2 in every round of the window and windows that do not satisfy S1 and S2 in any round. Therefore, we can rewrite the definition of *well-balanced* set of stations in terms of windows.

Definition 5.2 *A set of stations* $S(t)$ *is said to be* well-balanced *at time* t *if there exist* $h = c \cdot |S(t)| \log n$ *windows* $W_0, \dots, W_{h-1}$ *such that, for every* $0 \le g \le h-1$ *and every* $j \in W_g$ *(where* $s \le j \le t$ *and* c *is a sufficiently large integer constant) conditions* S1 *and* S2 *hold.*

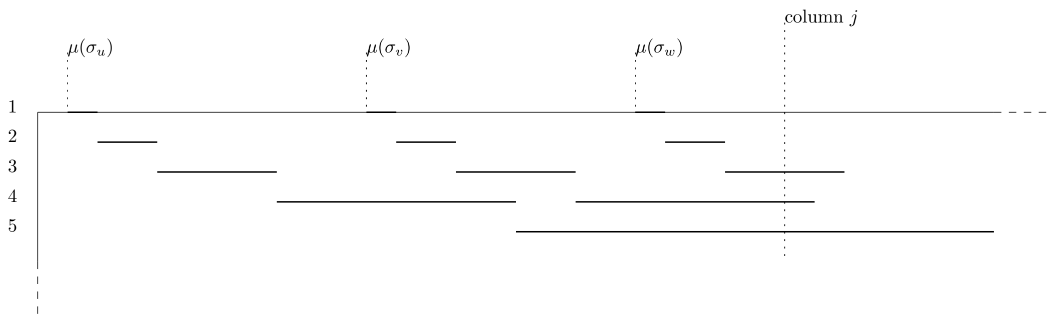

Figure 2: Three stations u, v and w wake up at different time slots: σ_u, σ_v and σ_w. In view of their different wake up times, at time slot j they transmit conditionally to transmission sets located in different rows of $\mathcal{M}$, but the same column j.

The above definition will become crucial in the proof of Lemma 5.4. Finally, we say that a station $w \in S_{i,j} \subseteq S(j)$ is *isolated* at time j if and only if

$$\bigcup_{i=1}^{\log n} (S_{i,j} \cap M_{i,j}) = \{w\}.$$

Definition 5.3 (Waking matrix) *A transmission matrix $\mathcal{M}$ is called a* waking matrix *if and only if for every $s \geq 0$, $1 \leq k \leq n$ and well-balanced set of stations $S(t)$ for round $t \geq s$ such that $|S(t)| \leq k$, there exists a time slot j, with $s \leq j \leq t$, and a station $w \in S(j)$ such that w is isolated at time slot j.*

5.3 Existence of a waking matrix

Let $\rho(j) = j \bmod \log\log n$. Let $\mathcal{M}$ be a randomly constructed transmission matrix of length $\ell = 2cn \log n \log\log n$, specified as follows: for a given $1 \leq i \leq \log n$ and $1 \leq j \leq \ell$, an entry $M_{i,j}$ is formed by letting $\text{Prob}[u \in M_{i,j}] = \frac{1}{2^{i+\rho(j)}}$ for every $u \in [n]$. All decisions on whether $u \in M_{i,j}$ are made *independently*. Our goal is to show that such a random matrix is a waking matrix with a positive probability, and therefore, by the probabilistic method, there is such a matrix.

The proof consists of two parts. First, we show a sufficient condition for $|S(t)|$ such that $S(t)$ is well-balanced in round t (*cf.* Theorem 5.1). Next, we estimate the probability of having a successful transmission within an interval $[s, t]$, for any well-balanced round t (*cf.* Lemmas 5.3 and 5.4) and then we conclude about $\mathcal{M}$ being a waking matrix (*cf.* Theorem 5.2).

Theorem 5.1 *Let t be the smallest time slot such that $t - s \geq 2c \cdot |S(t)| \log n \log\log n$. Then, $S(t)$ is well-balanced at time t.*

The proof of Theorem 5.1 is a straightforward combination of the following two lemmas, corresponding to properties S1 and S2, respectively, of the definition of well-balanced set $S(t)$.

Lemma 5.1 *Let t be the smallest time slot such that $t - s \geq 2c \cdot |S(t)| \log n \log\log n$. Then, there exist $c \cdot |S(t)| \log n \log\log n$ time slots j, with $s \leq j \leq t$, such that*

$$\sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^i} \leq \log n.$$

Proof: Let $H = \{s \le j \le t : \sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^i} > \log n\}$. Following protocol `wakeup`, any station $x \in S(t)$ woken up at time σ transmits conditionally to the transmission sets $M_{i,j}$ in row i of $\mathcal{M}$ for time slots $j = \mu(\sigma)+(m_0+\cdots+m_{i-1}), \ldots, \mu(\sigma)+(m_0+\cdots+m_i)-1$. Therefore, any station in $S(t)$ appears in the sets $S_{i,j}$ of the ith row of matrix $\mathcal{M}$ for not more than $m_i = c2^i \log n \log\log n$ columns, *i.e.*, for all rows i of the matrix, we must have

$$m_i|S(t)| \ge \sum_{j \in H} |S_{i,j}|.$$

Hence,

$$\begin{aligned}
\sum_{i=1}^{\log n} |S(t)| &\ge \sum_{i=1}^{\log n} \frac{\sum_{j\in H} |S_{i,j}|}{m_i} \\
&= \sum_{j\in H} \sum_{i=1}^{\log n} \frac{|S_{i,j}|}{m_i} \\
&= \frac{1}{c \log n \log\log n} \sum_{j\in H} \sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^i} \\
&> \frac{1}{c \log n \log\log n} \sum_{j\in H} \log n \qquad \left(\text{because } \textstyle\sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^i} > \log n \text{ for all } j \in H\right) \\
&= \frac{|H|}{c \log\log n} \ . \qquad (1)
\end{aligned}$$

From Eq. (1) we get $H \le c \log n \log\log n \cdot |S(t)|$. Recalling the assumption $t - s \ge 2c \cdot |S(t)| \log n \log\log n$, the lemma follows.

□

Lemma 5.2 *Suppose that* $j - s < 2c \cdot |S(j)| \log n \log\log n$ *for every* $j < t$. *Then, for every* $s \le j \le t$, *there exists* i, $1 \le i \le \log n$, *such that*

$$|S_{i,j}| \ge 2^{i-3}.$$

Proof: Suppose by contradiction that there exists a j, with $s \le j \le t$, such that $|S_{i,j}| < 2^{i-3}$ for every $i = 1, 2, \ldots, \log n$. Let l be such that $2^{l-1} \le S(j) \le 2^l$. Recalling the assumption $j - s < 2c \cdot |S(j)| \log n \log\log n$, we have $j - s < c2^{l+1} \log n \log\log n$. It follows that there is no station transmitting conditionally to transmission sets $M_{i,j}$ for $i > l+1$, *i.e.* $|S_{i,j}| = 0$ for $i > l+1$. By the contradiction hypothesis it follows that $|S_{i,j}| < 2^{i-3}$ for $i \le l+1$. Consequently, we have

$$\begin{aligned}
|S(j)| &= \sum_{1 \le i \le l+1} |S_{i,j}| \\
&< \sum_{1 \le i \le l+1} 2^{i-3} \\
&< 2^{l-1} \ ,
\end{aligned}$$

which is a contradiction with the definition of l. □

In the next two lemmas we estimate probabilities of isolating a station: first in any round, next in some number of rounds preceding a well-balanced round.

Lemma 5.3 *Let $S(t)$ be any set of stations that are awake at time t. The probability that there exists a station $w \in S(t)$ isolated at an arbitrary time slot $j \le t$ is at least*

$$\sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^{i+\rho(j)}} \left(\frac{1}{4}\right)^{\sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^{i+\rho(j)}}} .$$

Proof: Let $A(i,j)$ be the event "there exists $w \in S_{i,j}$ such that $S_{i,j} \cap M_{i,j} = \{w\}$", and let $B(i,j)$ be the event "for all l with $l \neq i, S_{l,j} \cap M_{l,j} = \emptyset$".

It is evident that a station in $S(t)$ is isolated at time j if and only if the event $\bigcup_{i=1}^{\log n} A(i,j) \cap B(i,j)$ arises.

The probability of $A(i,j)$ is the probability that there exists at least $w \in S_{i,j}$ such that $w \in M_{i,j}$ and $y \notin M_{i,j}$ for every $y \in S_{i,j} \setminus \{w\}$, while the probability of $B(i,j)$ corresponds to the probability that for all $l \neq i$ we have that $y \notin M_{l,j}$ for any $y \in S_{l,j}$. Hence,

$$\begin{aligned} \text{Prob}[A(i,j)] &= \frac{|S_{i,j}|}{2^{i+\rho(j)}} \left(1 - \frac{1}{2^{i+\rho(j)}}\right)^{|S_{i,j}|-1} \ge \frac{|S_{i,j}|}{2^{i+\rho(j)}} \left(1 - \frac{1}{2^{i+\rho(j)}}\right)^{|S_{i,j}|} ; \\ \text{Prob}[B(i,j)] &= \prod_{l=1, l\neq i}^{\log n} \left(1 - \frac{1}{2^{l+\rho(j)}}\right)^{|S_{l,j}|} . \end{aligned}$$

Since $A(i,j)$ and $B(i,j)$ are statistically independent, we can write:

$$\begin{aligned} \text{Prob}[A(i,j) \cap B(i,j)] &\ge \frac{|S_{i,j}|}{2^{i+\rho(j)}} \prod_{i=1}^{\log n} \left(1 - \frac{1}{2^{i+\rho(j)}}\right)^{|S_{i,j}|} \\ &= \frac{|S_{i,j}|}{2^{i+\rho(j)}} \prod_{i=1}^{\log n} \left(1 - \frac{1}{2^{i+\rho(j)}}\right)^{2^{i+\rho(j)} \frac{|S_{i,j}|}{2^{i+\rho(j)}}} \\ &\ge \frac{|S_{i,j}|}{2^{i+\rho(j)}} \left(\frac{1}{4}\right)^{\sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^{i+\rho(j)}}} . \end{aligned}$$

To complete the proof it is sufficient to observe that for any fixed j the events $A(i,j) \cap B(i,j)$, for $i = 1, 2, \ldots, \log n$, are mutually exclusive. $\square$

Lemma 5.4 *Let $S(t)$ be a well-balanced set of stations. There exist $c \cdot |S(t)| \log n$ time slots j, for $s \le j \le t$ such that*

$$\frac{1}{8} \le \sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^{i+\rho(j)}} \le 2 .$$

Proof: Since $S(t)$ is well-balanced, by definition there exist $c \cdot |S(t)| \log n \log\log n$ time slots j, with $s \le j \le t$ such that conditions S1 and S2 hold.

Recall now Definition 5.1 and Definition 5.2. By the hypothesis that $S(t)$ is well-balanced, there are $h = c \cdot |S(t)| \log n$ windows $W_0, \ldots, W_{h-1}$ such that, for every $0 \le g \le h-1$ and every $j \in W_g$ with $s \le j \le t$, the following conditions hold:

(a) $\sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^i} \leq \log n$;

(b) there exists $1 \leq i \leq \log n$ such that $\frac{|S_{i,j}|}{2^i} \geq \frac{1}{8}$.

These two conditions imply

$$\frac{1}{8} \leq \sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^i} \leq \log n. \tag{2}$$

Now fix $0 \leq g \leq h-1$ and consider the corresponding window W_g. By property P1 for any $1 \leq i \leq \log n$ we have that $S_{i,j}$ contains the same set of stations for all $j \in W_g$. This implies that the value of the sum $\sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^i}$ is the same for all $j \in W_g$. Since the value of $\rho(j)$ increases by 1 over successive indices of $j \in W_g$, ranging between a minimum of 0 (corresponding to $j = 0 \mod \log\log n$) and a maximum of $\log\log n - 1$ (for $j = \log\log n - 1 \mod \log\log n$), the sum $\sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^{i+\rho(j)}}$ will decrease by half over consecutive columns j of the same window W_g.

Hence, recalling (2), we have that for every $0 \leq g \leq h-1$, there exists $j \in W_g$ such that

$$\frac{1}{8} \leq \sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^{i+\rho(j)}} \leq 2.$$

□

Theorem 5.2 *There exists a waking matrix of length $O(n \log n \log\log n)$.*

Proof: Consider a sufficiently large integer constant $c > 0$ and a parameter $1 \leq k \leq n$. First observe that we may assume that $s = 0$; otherwise one could simply shift the whole execution to obtain this property, and shifting does not influence the property of isolating a node within t number of rounds after s, for any t. In order to prove the theorem, it is sufficient to show the following property:

> If t^* is the smallest value of t for which $S(t)$ is well-balanced and $|S(t^*)| \leq k$, then there is $w \in S(j)$, for some $j \leq t^*$, such that w is isolated at time j.

Consider a randomly constructed matrix $\mathcal{M}$, as defined in the beginning of Section 5.3: $\text{Prob}[u \in M_{i,j}] = \frac{1}{2^{i+\rho(j)}}$ for every $u \in [n]$, with decisions on whether $u \in M_{i,j}$ made independently. Fix sets $S(t)$, for every non-negative integer t. Let t^* be the smallest value of t for which $S(t)$ is well-balanced. By assumption we know that at most k stations can be woken up and therefore $|S(t^*)| \leq k$. Let us denote $|S(t^*)|$ by x^*. If such a t^* does not exist then the implication in the definition of k-waking matrix is automatically satisfied for the family $\{S(t)\}_{t\geq 0}$. Therefore we can restrict the analysis to families $\{S(t)\}_{t\geq 0}$ for which t^* is well defined. It follows from Theorem 5.1 that $t^* \leq 2ck \log n \log\log n$. Denote $2ck \log n \log\log n$ by λ. We call t^* the first *well-balanced round* of family $\{S(t)\}_{t\geq 0}$ and x^* the sparsity of this family.

By Lemma 5.4, there are at least $c \cdot |S(t^*)| \log n = cx^* \log n$ slots $j \leq t^*$ such that $\frac{1}{8} \leq \sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^{i+\rho(j)}} \leq 2$. Let J be the set of such values j. Applying Lemma 5.3 to each slot $j \in J$ we get the following: the probability that there exists a station $w \in S(j)$ isolated at time slot $j \leq t^*$ is at least

$$\sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^{i+\rho(j)}} \left(\frac{1}{4}\right)^{\sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^{i+\rho(j)}}} \geq \sum_{i=1}^{\log n} \frac{|S_{i,j}|}{2^{i+\rho(j)}} \left(\frac{1}{4}\right)^2 \geq \frac{1}{128} .$$

Therefore, the probability that in none of the time slots in J there is an isolated station is at most $\left(\frac{127}{128}\right)^{|J|}$, which is, by the lower bound on $|J|$, smaller or equal to

$$\left(\frac{127}{128}\right)^{cx^* \log n} = n^{-cx^* \log(128/127)} \ . \tag{3}$$

The same upper bound automatically holds for the event that in none of the time slots in $[0, t^*]$ there is any isolated station.

Fix positive integers $t^* \leq 2ck \log n \log \log n$ and $x^* \leq k$. Now we estimate the number of possible families of sets $\{S(t)\}_{t \geq 0}$ with the first well-balanced round t^* and sparsity x^*. This family is monotonically non-decreasing, in the sense of set inclusion, therefore each such family can be uniquely encoded by a subset of n of size x^* corresponding to set $S(t^*)$ and a function from $[x^*]$ to $[t^* + 1]$ corresponding to the $t^* + 1$ possible wakeup times for each of the stations in $S(t^*)$ during the interval $[0, t^*]$. Therefore, the number of such configurations corresponds to the number of all the possible subsets of x^* elements out of n possible choices combined with all the possible $t^* + 1$ wake-up times for each of the x^* stations. Recalling that $\binom{n}{y} \leq (\frac{ne}{y})^y$, for any $1 \leq y \leq n$, the number of configurations can be upper bounded as follows:

$$\begin{aligned} \binom{n}{x^*}(t^*+1)^{x^*} &\leq \left(\frac{ne}{x^*}\right)^{x^*} (cn^3)^{x^*} \ \text{(since } t^* + 1 \leq 2ck \log n \log \log n \leq cn^3\text{)} \\ &= \left(\frac{1}{x^*}\right)^{x^*} (cen^4)^{x^*} \leq (3cn^4)^{x^*}. \end{aligned} \tag{4}$$

In order to get the final result, we apply a union bound over all possible different values of parameter x^*, parameter t^*, and all possible correspondent different families $\{S(t)\}_{t=0}^{t^*}$ with the first well-balanced round t^* and sparsity x^* (with respect to the events that no station is isolated in the time interval $[0, t^*]$). By using n as an upper bound on the possible range of parameter x^*, and value $3cn^3 > cn^3$ as an upper bound on the possible range of parameter t^*, we get that the probability that there is a family $\{S(t)\}_{t \geq 0}$ with no isolated station before its first well-balanced round of sparsity at most k is smaller than

$$\begin{aligned} n \cdot 3cn^3 \cdot (3cn^4)^{x^*} \cdot n^{-cx^* \log(128/127)} &= (3cn^4)^{x^*+1} \cdot n^{-cx^* \log(128/127)} \\ &\leq n^{(4+\log(3c)) \cdot (x^*+1)} \cdot n^{-cx^* \log(128/127)} \\ &\leq n^{x^*(2(4+\log(3c)) - c \log(128/127))} \\ &< 1 \ \text{(for sufficiently large } c\text{)}, \end{aligned}$$

where in the above estimates we used the upper bound on the number of schedules obtained in (4) and the probability that for a fixed such family there is no isolated station in the time interval $[0, t^*]$ obtained in (3). By the probabilistic method, there must be a deterministic matrix $\mathcal{M}$ such that for any family $\{S(t)\}_{t \geq 0}$ it isolates some station by the first well-balanced round of this family. This completes the proof. □

5.4 Correctness of the algorithm

We are ready to show that if every station executes protocol `wakeup`, then the wake-up problem is solved in at most $O(k \log n \log \log n)$ time slots, where k is the number of awaken stations.

Algorithm `wakeup`(n):
Every station is provided with a $(\log n \times \ell)$ *waking matrix* $\mathcal{M}$*, whose existence is guaranteed by Theorem 5.2. Every station* u*, woken up at time* σ*, executes protocol* `wakeup`(u, σ).

Theorem 5.3 *Algorithm* `wakeup`(n) *solves the wake-up problem under the dynamic scenario by using at most* $O(k \log n \log \log n)$ *time slots, where* k *is the number of stations awaken during the execution.*

Proof: Let s be the first time slot such that $S(s) > 0$. According to algorithm `wakeup`(n), each station u woken up at time $\sigma \geq s$ executes protocol `wakeup`(u, σ). First we can notice that there will be a time slot t, $s < t \leq 2ck \log n \log \log n$, such that $S(t)$ is well-balanced. Indeed, in view of the assumption that $|S(t)| \leq k$ for all t, eventually for $t = s + 2ck \log n \log \log n$ (if not before) we must have $t - s \geq 2c \cdot |S(t)| \log n \log \log n$ which, by Theorem 5.1, implies that $S(t)$ is well-balanced. Let $t \leq s + 2ck \log n \log \log n$ be the smallest time slot such that $S(t)$ is well-balanced. Since all stations following the algorithm transmit according to the waking matrix with properties guaranteed by Theorem 5.2, there exists a station $w \in S(t)$ isolated at some time $j \leq t$. This completes the proof. □

6 A note on randomized solutions

Kushilevitz and Mansour [28] showed that, when a global clock is available, if m stations wake up simultaneously, for $m \in \{2^0, 2^1, \ldots, 2^{\log k}\}$, then the expected number of time slots until the first successful transmission is $\Omega(\log k)$ for any randomized protocol. This lower bound is independent of the time at which the stations wake up, *i.e.*, the knowledge of s does not play any role in the proof, and holds even in the case when k is known. Therefore, it holds for all three scenarios considered in our work.

As for the upper bounds, Jurdzinski and Stachowiak [24] considered the case when each of the n stations can spontaneously and independently wake up at any moment. For the case when a global clock is available and n is known, they designed the following algorithm, called *Repeated Probability Decrease* (RPD). Let $\ell = 2\lceil \log n \rceil$. Each station, starting from the time it wakes up, transmits in round σ with probability $2^{-1-\sigma \bmod \ell}$. The authors showed that this algorithm accomplishes the wakeup in $O(\log n)$ expected time. It implies this same upper bound for all three scenarios considered in the present paper. However, in scenario B, when k is known, an optimum result could be obtained by choosing $\ell = 2\lceil \log k \rceil$ in the algorithm RPD; mainly, algorithm RPD with such value of ℓ gives an expected time complexity of $O(\log k)$, so matching the lower bound by Kushilevitz and Mansour [28].

However, in scenarios A and C the gap between the upper bound $O(\log n)$ and the lower bound $\Omega(\log k)$ on expected wake-up time remains still open.

It is also worth noting that our deterministic non-explicitly constructive solution is actually based on de-randomization of some randomized algorithm; more precisely, an algorithm instantiated by a random waking matrix. Our analysis shows that with probability exponentially close to 1 this randomized algorithm achieves wake-up in $O(k \log n \log \log n)$ rounds. The time performance is far from the optimal for randomized solutions, but the probability is exponentially close to 1, which is not the case in faster Monte Carlo solutions.

7 Conclusions

In this work we provided deterministic solutions to the wake-up problem on a multiple-access channel with non-synchronized awakening and global clock, under three slightly different scenarios. In two of them, the solutions are asymptotically optimal, while in the hardest of the settings there is only an $O(\log \log n)$ factor away from the best known lower bound. Closing this gap is the first remaining open problem. The second natural twist is to provide an efficient implementation of our protocol. This could require an explicit construction of our *waking matrices* or even an entirely new (constructive) solution.

Another interesting open question is whether global clock helps in the wake-up task — the best deterministic solution without global clock is nearly logarithmically worse than the performance of our algorithm (in scenario C), and we conjecture that this gap cannot be removed. Finally, as noted in the last section, there are still remaining gaps for randomized algorithms.

References

[1] N. Abramson, The Aloha system — Another alternative for computer communications, In *Proceedings, Fall Joint Computer Conf., AFIPS Conf* , 1970, vol. 37.

[2] D. Bertsekas and R. Gallager, *Data Networks.* Prentice Hall, 1991.

[3] M. Bienkowski, M. Klonowski, M. Korzeniowski, and D.R. Kowalski, Dynamic sharing of a multiple access channel, in *Proceedings, 27th International Symposium on Theoretical Aspects of Computer Science (STACS)*, 2010 : 83-94.

[4] J. Capetanakis, Tree algorithms for packet broadcast channels. *IEEE Transactions on Information Theory* 25 (1979) 505 - 515.

[5] B. S. Chlebus, G. De Marco, M. Talo, Naming a Channel with Beeps, *Fundamenta Informaticae*, 2017, to appear.

[6] B. S. Chlebus, G. De Marco, D. R. Kowalski, Scalable wake-up of multi-channel single-hop radio networks, *Theoretical Computer Science*, vol. 615, pp. 23-44, 2016.

[7] B. S. Chlebus, G. De Marco, D. R. Kowalski, Scalable Wake-up of Multi-channel Single-Hop Radio Networks in *Proceedings, 18th International Conference on Principles of Distributed Systems (OPODIS 2014)*, Cortina d'Ampezzo, Italy, December 16-19, 2014.

[8] B.S. Chlebus, L. Gasieniec, A. Gibbons, A. Pelc, and W. Rytter, Deterministic broadcasting in unknown radio networks, *Distributed Computing* 15 (2002) 27-38.

[9] B.S. Chlebus, L. Gasieniec, D. Kowalski, and T. Radzik, On the wake-up problem in radio networks, in *Proceedings, 32nd International Colloquium on Automata, Languages and Programming (ICALP)*, 2005, pp. 347 - 359.

[10] B.S. Chlebus and D. Kowalski, A better wake-up in radio networks, in *Proceedings, 23rd ACM Symp. on Principles of Distributed Computing (PODC)*, 2004, 266 - 274.

[11] B.S. Chlebus, D.R. Kowalski, and T. Radzik, Many-to-many communication in radio networks, *Algorithmica* 54 (2009) 118 - 139.

[12] M. Chrobak, L. Gasieniec, and D.R. Kowalski, The wake-up problem in multihop radio networks, *SIAM Journal on Computing* 36 (2007) 1453 - 1471.

[13] M. Chrobak, L. Gasieniec, and W. Rytter, Fast broadcasting and gossiping in radio networks, *Journal of Algorithms* 43 (2002) 177 - 189.

[14] A.E.F. Clementi, A. Monti, and R. Silvestri, Distributed broadcast in radio networks of unknown topology, *Theoretical Computer Science* 302 (2003) 337 - 364.

[15] G. De Marco, Distributed Broadcast in Unknown Radio Networks, *19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2008)*, San Francisco, California, USA, January 2008.

[16] G. De Marco, Distributed Broadcast in Unknown Radio Networks. *SIAM J. Comput.* 39(6): 2162-2175 (2010)

[17] G. De Marco, D. Kowalski, Contention Resolution in a Non-synchronized Multiple Access Channel, *27th Annual IEEE International Symposium on Parallel and Distributed Processing, (IPDPS 2013)* Cambridge, MA, USA, May 2013.

[18] G. De Marco, D. R. Kowalski, Searching for a Subset of Counterfeit Coins: Randomization vs Determinism and Adaptiveness vs Non-Adaptiveness, *Random Structures and Algorithms,* vol. 42 (1), pp. 97-109, 2013.

[19] G. De Marco, D. R. Kowalski, Towards Power-Sensitive Communication on a Multiple-Access Channel. *30th International Conference on Distributed Computing Systems (ICDCS 2010)*, Genoa, Italy, May 2010.

[20] G. De Marco, M. Pellegrini, G. Sburlati, Faster deterministic wakeup in multiple access channels. *Discrete Applied Mathematics* 155(8): 898-903 (2007)

[21] R.G. Gallager, A perspective on multiaccess channels, *IEEE Transactions on Information Theory* 31 (1985) 124 - 142.

[22] L. Gasieniec, A. Pelc, and D. Peleg, The wakeup problem in synchronous broadcast systems, *SIAM Journal on Discrete Mathematics* 14 (2001) 207 - 222.

[23] A.G. Greenberg and S. Winograd, A lower bound on the time needed in the worst case to resolve conflicts deterministically in multiple access channels, *Journal of ACM* 32 (1985) 589 - 596.

[24] T. Jurdzinski and G. Stachowiak, Probabilistic algorithms for the wakeup problem in single-hop radio networks, *Theory Comput. Syst.*, vol. 38, (3), (2005), pp. 347–367.

[25] J. Komlós and A.G. Greenberg, *An Asymptotically Optimal Nonadaptive Algorithm for Conflict Resolution in Multiple-Access Channels*, IEEE Trans. on Information Theory, vol. 31, (2), (1985), pp. 302 - 306

[26] W.H. Kautz and R.R.C. Singleton, Nonrandom binary superimposed codes, *IEEE Transactions on Information Theory* 10 (1964) 363 - 377.

[27] D. Kowalski, On selection problem in radio networks, in *Proceedings, 24th ACM Symposium on Principles of Distributed Computing (PODC),* 2005, pp. 158 - 166.

[28] E. Kushilevitz and Y. Mansour, An $\Omega(D \log(N/D))$ lower bound for broadcast in radio networks, *SIAM Journal on Computing* 27 (1998) 702 - 712.