

Contention Resolution in Multiple-Access Channels: *k*-Selection in Radio Networks*

Antonio Fernández Anta¹ and Miguel A. Mosteiro^{1,2}

¹ LADyR, GSyC, Universidad Rey Juan Carlos, 28933 Móstoles, Madrid, Spain
anto@gsyc.es

² Department of Computer Science, Rutgers University, Piscataway, NJ 08854, USA
mosteiro@cs.rutgers.edu

Abstract. In this paper, contention resolution among k contenders on a multiple-access channel is explored. The problem studied has been modeled as a k -Selection in Radio Networks, in which every contender has to have exclusive access at least once to a shared communication channel. The randomized adaptive protocol presented shows that, for a probability of error 2ε , all the contenders get access to the channel in time $(e + 1 + \xi)k + O(\log^2(1/\varepsilon))$, where $\varepsilon \leq 1/(n + 1)$, $\xi > 0$ is any constant arbitrarily close to 0, and n is the total number of potential contenders. The above time complexity is asymptotically optimal for any significant ε . The protocol works even if the number of contenders k is unknown and collisions can not be detected.

1 Introduction

A recurrent question, in settings where a resource must be shared among many contenders, is how to make that resource available to all of them. The problem is particularly challenging if not even the number of contenders is known. The broad spectrum of settings where answers to such a question are useful makes its study a fundamental task. An example of such contention is the problem of broadcasting information in a multiple-access channel. A multiple-access channel is a synchronous system that allows a message to be delivered to many recipients at the same time using a channel of communication but, due to the shared nature of the channel, the simultaneous introduction of messages from multiple sources produce a conflict that precludes any message from being delivered to any recipient. In Radio Networks,³ one of the instances of such a question is the problem known in the literature [3] as *Selection*. In its general version, the k -*Selection* problem [11], also known as *all-broadcast*, is solved when an unknown size- k subset of n network nodes have been able to access a unique shared channel of communication, each of them at least once. The k -Selection problem in Radio Networks and related problems have been well-studied for settings where a tight upper bound on k is known. In this paper, a randomized adaptive protocol for k -Selection in Radio Networks is presented, assuming that such a knowledge is not available, the arrival of messages is batched, and conflicts to access the channel cannot be detected by all nodes. To our knowledge, this is the first k -Selection protocol in the Radio Networks literature that works in such conditions and it is asymptotically

* This research was supported in part by Comunidad de Madrid grant S-0505/TIC/0285; Spanish MEC grant TIN2005-09198-C02-01, TIN2008-06735-C02-01; NSF grant CCF 0632838; and EU Marie Curie International Reintegration Grant IRG 210021.

³ As pointed out in [3], the historical developments justify the use of Radio Network to refer to any communication network where the channel is contended, even if radio communication is not actually used.

optimal for any sensible error-probability bound (up to inverse exponential in k). This protocol improves over previous work in adversarial packet contention-resolution thanks to the adaptive nature of the protocol and the knowledge of n . Given that the error probability is parametrized, this protocol can be also applied to solve k -Selection in multiple neighborhoods of a multi-hop Radio Network.

Notation and Model Most of the following assumptions and notation are folklore in the Radio Networks literature. For details and motivation, see the survey of Chlebus [3]. We study the k -Selection problem in a Radio Network comprised of n labeled stations called *nodes*. Each node is assumed to be potentially reachable from any other node in one communication step, hence, the network is characterized as *single-hop* or *one-hop* indistinctively. Before running the protocol, nodes have no information besides n and their own label, which is assumed to be unique but arbitrary.⁴ Time is supposed to be slotted in *communication steps*. Assuming that the computation time-cost is negligible in comparison with the communication time-cost, time efficiency is studied in terms of communication steps only. The piece of information assigned to a node in order to deliver it to other nodes is called a *message*. The assignment of a message is due to an external agent and such an event is called a *message arrival*. Communication among nodes is carried out by means of radio broadcast on a shared channel. If exactly one node transmits at a communication step, such a transmission is called *successful* or *non-colliding*, we say that the message was *delivered*, and all other nodes *receive* such a message. If more than one message is transmitted at the same time, a *collision* occurs, the messages are garbled, and nodes only receive *interference noise*. If no message is transmitted in a communication step, nodes receive only *background noise*. In this work, nodes can not distinguish between interference noise and background noise, thus, the channel is called *without collision detection*. Each node is in one of two states, *active* if it holds a message to deliver, or *idle* otherwise. In contrast with *oblivious* protocols, where the sequence of transmissions of a node does not depend on the transmissions received, the *adaptive* protocol presented in this paper exploits the information implicit on the occurrence of a successful transmission. In the randomized protocol presented here all active nodes use the same probability in the same communication step, a class of protocols usually called *fair*. Therefore, it is also a *uniform* protocol, i.e., all active nodes use the same protocol. As in for instance [1, 7, 11], we assume that all the k messages arrive in a *batch*, i.e. in the same communication step, a problem usually called *static* k -Selection,⁵ and that each node becomes idle upon delivering its message.

Problem Definition Given a Radio Network where a subset K of the set of n network nodes, such that $|K| = k$, are activated by message arrivals, the k -Selection problem is solved when each node in K has delivered its message. The definition given pertains to the general version of the problem where messages may arrive at different times, although in this paper we study only simultaneous, or *batched*, arrivals.

Related Work Regarding deterministic solutions, the k -Selection problem was shown to be in $O(k \log(n/k))$ already in the 70's by giving adaptive protocols that make use of collision detection [2, 8, 14]. In all these results the algorithmic technique, known as *tree algorithms*,

⁴ Notice that our protocol does not make any use of the identity of a message originator. Thus, it can be used even in settings where nodes are not labeled or labels are not unique.

⁵ A *dynamic* counterpart where messages arrive at different times was also studied [11].

relies on modeling the protocol as a complete binary tree where the messages are placed at the leaves. Later, Greenberg and Winograd [6] showed a lower bound for that class of protocols of $\Omega(k \log_k n)$. Regarding oblivious algorithms, Komlòs and Greenberg [10] showed the existence of $O(k \log(n/k))$ solutions even without collision detection but requiring knowledge of k and n . More recently, Clementi, Monti, and Silvestri [4] showed a lower bound of $\Omega(k \log(n/k))$, which also holds for adaptive algorithms if collision detection is not available. In [11], Kowalski presented the construction of an oblivious deterministic protocol that, using the explicit selectors of Indyk [9], gives a $O(k \text{ polylog } n)$ upper bound without collision detection.

In the following results, availability of collision detection is assumed. Martel presented in [13] a randomized adaptive protocol for k -Selection that works in $O(k + \log n)$ time in expectation⁶. As argued by Kowalski in [11], this protocol can be improved to $O(k + \log \log n)$ in expectation using Willard's expected $O(\log \log n)$ selection protocol of [17]. In the same paper, Willard shows that, for any given protocol, there exists a choice of $k \leq n$ such that selection takes $\Omega(\log \log n)$ expected time for the class of fair selection protocols. For the case in which n is not known, in the same paper a $O(\log \log k)$ expected time selection protocol is described, again, making use of collision detection. If collision detection is not available, using the techniques of Kushilevitz and Mansour in [12], it can be shown that, for any given protocol, there exists a choice of $k \leq n$ such that $\Omega(\log n)$ is a lower bound in the expected time to get even the first message delivered.

A frequent challenging difficulty to overcome in resolving collisions is to determine which is the best probability of transmission to be used by the contenders when their number is unknown. The method of choice is then to increase or decrease such probability based on the success or failure of successive trials. When the probability of transmission is increased it is said that a *back-on* strategy is used, whereas *back-off* is the term used when such probability is decreased. A combination of both strategies is usually called *back-on/back-off*. Monotonic back-off strategies for contention resolution of batched arrivals of k packets on simple multiple access channels, a problem that can be seen as k -Selection, have been analyzed in [1]. The best strategy shown is the so-called *loglog-iterated back-off* with a makespan in $\Theta(k \log \log k / \log \log \log k)$ with probability at least $1 - 1/k^c$, $c > 0$, which does not use any knowledge of k or n .

Regarding related problems, extending previous work on tree algorithms, Greenberg and Leiserson [7] presented randomized routing strategies in fat-trees for bounded number of messages. Choosing appropriate constant capacities for the edges of the fat-tree, the problem could be seen as k -Selection. However, that choice implies a logarithmic congestion parameter which yields an overall $O(k \text{ polylog } n)$ time. In [5], Gerèb-Graus and Tsantilas presented an algorithm that solves the problem of realizing arbitrary h -relations in an n -node network, with probability at least $1 - 1/n^c$, $c > 0$, in $\Theta(h + \log n \log \log n)$ steps. In an h -relation, each processor is the source as well as the destination of h messages. Making $h = k$ this protocol can be used to solve k -Selection. However, it requires that nodes know h .

Results and Outline In this paper, a randomized adaptive protocol for k -Selection, in a one-hop Radio Network without collision detection, that does not require knowledge of the number of contenders k , is presented. Assuming that $\varepsilon \leq 1/(n + 1)$, the protocol is shown to solve the problem in $(e + 1 + \xi)k + O(\log^2(1/\varepsilon))$ communication steps, where $\xi > 0$ is any constant arbitrarily close to 0 with probability at least $1 - 2\varepsilon$. Given that the error probability

⁶ Throughout this paper, \log means \log_2 unless otherwise stated.

is parametric, this protocol can be applied to multiple neighborhoods of a multi-hop Radio Network, adjusting the error probability in each one-hop neighborhood appropriately. To our knowledge, $O(k \log \log k / \log \log \log k)$ [1] is the best upper bound in the literature for a protocol suitable to solve k -Selection in Radio Networks (although they propose it for packet contention resolution), that works without knowledge of k , under batched arrivals, and without collision detection. By exploiting back-on/back-off and the knowledge of n , our protocol improves their time complexity. Given that k is a lower bound for this problem, the protocol is optimal (modulo a small constant factor) if $\varepsilon \in \Omega(2^{-\sqrt{k}})$. In Section 2 the details of the protocol are presented and they are analyzed in Section 3.

2 Protocol

The protocol comprises two different algorithms. Each of them is particularly suited for one of two scenarios, depending on the number of messages left to deliver. The algorithm called *BT* solves the problem for the case when that number is below a threshold (that will be defined later). The algorithm called *AT* is suited to reduce that number from the initial k to a value below that threshold. The BT algorithm uses the well-known technique of repeating transmissions with the same appropriately-suited probability until the problem is solved. The AT algorithm on the other hand is adaptive by repeatedly increasing an estimation of the messages left and decreasing such an estimation by roughly one each time a message is delivered. (Even if that successful transmission is due to the BT algorithm.) An illustration of the estimation progress is depicted in Figure 1. Further details can be seen in Algorithm 1. Both algorithms




Fig. 1. Illustration of estimate progress.

are executed interleaving their communication steps (see Task 1 in Algorithm 1). For clarity, each communication step is referred to by using the name of the algorithm executed at that step. The following notation used in the algorithm is defined for clarity: $\beta \triangleq e + \xi_\beta$, $\delta \triangleq 1 + \xi_\delta$, $\tau \triangleq 300\beta \ln(1/\varepsilon)$, $\varepsilon \triangleq$ error probability, $0 < \xi_\delta < 1$, $0 < \xi_\beta < 0.27$ and $0 < \xi_t \leq 1/2$ are constants arbitrarily close to 0, and $1/\xi_t \in \mathbb{N}$.

3 Analysis

For clarity, each of the algorithms comprising the protocol are first analyzed separately and later put together in the main theorem. Consider first the AT algorithm. (Refer to Algorithm 1.) Let $\tilde{\kappa}$ be called the *density estimator*. Let a *round* be the sequence of AT-steps between increasings of the density estimator (Line 14). Let the rounds be numbered as $r \in \{1, 2, \dots\}$

Algorithm 1: Pseudocode for node x .

```
1 upon message arrival do
2    $t \leftarrow \tau$ 
3    $\tilde{\kappa} \leftarrow \tau$ 
4   start tasks 1, 2 and 3
5 Task 1
6   foreach communication-step = 1, 2, ... do
7     if communication-step  $\equiv 1 \pmod{1/\xi_t}$  then // BT-step
8       transmit  $\langle x, \text{message} \rangle$  with prob  $1/\tau$ 
9     else // AT-step
10      transmit  $\langle x, \text{message} \rangle$  with prob  $1/\tilde{\kappa}$ 
11       $t \leftarrow t - 1$ 
12      if  $t \leq 0$  then
13         $t \leftarrow \tau$ 
14         $\tilde{\kappa} \leftarrow \tilde{\kappa} + \tau$ 
15 Task 2
16 upon reception from other node do
17    $\tilde{\kappa} \leftarrow \max\{\tilde{\kappa} - \delta, \tau\}$ 
18    $t \leftarrow t + \beta$ 
19 Task 3
20 upon message delivery stop
```

and the AT-steps within a round as $t \in \{1, 2, \dots\}$. (E.g., round 1 is the sequence of AT-steps from initialization until Line 14 of the algorithm is executed for the first time.) Let $\kappa_{r,t}$, called the *density*, be the number of messages not delivered yet (i.e., the number of active nodes) at the beginning of AT-step t of round r . Let $\tilde{\kappa}_{r,t}$ be the density estimator used at the AT-step t of round r . Let $X_{r,t}$ be an indicator random variable such that, $X_{r,t} = 1$ if a message is delivered at the AT-step t of round r , and $X_{r,t} = 0$ otherwise. Then, $Pr(X_{r,t} = 1) = (\kappa_{r,t}/\tilde{\kappa}_{r,t})(1 - 1/\tilde{\kappa}_{r,t})^{\kappa_{r,t}-1}$. Also, for a round r , let the number of messages delivered in the interval of AT-steps $[1, t]$ of r , including those delivered in BT steps, be $\sigma_{r,t}$. The following intermediate results will be useful. First, we state the following useful fact.

Fact 1 [15, §2.68]

$$e^{x/(1+x)} \leq 1 + x \leq e^x, 0 < |x| < 1.$$

Lemma 1. For any round r where $\tilde{\kappa}_{r,1} \leq \kappa_{r,1} - \gamma$, $\gamma \geq \delta(2 - \delta)/(\delta - 1) \geq 0$, $Pr(X_{r,t} = 1)$ is monotonically non-increasing with respect to t for $\delta + 1 < \tilde{\kappa}_{r,t} \leq \kappa_{r,t}$, and $\delta < (\kappa_{r,t} - \gamma)(\kappa_{r,t} - \gamma - 1)/(\kappa_{r,t} - \gamma + 1)$.

Proof. We want to show conditions such that for any t in round r , $Pr(X_{r,t} = 1) \geq Pr(X_{r,t+1} = 1)$. If $\kappa_{r,t} = \kappa_{r,t+1}$ the claim holds trivially. Then, let us assume instead that $\kappa_{r,t} > \kappa_{r,t+1}$. We want to show that

$$\frac{\kappa_{r,t}}{\tilde{\kappa}_{r,t}} \left(1 - \frac{1}{\tilde{\kappa}_{r,t}}\right)^{\kappa_{r,t}-1} \geq \frac{\kappa_{r,t+1}}{\tilde{\kappa}_{r,t+1}} \left(1 - \frac{1}{\tilde{\kappa}_{r,t+1}}\right)^{\kappa_{r,t+1}-1}.$$

Due to the BT-step between two consecutive AT-steps, at most two messages are delivered in the interval $[t, t + 1]$ of r . Thus, replacing appropriately, we want to show that the following

inequalities hold.

$$\frac{\kappa_{r,t}}{\tilde{\kappa}_{r,t}} \left(1 - \frac{1}{\tilde{\kappa}_{r,t}}\right)^{\kappa_{r,t}-1} \geq \frac{\kappa_{r,t}-1}{\tilde{\kappa}_{r,t}-\delta} \left(1 - \frac{1}{\tilde{\kappa}_{r,t}-\delta}\right)^{\kappa_{r,t}-2}, \quad (1)$$

$$\frac{\kappa_{r,t}}{\tilde{\kappa}_{r,t}} \left(1 - \frac{1}{\tilde{\kappa}_{r,t}}\right)^{\kappa_{r,t}-1} \geq \frac{\kappa_{r,t}-2}{\tilde{\kappa}_{r,t}-2\delta} \left(1 - \frac{1}{\tilde{\kappa}_{r,t}-2\delta}\right)^{\kappa_{r,t}-3}. \quad (2)$$

Reordering 1,

$$\frac{\tilde{\kappa}_{r,t}-\delta-1}{\tilde{\kappa}_{r,t}} \left(\frac{\tilde{\kappa}_{r,t}-1}{\tilde{\kappa}_{r,t}} \frac{\tilde{\kappa}_{r,t}-\delta}{\tilde{\kappa}_{r,t}-\delta-1}\right)^{\kappa_{r,t}-1} \geq \frac{\kappa_{r,t}-1}{\kappa_{r,t}}. \quad (3)$$

Using calculus, it can be seen that the left-hand side of 3 is monotonically non-increasing for $\delta+1 < \tilde{\kappa}_{r,t} \leq \kappa_{r,t}$. The details are omitted for brevity. Then, given that $\tilde{\kappa}_{r,t} = \tilde{\kappa}_{r,1} - \sigma_{r,t} \leq \kappa_{r,1} - \sigma_{r,t} - \gamma \leq \kappa_{r,t} - \gamma$, it is enough to show

$$\frac{\kappa_{r,t}}{\kappa_{r,t}-1} \cdot \frac{\kappa_{r,t}-\gamma-\delta-1}{\kappa_{r,t}-\gamma} \cdot \left(\frac{\kappa_{r,t}-\gamma-1}{\kappa_{r,t}-\gamma} \frac{\kappa_{r,t}-\gamma-\delta}{\kappa_{r,t}-\gamma-\delta-1}\right)^{\kappa_{r,t}-1} \geq 1. \quad (4)$$

Again using calculus, it can be seen that the left-hand side of Inequality 4 is monotonically non-increasing on $\kappa_{r,t}$ for $\gamma \geq \delta(2-\delta)/(\delta-1)$ and $\delta < (\kappa_{r,t}-\gamma)(\kappa_{r,t}-\gamma-1)/(\kappa_{r,t}-\gamma+1)$. The details are omitted for brevity. Then, it is enough to show that, in the limit, the left-hand side of Inequality 4 tends to 1. Which can be verified using standard calculus techniques. The details are omitted for brevity. Using the same techniques, Inequality 2 can be shown to hold.

Lemma 2. For any round r where $\kappa_{r,1} - \gamma - \tau \leq \tilde{\kappa}_{r,1} < \kappa_{r,1} - \gamma$, $\gamma \geq 0$ and for any AT-step t in r such that

$$\sigma_{r,t} \leq \kappa_{r,1} \frac{\ln \beta - 1}{\delta \ln \beta - 1} - \frac{(\gamma + \tau + 1) \ln \beta - 1}{\delta \ln \beta - 1},$$

the probability of a successful transmission is at least $\Pr(X_{r,t} = 1) \geq 1/\beta$.

Proof. We want to show $(\kappa_{r,t}/\tilde{\kappa}_{r,t})(1-1/\tilde{\kappa}_{r,t})^{\kappa_{r,t}-1} \geq 1/\beta$. Given that nodes are active until their message is delivered, it is enough to show

$$\frac{\kappa_{r,1}-\sigma_{r,t}}{\tilde{\kappa}_{r,1}-\delta\sigma_{r,t}} \left(1 - \frac{1}{\tilde{\kappa}_{r,1}-\delta\sigma_{r,t}}\right)^{\kappa_{r,1}-1-\sigma_{r,t}} \geq 1/\beta. \quad (5)$$

Using calculus, it can be seen that the left hand side of Inequality 5 is monotonically non-decreasing with respect to $\tilde{\kappa}_{r,1}$ under the conditions of the Lemma. The details are omitted for brevity. Then, it is enough to prove Inequality 5 for $\tilde{\kappa}_{r,1} = \kappa_{r,1} - \gamma - \tau$.

$$\begin{aligned} \frac{\kappa_{r,1}-\sigma_{r,t}}{\kappa_{r,1}-\gamma-\tau-\delta\sigma_{r,t}} \cdot \left(1 - \frac{1}{\kappa_{r,1}-\gamma-\tau-\delta\sigma_{r,t}}\right)^{\kappa_{r,1}-1-\sigma_{r,t}} &\geq 1/\beta \\ \left(1 - \frac{1}{\kappa_{r,1}-\gamma-\tau-\delta\sigma_{r,t}}\right)^{\kappa_{r,1}-1-\sigma_{r,t}} &\geq 1/\beta. \end{aligned}$$

Given that $\sigma_{r,t} \leq \kappa_{r,1} \frac{\ln \beta - 1}{\delta \ln \beta - 1} - \frac{(\gamma + \tau + 1) \ln \beta - 1}{\delta \ln \beta - 1} < (\tilde{\kappa}_{r,1} - (\gamma + \tau + 1))/\delta$, using Fact 1, we want

$$\exp\left(\frac{\kappa_{r,1} - \sigma_{r,t} - 1}{\kappa_{r,1} - \gamma - \tau - \delta \sigma_{r,t} - 1}\right) \leq \beta$$

$$\frac{\kappa_{r,1} - \sigma_{r,t} - 1}{\kappa_{r,1} - \gamma - \tau - \delta \sigma_{r,t} - 1} \leq \ln \beta.$$

Manipulating the last expression, it can be seen that the lemma holds.

The following lemma, shows the efficiency and correctness of the AT-algorithm.

Lemma 3. *If the number of messages to deliver is more than*

$$M = 2 \frac{\delta \ln \beta - 1}{\ln \beta - 1} \left(\sum_{j=1}^5 (5/6)^{j-1} \tau \right) + \frac{((\delta(2 - \delta)/(\delta - 1)) + \tau + 1) \ln \beta - 1}{\ln \beta - 1} \in O(\log(1/\varepsilon)),$$

after running the AT-algorithm for $(e + \xi_\beta + 1 + \xi_\delta)k - \tau$ steps, where ξ_β and ξ_δ are constants arbitrarily close to 0, the number of messages left to deliver is reduced to at most M with probability at least $1 - \varepsilon$, for $\varepsilon \leq 1/(n + 1)$.

Proof. Consider the first round r such that

$$\kappa_{r,1} - \gamma - \tau \leq \tilde{\kappa}_{r,1} < \kappa_{r,1} - \gamma, \gamma = \delta(2 - \delta)/(\delta - 1). \quad (6)$$

By definition of the AT algorithm, unless the number of messages left to deliver is reduced to at most M before, such a round exists. To see why, notice in Algorithm 1 that the density estimator is either increased by τ in Line 14, or decreased by δ in Line 17, or assigned τ in Line 3 or 17. After the first assignment, we have $\tilde{\kappa}_{1,1} = \tau < \kappa_{1,1} - \gamma - \tau$, because $\kappa_{1,1} > M > 2\tau + \gamma$. We show now that condition 6 of r can not be satisfied right after decreasing the density estimator in Line 17. Consider two consecutive steps $t', t' + 1$ of some round r' such that still $\tilde{\kappa}_{r',t'} < \kappa_{r',t'} - \gamma - \tau$. If, upon a success at step t' of r' , $\tilde{\kappa}_{r',t'+1} = \tau$ by the assignment in Line 17, and $\kappa_{r',t'+1} - \gamma - \tau \leq \tilde{\kappa}_{r',t'+1}$, then $\kappa_{r',t'+1} \leq \tau + \gamma + \tau < M$ and we are done. If on the other hand $\tilde{\kappa}_{r',t'+1} = \tilde{\kappa}_{r',t'} - \delta$ by the assignment in Line 17, then $\tilde{\kappa}_{r',t'+1} = \tilde{\kappa}_{r',t'} - \delta < \kappa_{r',t'} - \gamma - \tau - \delta < \kappa_{r',t'+1} - \gamma - \tau$. Thus, the only way in which the density estimator gets inside the aforementioned range is by the increase in Line 14 and therefore round r exists.

We show now that, before leaving round r , at least τ messages are delivered with high probability so that in some future round $r'' > r$ the condition $\kappa_{r'',1} - \gamma - \tau \leq \tilde{\kappa}_{r'',1} < \kappa_{r'',1} - \gamma$ holds again. In order to do that, we divide round r in consecutive sub-rounds of size $\tau, 5/6\tau, (5/6)^2\tau, \dots$ (The fact that a number of steps is an integer is omitted throughout for clarity.) More specifically, the sub-round S_1 is the set of AT-steps in the interval $(0, \tau]$ and, for $i \geq 2$, the sub-round S_i is the set of steps in the interval $((5/6)^{i-2}\tau, (5/6)^{i-1}\tau]$. Thus, denoting $|S_i| = \tau_i$ for all $i \geq 1$, it is $\tau_1 = \tau$ and $\tau_i = (5/6)\tau_{i-1}$ for $i \geq 2$. For each $i \geq 1$, let Y_i be a random variable such that $Y_i = \sum_{t \in S_i} X_{r,t}$. Even if no message is delivered, round r still has at least the sub-round S_1 by definition of the algorithm. Given that, according with Algorithm 1, each message delivered delays the end of round r in $\beta = e + \xi_\beta$ AT-steps, for $i \geq 2$, the existence of sub-round S_i is conditioned on $Y_{i-1} \geq 5\tau_{i-1}/(6\beta)$. We show now that with big enough probability round r has 5 sub-rounds and at least τ messages are delivered. Even if messages are delivered in every step of the 5 sub-rounds (including

messages delivered in BT-steps), given that $\kappa_{r,1} > M$, the total number of messages delivered is less than $\kappa_{r,1} \frac{\ln \beta - 1}{\delta \ln \beta - 1} - \frac{(\gamma + \tau + 1) \ln \beta - 1}{\delta \ln \beta - 1}$ because $\gamma = \delta(2 - \delta)/(\delta - 1)$. Thus, Lemma 2 can be applied and the expected number of messages delivered in S_i is $E[Y_i] \geq \tau_i/\beta$. In order to use Lemma 1, we verify first its preconditions. If, at any step t , $\kappa_{r,t} \leq M$, we are done. Otherwise, we know that $\kappa_{r,t} \geq \tilde{\kappa}_{r,t} > \delta + 1$ and $(\kappa_{r,t} - \gamma)(\kappa_{r,t} - \gamma - 1)/(\kappa_{r,t} - \gamma + 1) > \delta$. Given that $\gamma = \delta(2 - \delta)/(\delta - 1)$, by Lemma 1, the random variables $X_{r,i}$ are not positively correlated, therefore, in order to bound from below the number of successful transmissions we can use the following Chernoff-Hoeffding bound [16]. For $0 < \varphi < 1$,

$$\begin{cases} Pr(Y_1 \leq (1 - \varphi)\tau_1/\beta) \leq e^{-\varphi^2 \tau_1/(2\beta)} \\ Pr(Y_i \leq (1 - \varphi)\tau_i/\beta | Y_{i-1} \geq 5\tau_{i-1}/(6\beta)) \leq e^{-\varphi^2 \tau_i/(2\beta)}, \forall i : 2 \leq i \leq 5. \end{cases}$$

Taking $\varphi = 1/6$,

$$\begin{cases} Pr(Y_1 \leq 5\tau_1/(6\beta)) \leq e^{-\varphi^2 300 \ln(1/\varepsilon)/2} \\ Pr(Y_i \leq 5\tau_i/(6\beta) | Y_{i-1} \geq 5\tau_{i-1}/(6\beta)) \leq e^{-\varphi^2 (5/6)^{i-1} 300 \ln(1/\varepsilon)/2}, \forall i : 2 \leq i \leq 5. \end{cases}$$

$$\begin{cases} Pr(Y_1 \leq 5\tau_1/(6\beta)) < e^{-2 \ln(1/\varepsilon)} \\ Pr(Y_i \leq 5\tau_i/(6\beta) | Y_{i-1} \geq 5\tau_{i-1}/(6\beta)) < e^{-2 \ln(1/\varepsilon)}, \forall i : 2 \leq i \leq 5. \end{cases}$$

Given that $\varepsilon \leq 1/(n + 1)$ and $k \leq n$, then it holds that $\varepsilon^2 + k\varepsilon \leq 1$ which implies that $\ln(1/\varepsilon) \geq \ln(\varepsilon + k)$, therefore $e^{-2 \ln(1/\varepsilon)} \leq e^{-\ln(\varepsilon + k) - \ln(1/\varepsilon)} = \varepsilon/(\varepsilon + k)$. So, more than $(5/(6(e + \xi_\beta)))\tau_i$ messages are delivered in any sub-round S_i with probability at least $1 - \varepsilon/(\varepsilon + k)$. Given that each success delays the end of round r in $\beta = e + \xi_\beta$ AT-steps, we know that, for $1 \leq i \leq 4$, sub-round S_{i+1} exists with probability at least $1 - \varepsilon/(\varepsilon + k)$. If, after any sub-round, the number of messages left to deliver is at most M , we are done. Otherwise, conditioned on these events, the total number of messages delivered over the 5 sub-rounds is at least $\sum_{j=1}^5 Y_j > \sum_{j=1}^5 (5/(6(e + \xi_\beta)))^j (e + \xi_\beta)^{j-1} \tau = (\tau/(e + \xi_\beta)) \sum_{j=1}^5 (5/6)^j > \tau$ because $\xi_\beta < 0.27$.

Thus, the same analysis can be repeated over the next round r'' such that $\kappa_{r'',1} - \gamma - \tau \leq \tilde{\kappa}_{r'',1} < \kappa_{r'',1} - \gamma$. Unless the number of messages left to deliver is reduced to at most M before, such a round r'' exists by the same argument used to prove the existence of round r . The same analysis is repeated over various rounds until all messages have been delivered or the number of messages left is at most M . Then, using conditional probability, the overall probability of success is at least $(1 - \varepsilon/(\varepsilon + k))^k$. Using Fact 1 twice, that probability is at least $1 - \varepsilon$.

It remains to be shown the time complexity of the AT algorithm. The difference between the number of messages to deliver and the density estimator right after initialization is at most $k - \tau$. This difference is increased with each message delivered by at most $\delta - 1$ and reduced at the end of each round by τ . Therefore, the total number of rounds is at most $(k - \tau + (\delta - 1)k)/\tau = \delta k/\tau - 1$. Each message delivered adds only a constant factor β to the total time, whereas the other steps in each round add up to τ . Therefore, the total time is at most $(\beta + \delta)k - \tau = (e + \xi_\beta + 1 + \xi_\delta)k - \tau$.

The time efficiency and correctness of the BT algorithm is established in the following lemma. The proof, omitted for brevity, is a straightforward computation of the probability of some message not being delivered.

Lemma 4. *If the number of messages left to deliver is at most*

$$M = 2 \frac{\delta \ln \beta - 1}{\ln \beta - 1} \left(\sum_{j=1}^5 (5/6)^{j-1} \tau \right) + \frac{((\delta(2-\delta)/(\delta-1)) + \tau + 1) \ln \beta - 1}{\ln \beta - 1},$$

there exists a constant $c > 0$ such that, after running the BT-algorithm for $c \log^2(1/\varepsilon)$ steps, all messages are delivered with probability at least $1 - \varepsilon$.

The following theorem establishes the main result.

Theorem 2. *For any one-hop Radio Network, under the model detailed in Section 1, Algorithm 1 solves the k -selection problem within $(e + 1 + \xi)k + O(\log^2(1/\varepsilon))$ communication steps, where $\xi > 0$ is any constant arbitrarily close to 0, with probability at least $1 - 2\varepsilon$ for $\varepsilon \leq 1/(n + 1)$.*

Proof. From Lemmas 3 and 4, and the definition of the algorithm, the total time is $(e + 1 + \xi_\delta + \xi_\beta)k/(1 - \xi_t) + O(\log^2(1/\varepsilon))$. Given that ξ_β , ξ_δ , and ξ_t are positive constants arbitrarily close to 0, the claim follows.

4 Conclusions and Open Problems

The general problem of enabling an unknown number of contenders the access to a shared resource was studied in this paper. The results obtained pertain to a problem of broadcasting information in a multiple-access radio-channel, but they may be straightforwardly applied to any setting that supports the same model. The specific problem studied here, k -Selection in Radio Networks, was previously studied in the literature, but assuming that a tight upper bound on the number of contenders is known. Thus, a crucial contribution of this paper was the removal of such assumption, consequently widening the scope of application of the protocol presented. Furthermore, we have assumed that messages are assigned to all nodes at the same time, increasing the potential contention for the channel with respect to scenarios where messages might arrive sparsely. To avoid collisions resulting from that contention it would be useful to have a mechanism to detect them at each node. However, we studied a more challenging scenario where only the transmitter of a message knows if it was the only one to access the channel in a time slot or not. Nonetheless, even under all these challenging conditions, the bound shown is asymptotically optimal for any sensible error-probability bound. To the best of our knowledge, the k -Selection protocol presented in this paper is the first in the Radio Networks literature that works in such conditions and is optimal.

A number of possible extensions of this work arise as natural questions that are left for future work. First, different patterns of message arrivals comprising specific application scenarios, such as Poisson arrivals and others, may also yield optimal bounds. Also, the protocol presented here improves over previous work in adversarial packet contention-resolution thanks to the adaptive nature of the protocol and the knowledge of n . Therefore, the question of how to solve the problem optimally in settings where nodes don't even know n or the feasibility of a non-adaptive optimal protocol are also important. Finally, the experimental evaluation of the protocol presented here, or others resulting from the above mentioned future work, would be useful for comparison with heuristics currently in use.

References

1. M. A. Bender, M. Farach-Colton, S. He, B. C. Kuszmaul, and C. E. Leiserson. Adversarial contention resolution for simple channels. In *17th Ann. ACM Symp. on Parallel Algorithms and Architectures*, pages 325–332, 2005.
2. J. Capetanakis. Tree algorithms for packet broadcast channels. *IEEE Trans. Inf. Theory*, IT-25(5):505–515, 1979.
3. B. S. Chlebus. Randomized communication in radio networks. In P. M. Pardalos, S. Rajasekaran, J. H. Reif, and J. D. P. Rolim, editors, *Handbook on Randomized Computing*, volume 1, pages 401–456. Kluwer Academic Publishers, 2001.
4. A. Clementi, A. Monti, and R. Silvestri. Selective families, superimposed codes, and broadcasting on unknown radio networks. In *Proc. of the 12th Ann. ACM-SIAM Symp. on Discrete Algorithms*, pages 709–718, 2001.
5. M. Gerèb-Graus and T. Tsantilas. Efficient optical communication in parallel computers. In *4th Ann. ACM Symp. on Parallel Algorithms and Architectures*, pages 41–48, 1992.
6. A. 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 the ACM*, 32:589–596, 1985.
7. R. I. Greenberg and C. E. Leiserson. Randomized routing on fat-trees. *Advances in Computing Research*, 5:345–374, 1989.
8. J. F. Hayes. An adaptive technique for local distribution. *IEEE Trans. Comm.*, COM-26:1178–1186, 1978.
9. P. Indyk. Explicit constructions of selectors and related combinatorial structures, with applications. In *Proc. of the 13th Ann. ACM-SIAM Symp. on Discrete Algorithms*, pages 697–704, 2002.
10. J. Komlòs and A. Greenberg. An asymptotically nonadaptive algorithm for conflict resolution in multiple-access channels. *IEEE Trans. Inf. Theory*, 31:303–306, 1985.
11. D. R. Kowalski. On selection problem in radio networks. In *Proc. 24th Ann. ACM Symp. on Principles of Distributed Computing*, pages 158–166, 2005.
12. E. Kushilevitz and Y. Mansour. An $\Omega(D \log(N/D))$ lower bound for broadcast in radio networks. *SIAM Journal on Computing*, 27(3):702–712, 1998.
13. C. U. Martel. Maximum finding on a multiple access broadcast network. *Inf. Process. Lett.*, 52:7–13, 1994.
14. V. Mikhailov and B. S. Tsybakov. Free synchronous packet access in a broadcast channel with feedback. *Problemy Peredachi Inform*, 14(4):32–59, 1978.
15. D. S. Mitrinović. *Elementary Inequalities*. P. Noordhoff Ltd. - Groningen, 1964.
16. M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005.
17. D. E. Willard. Log-logarithmic selection resolution protocols in a multiple access channel. *SIAM Journal on Computing*, 15:468–477, 1986.