

Unbounded Contention Resolution in Multiple-Access Channels*

Antonio Fernández Anta¹,
Miguel A. Mosteiro^{2,3}, and Jorge Ramón Muñoz³

¹ Institute IMDEA Networks, Madrid, Spain
`antonio.fernandez@imdea.org`

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

³ LADyR, GSyC, Universidad Rey Juan Carlos, Móstoles, Spain
`jorge.ramon@madrimasd.net`

Abstract. A frequent problem in settings where a unique resource must be shared among users is how to resolve the contention that arises when all of them must use it, but the resource allows only for one user each time. The application of efficient solutions for this problem spans a myriad of settings such as radio communication networks or databases. For the case where the number of users is unknown, recent work has yielded fruitful results for local area networks and radio networks, although either a (possibly loose) upper bound on the number of users needs to be known [7], or the solution is suboptimal [2], or it is only implicit [11] or embedded [6] in other problems, with bounds proved only asymptotically. In this paper, under the assumption that collision detection or information on the number of contenders is not available, we present a novel protocol for contention resolution in radio networks, and we recreate a protocol previously used for other problems [11, 6], tailoring the constants for our needs. In contrast with previous work, both protocols are proved to be optimal up to a small constant factor and with high probability for big enough number of contenders. Additionally, the protocols are evaluated and contrasted with the previous work by extensive simulations. The evaluation shows that the complexity bounds obtained by the analysis are rather tight, and that both protocols proposed have small and predictable complexity for many system sizes (unlike previous proposals).

1 Introduction

The topic of this work is the resolution of contention in settings where an unknown number of users must access a single shared resource, but multiple simultaneous accesses are not feasible. The scope of interest in this problem is

* A brief announcement of this work has been presented at PODC 2011. This work is supported in part by the Comunidad de Madrid grant S2009TIC-1692, Spanish MICINN grant TIN2008-06735-C02-01, and National Science Foundation grant CCF-0937829.

wide, ranging from radio and local area networks to databases and transactional memory. (See [2] and the references therein.)

A common theme in protocols used for this problem is the adaptive adjustment of some user variable that reflects its eagerness in trying to access the shared resource. Examples of such variable are the probability of transmitting a message in a radio network or the frequency of packet transmission in a local area network. When such adjustment reduces (resp. increases) the contention, the technique is called *back-off* (resp. *back-on*). Combination of both methods are called *back-on/back-off*. Protocols used may be further characterized by the rate of adjustment. E.g., *exponential back-off*, *polynomial back-on*, etc. In particular, exponential back-off is widely used and it has proven to be efficient in practical applications where statistical arrival of contenders is expected. Nevertheless, worst case arrival patterns, such as bursty or *batched* arrivals, are frequent [18, 12].

A technique called LOGLOG-ITERATED BACK-OFF was shown to be within a sublogarithmic factor from optimal with high probability in [2].¹ The protocol was presented in the context of packet contention resolution in local area networks for batched arrivals. Later on, also for batched arrivals, we presented a back-on/back-off protocol in [7], instantiated in the k -selection problem in Radio Networks (defined in Section 2). The latter protocol, named here LOG-FAILS ADAPTIVE, is asymptotically optimal for any significant probability of error, but additionally requires that some upper bound (possibly loose) on the number of contenders is known. In the present paper, we remove such requirement. In particular, we present and analyze a protocol that we call ONE-FAIL ADAPTIVE for k -selection in Radio Networks. We also recreate and analyze another protocol for k -selection, called here EXP BACK-ON/BACK-OFF, which was previously embedded in protocols for other problems and analyzed only asymptotically [11, 6]. Our analysis shows that ONE-FAIL ADAPTIVE and EXP BACK-ON/BACK-OFF, both of independent interest, resolve contention among an unknown and unbounded² number of contenders with high probability in optimal time up to constants. Additionally, by means of simulations, we evaluate and contrast the average performance of all four protocols. The simulations show that the complexity bounds obtained in the analysis (with high probability) for these protocols are rather tight for the input sizes considered. Additionally, they show that they are faster than LOGLOG-ITERATED BACK-OFF and more predictable for all network sizes than LOG-FAILS ADAPTIVE.

Roadmap: The rest of the paper is organized as follows. In the following section the problem, model, related work and results are detailed. In Section 3, we introduce ONE-FAIL ADAPTIVE and its analysis. EXP BACK-ON/BACK-OFF is detailed and analyzed in Section 4. The results of the empirical contrast of all

¹ For k contenders, we define *with high probability* to mean with probability at least $1 - 1/k^c$ for some constant $c > 0$.

² We use the term unbounded to reflect that not even an upper bound on the number of contenders is known. This should not be confused with the infinitely-many users model where there are countably infinitely many stations. [4]

four protocols is given in Section 5 and we finish with concluding remarks and open problems in Section 6.

2 Preliminaries

A well-studied example of unique-resource 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. The particular model of multiple-access channel we consider here is the Radio Network, a model of communication network where the channel is contended (even if radio communication is not actually used [4]). We first precise our model of Radio Network as follows.

The Model: We consider a Radio Network comprised of n 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, not even the number of nodes n or their own label. 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. As in [2, 16, 11], we assume that a node becomes idle upon delivering its message, for instance when an explicit acknowledgement is received (like in the IEEE 802.11 Medium Access Control protocol [1]). For settings where the channel does not provide such functionality, such as Sensor Networks, a hierarchical infrastructure may be predefined to achieve it [6], or a leader can be elected as the node responsible for acknowledging successful transmissions [22].

One of the problems that require contention resolution in Radio Networks is the problem known in the literature as *all-broadcast* [4], or *k-selection* [16]. In *k-selection*, a set of k out of n network nodes have to access a unique shared

channel of communication, each of them at least once. As in [2, 16, 11], in this paper we study k -selection when all messages arrive simultaneously, or in a *batch*. Under this assumption the k -selection problem is called *static*. A *dynamic* counterpart where messages arrive at different times was also studied [16].

The Problem: Given a Radio Network where k network nodes are activated by a message that arrives simultaneously to all of them, the *static k -selection* problem is solved when each node has delivered its message.

Related Work: A number of fruitful results for contention resolution have been obtained assuming availability of collision detection. Martel presented in [19] a randomized adaptive protocol for k -Selection that works in $O(k + \log n)$ time in expectation³. As argued by Kowalski in [16], this protocol can be improved to $O(k + \log \log n)$ in expectation using Willard's expected $O(\log \log n)$ selection protocol of [23]. 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 (i.e., protocols where all nodes use the same probability of transmission to transmit in any given time slot). 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 [17], 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.

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 [3, 13, 20]. In all these results the algorithmic technique, known as *tree algorithms*, relies on modeling the protocol as a complete binary tree where the messages are placed at the leaves. Later, Greenberg and Winograd [10] showed a lower bound for that class of protocols of $\Omega(k \log_k n)$. Regarding oblivious algorithms, Komlós and Greenberg [15] 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 [5] showed a lower bound of $\Omega(k \log(n/k))$, which also holds for adaptive algorithms if collision detection is not available. In [16], Kowalski presented the construction of an oblivious deterministic protocol that, using the explicit selectors of Indyk [14], gives a $O(k \text{ polylog } n)$ upper bound without collision detection.

In [9], 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 at most h messages. Making $h = k$ this protocol can be used to solve static k -selection. However, it requires that nodes know h .

Extending previous work on tree algorithms, Greenberg and Leiserson [11] presented randomized routing strategies in fat-trees for bounded number of

³ Throughout this paper, \log means \log_2 unless otherwise stated.

messages. Underlying their algorithm lies a sawtooth technique used to “guess” the appropriate value for some critical parameter (load factor), that can be used to “guess” the number of contenders in static k -selection. Furthermore, modulo choosing the appropriate constants, EXP BACK-ON/BACK-OFF uses the same sawtooth technique. Their algorithm uses n and it is analyzed only asymptotically.

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 static k -selection, have been analyzed in [2]. In that paper, it is shown that *r-exponential back-off*, a monotonic technique used widely that has proven to be efficient for many practical applications is in $\Theta(k \log^{\log r} k)$ for batched arrivals. 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 . In the same paper, the sawtooth technique used in [11] is informally described in a paragraph while pointing out that it yields linear time for contention resolution thanks to non-monotonicity, but no analysis is provided.

Later on, Farach-Colton and Mosteiro presented an optimal protocol for Gossiping in Radio Networks in [6]. The sawtooth technique embedded in [11] is used in that paper as a subroutine to resolve contention in linear time as in EXP BACK-ON/BACK-OFF. However, the algorithm makes use of n to achieve the desired probability of success and the analysis is only asymptotical.

A randomized adaptive protocol for static k -selection in a one-hop Radio Network without collision detection was presented in [7]. The protocol is shown to solve the problem in $(e + 1 + \xi)k + O(\log^2(1/\varepsilon))$ steps with probability at least $1 - 2\varepsilon$, where $\xi > 0$ is an arbitrarily small constant and $0 < \varepsilon \leq 1/(n + 1)$. Modulo a constant factor, the protocol is optimal if $\varepsilon \in \Omega(2^{-\sqrt{n}})$. However, the algorithm makes use of the value of ε , which must be upper bounded as above in order to guarantee the running time. Therefore, knowledge of n is required.

Our Results: In this paper, we present a novel randomized protocol for static k -selection in a one-hop Radio Network, and we recreate a previously used technique suiting the constants for our purpose and without making use of n . Both protocols work without collision detection and do not require information about the number of contenders. As mentioned, these protocols are called ONE-FAIL ADAPTIVE and EXP BACK-ON/BACK-OFF. It is proved that ONE-FAIL ADAPTIVE solves static k -selection within $2(\delta + 1)k + O(\log^2 k)$ steps, with probability at least $1 - 2/(1 + k)$, for $e < \delta \leq \sum_{j=1}^5 (5/6)^j$. On the other hand, EXP BACK-ON/BACK-OFF is shown to solve static k -selection within $4(1 + 1/\delta)k$ steps with probability at least $1 - 1/k^c$ for some constant $c > 0$, $0 < \delta < 1/e$, and big enough k . Given that k is a lower bound for this problem, both protocols are optimal (modulo a small constant factor) for big enough number of contenders.

Observe that the bounds and the probabilities obtained are given as functions of the parameter k , as done in [2], since this is the input parameter of our version of the problem. A fair comparison with the results obtained as function of k and n would require that k is large enough, so that $n = \Omega(k^c)$, for some constant c .

Both protocols presented are of interest because, although protocol EXP BACK-ON/BACK-OFF is simpler, ONE-FAIL ADAPTIVE achieves a better multiplicative factor, although the constant in the sublinear additive factor may be big for small values of k .

Additionally, results of the evaluation by simulation of the average behavior of ONE-FAIL ADAPTIVE and EXP BACK-ON/BACK-OFF and a comparison with LOG-FAILS ADAPTIVE and LOGLOG-ITERATED BACK-OFF are presented. Both algorithms ONE-FAIL ADAPTIVE and EXP BACK-ON/BACK-OFF run faster than LOGLOG-ITERATED BACK-OFF on average, even for small values of k . Although LOGLOG-ITERATED BACK-OFF has higher asymptotic complexity, one may have expected that it may run fast for small networks. On the other hand, the knowledge on a bound of k assumed by LOG-FAILS ADAPTIVE seems to provide an edge with respect to ONE-FAIL ADAPTIVE and EXP BACK-ON/BACK-OFF for large values of k . However, LOG-FAILS ADAPTIVE has a much worse behavior than the proposed protocols for small to moderate network sizes ($k \leq 10^5$). In any case, for all values of k simulated, ONE-FAIL ADAPTIVE and EXP BACK-ON/BACK-OFF have a very stable and efficient behavior.

3 One-Fail Adaptive

As in LOG-FAILS ADAPTIVE [7], ONE-FAIL ADAPTIVE is composed by two interleaved randomized algorithms, each intended to handle the communication for different levels of contention. One of the algorithms, which we call *AT*, is intended for delivering messages while the number of nodes contending for the channel is above some logarithmic threshold (to be defined later). The other algorithm, called *BT*, has the purpose of handling message delivery after that number is below that threshold. Nonetheless, a node may transmit using the BT (resp. AT) algorithm even if the number of messages left to deliver is above (resp. below) that threshold.

Both algorithms, AT and BT, are based on transmission trials with certain probability and what distinguishes them is just the specific probability value used. It is precisely the particular values of probability used in each algorithm what differentiates ONE-FAIL ADAPTIVE from LOG-FAILS ADAPTIVE. For the BT algorithm, the probability of transmission is inversely logarithmic on the number of messages already transmitted, while in LOG-FAILS ADAPTIVE that probability was fixed. For the AT algorithm the probability of transmission is the inverse of an estimation on the number of messages left to deliver. In ONE-FAIL ADAPTIVE this estimation is updated continuously, whereas in LOG-FAILS ADAPTIVE it was updated after some steps without communication. These changes yield a protocol still linear, but now it is not necessary to know n . Further details can be seen in Algorithm 1.

For clarity, Algorithms AT and BT are analyzed separately taking into account in both analyses the presence of the other. We show the efficiency of the AT algorithm in producing successful transmissions while the number of messages left is above some logarithmic threshold, and the efficiency of the BT algorithm

Algorithm 1. ONE-FAIL ADAPTIVE. Pseudocode for node x . δ is a constant such that $e < \delta \leq \sum_{j=1}^5 (5/6)^j$.

```

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

```

handling the communication after that threshold is crossed. For the latter, we use standard probability computations to show our time upper bound. For the AT algorithm, we use concentration bounds to show that the messages are delivered with large enough probability, while the density estimator $\tilde{\kappa}$ does not exceed the actual number of messages left. This second proof is more involved since it requires some preliminary lemmas. We establish here the main theorem, which is direct consequence of the lemmata described that can be found in [8].

Theorem 1. For any $e < \delta \leq \sum_{j=1}^5 (5/6)^j$ and for any one-hop Radio Network under the model detailed in Section 1, ONE-FAIL ADAPTIVE solves static k -selection within $2(\delta + 1)k + O(\log^2 k)$ communication steps, with probability at least $1 - 2/(1 + k)$.

4 Exp Back-on/Back-off

The algorithm presented in this section is based in contention windows. That is, each node repeatedly chooses uniformly one time slot within an interval, or *window*, of time slots to transmit its message. Regarding the size of such window, our protocol follows a back-on/back-off strategy. Namely, the window is increased in an outer loop and decreased in an inner loop, as detailed in Algorithm 2.

The intuition for the algorithm is as follows. Let m be the number of messages left at a given time right before using a window of size w . We can think of

Algorithm 2. Window size adjustment in EXP BACK-ON/BACK-OFF. $0 < \delta < 1/e$ is a constant.

```

1 for  $i = \{1, 2, \dots\}$  do
2    $w \leftarrow 2^i$ 
3   while  $w \geq 1$  do
4     Choose uniformly a step within the next  $w$  steps
5      $w \leftarrow w \cdot (1 - \delta)$ 

```

the algorithm as a random process where m balls (modelling the messages) are dropped uniformly in w bins (modelling time slots). We will show that, if $m \leq w$, for large enough m , with high probability, at least a constant fraction of the balls fall alone in a bin. Now, we can repeat the process removing this constant fraction of balls and bins until all balls have fallen alone. Since nodes do not know m , the outer loop increasing the size of the window is necessary. The analysis follows.

Lemma 1. *For $k \geq m \geq (2e/(1 - e\delta)^2)(1 + (\beta + 1/2) \ln k)$, $0 < \delta < 1/e$, $m \leq w$, and $\beta > 0$, if m balls are dropped in w bins uniformly at random, the probability that the number of bins with exactly one ball is less than δm is at most $1/k^\beta$.*

Proof. Since a bigger number of bins can only reduce the number of bins with more than one ball, if the claim holds for $w = m$ it also holds for $w > m$. Thus, it is enough to prove the first case. The probability for a given ball to fall alone in a given bin is $(1/m)(1 - 1/m)^{m-1} \geq 1/(em)$. Let X_i be a random variable that indicates if there is exactly one ball in bin i . Then, $\Pr(X_i = 1) \geq 1/e$. To handle the dependencies that arise in balls and bins problems, we approximate the joint distribution of the number of balls in all bins by assuming the load in each bin is an independent Poisson random variable with mean 1. Let X be a random variable that indicates the total number of bins with exactly one ball. Then, $\mu = E[X] = m/e$. Using Chernoff-Hoeffding bounds [21], $\Pr(X \leq \delta m) \leq \exp\left(-m(1 - e\delta)^2/(2e)\right)$, because $0 < \delta < 1/e$.

As shown in [21], any event that takes place with probability p in the Poisson case takes place with probability at most $pe\sqrt{m}$ in the exact case. Then, we want to show that $\exp\left(-m(1 - e\delta)^2/(2e)\right) e\sqrt{m} \leq k^{-\beta}$, which is true for $m \geq \frac{2e}{(1 - e\delta)^2} \left(1 + \left(\frac{1}{2} + \beta\right) \ln k\right)$. \square

Theorem 2. *For any constant $0 < \delta < 1/e$, EXP BACK-ON/BACK-OFF solves static k -selection within $4(1 + 1/\delta)k$ steps with probability at least $1 - 1/k^c$, for some constant $c > 0$ and big enough k .*

Proof. Consider an execution of the algorithm on k nodes. Let a round be the sequence of time steps corresponding to one iteration of the inner loop of Algorithm 2, i.e. the time steps of a window. Let a phase be the sequence of rounds corresponding to one iteration of the outer loop of Algorithm 2, i.e. when the window is monotonically reduced.

Consider the first round when $k \leq w < 2k$. Assume no message was transmitted successfully before. (Any messages transmitted could only reduce the running time.) By Lemma 1, we know that, for $0 < \delta < 1/e$ and $\beta > 0$, at least δk messages are transmitted in this round with probability at least $1 - 1/k^\beta$, as long as $k \geq \tau$, where $\tau \triangleq (2e/(1 - e\delta)^2)(1 + (\beta + 1/2) \ln k)$.

Conditioned on this event, for some $\delta_1 \geq \delta$ fraction of messages transmitted in the first round, using the same lemma we know that in the following round at least $\delta(1 - \delta_1)k$ messages are transmitted with probability at least $1 - 1/k^\beta$, as long as $(1 - \delta_1)k \geq \tau$. This argument can be repeated for each subsequent round until the number of messages left to be transmitted is less than τ . Furthermore, given that the size of the window is monotonically reduced within a phase until $w = 1$, even if the fraction of messages transmitted in each round is just δ , the overall probability of reducing the number of messages left from k to τ within this phase is at least $(1 - 1/k^\beta)^{\log_{1/(1-\delta)}(2k)}$.

Consider now the first round of the following phase, i.e. when $2k \leq w < 4k$. Assume that at most τ nodes still hold a message to be transmitted. Using the union bound, the probability that two or more of m nodes choose a given step in a window of size w is at most $\binom{m}{2}/w^2$. Applying again the union bound, the probability that in any step two or more nodes choose to transmit is at most $\binom{m}{2}/w \leq \binom{\tau}{2}/(2k) = \tau(\tau + 1)/(4k)$.

Therefore, using conditional probability, in order to complete the proof, it is enough to show that

$$\begin{aligned} & \left(1 - \frac{\tau(\tau + 1)}{4k}\right) \left(1 - \frac{1}{k^\beta}\right)^{\log_{1/(1-\delta)}(2k)} \geq 1 - \frac{1}{k^c}, \text{ for some constant } c > 0 \\ \exp\left(-\frac{\tau(\tau + 1)}{4k - \tau(\tau + 1)} - \frac{\log_{1/(1-\delta)}(2k)}{k^\beta - 1}\right) & \geq \exp\left(-\frac{1}{k^c}\right) \\ \frac{\tau(\tau + 1)}{4k - \tau(\tau + 1)} + \frac{\log_{1/(1-\delta)}(2k)}{k^\beta - 1} & \leq \frac{1}{k^c}. \end{aligned} \quad (1)$$

Given that δ is a constant and fixing $\beta > 0$ as a constant, Inequality 1 is true for some constant $c < \min\{1, \beta\}$, for big enough k . Telescoping the number of steps up to the first round when $w = 4k$, the running time is less than $4k + 2k \sum_{i=0}^{\infty} \sum_{j=0}^{\infty} (1 - \delta)^j / 2^i = 4(1 + 1/\delta)k$. \square

5 Evaluation

In order to evaluate the expected behavior of the algorithms ONE-FAIL ADAPTIVE and EXP BACK-ON/BACK-OFF, and compare it with the previously proposed algorithms LOGLOG-ITERATED BACK-OFF and LOG-FAILS ADAPTIVE, we have simulated the four algorithms. The simulations measure the number of steps that the algorithms take until the static k -selection problem has been solved, i.e., each of the k activated nodes of the Radio Network has delivered its message, for different values of k . Several of the algorithms have parameters that can be adapted. The value of these parameters is the same for all the simulations of




Fig. 1. Number of steps to solve static k -selection, per number of nodes k

Table 1. Ratio steps/nodes as a function of the number of nodes k

k	10	10^2	10^3	10^4	10^5	10^6	10^7	Analysis
LOG-FAILS ADAPTIVE $\xi_t = 1/2$	46.4	1292.4	181.9	26.6	9.4	8.0	7.8	7.8
LOG-FAILS ADAPTIVE $\xi_t = 1/10$	26.3	3289.2	593.8	50.3	11.5	4.5	4.4	4.4
ONE-FAIL ADAPTIVE	4.0	6.9	7.4	7.4	7.4	7.4	7.4	7.4
EXP BACK-ON/BACK-OFF	4.0	5.5	5.2	7.2	6.6	5.6	7.9	14.9
LOGLOG-ITERATED BACK-OFF	5.6	8.6	9.6	9.2	10.5	10.5	10.1	$\Theta\left(\frac{\log \log k}{\log \log \log k}\right)$

the same algorithm (except the parameter ε of LOG-FAILS ADAPTIVE that has to depend on k). For EXP BACK-ON/BACK-OFF the parameter is chosen to be $\delta = 0.366$. For ONE-FAIL ADAPTIVE the parameter is chosen to be $\delta = 2.72$. For LOG-FAILS ADAPTIVE, the parameters (see their meaning in [7]) are chosen to be $\xi_\delta = \xi_\beta = 0.1$ and $\varepsilon \approx 1/(k + 1)$, while two values of ξ_t have been used, $\xi_t = 1/2$ and $\xi_t = 1/10$. Finally, LOGLOG-ITERATED BACK-OFF is simulated with parameter $r = 2$ (see [2]).

Figure 1 presents the average number of steps taken by the simulation of the algorithms. The plot shows the the average of 10 runs for each algorithm as a function of k . In this figure it can be observed that LOG-FAILS ADAPTIVE takes significantly larger number of steps than the other algorithms for moderately small values of k (up to 10^5). Beyond $k = 10^5$ all algorithms seem to have a similar behavior.

A higher level of detail can be obtained by observing Table 1, which presents the ratio obtained by dividing the number of steps (plotted in Figure 1) by the value of k , for each k and each algorithm. In this table, the bad behavior of LOG-FAILS ADAPTIVE for moderate values of k can be observed, with values of the ratio well above those for large k . It seems like the value of ξ_t used has an impact in this ratio, so that the smaller value $\xi_t = 1/10$ causes larger ratio values. Surprisingly, for large values of k ($k \geq 10^6$), the ratios observed are almost exactly the constant factors of k obtained from the analysis [7]. (Recall that all the analyses we refer to are with high probability while the simulation results

are averages.) This may indicate that the analysis with high probability is very tight and that the term $O(\log^2(1/\varepsilon))$ that appears in the complexity expression is mainly relevant for moderate values of k . The ratio obtained for large k by LOG-FAILS ADAPTIVE with $\xi_t = 1/10$ is the smallest we have obtained in the set of simulations. LOGLOG-ITERATED BACK-OFF, on its hand, seems to have a constant ratio of around 10. In reality this ratio is not constant but, since it is sublogarithmic, this fact can not be observed for the (relatively small) values of k simulated.

Regarding the ratios obtained for the algorithms proposed in this paper, they seem to show that the constants obtained in the analyses (with high probability) are very accurate. Starting at moderately large values of k (10^3 and up) the ratio for ONE-FAIL ADAPTIVE becomes very stable and equal to the value of 7.4 obtained in the analysis. The ratios for the EXP BACK-ON/BACK-OFF simulations, on their hand, move between 4 and 8, while the analysis for the value of δ used yields a constant factor of 14.9. Hence, the ratios are off by only a small constant factor. To appreciate these values it is worth to note that the smallest ratio expected by any algorithm in which nodes use the same probability at any step is e , so these values are only a small factor away from this optimum ratio. In summary, the algorithms proposed here have small and stable ratios for all values of k considered.

6 Conclusions and Open Problems

In this work, we have shown optimal randomized protocols (up to constants) for static k -selection in Radio Networks that do not require any knowledge on the number of contenders. Future work includes the study of the dynamic version of the problem when messages arrive at different times under the same model, either assuming statistical or adversarial arrivals. The stability of monotonic strategies (exponential back-off) has been studied in [2]. In light of the improvements obtained for batched arrivals, the application of non-monotonic strategies to the dynamic problem is promising.

References

1. Balador, A., Movaghar, A., Jabbehdari, S.: History based contention window control in IEEE 802.11 MAC protocol in error prone channel. *Journal of Computer Science* 6, 205–209 (2010)
2. Bender, M.A., Farach-Colton, M., He, S., Kuszmaul, B.C., Leiserson, C.E.: Adversarial contention resolution for simple channels. In: 17th Ann. ACM Symp. on Parallel Algorithms and Architectures, pp. 325–332 (2005)
3. Capetanakis, J.: Tree algorithms for packet broadcast channels. *IEEE Trans. Inf. Theory* IT-25(5), 505–515 (1979)
4. Chlebus, B.S.: Randomized communication in radio networks. In: Pardalos, P.M., Rajasekaran, S., Reif, J.H., Rolim, J.D.P. (eds.) *Handbook on Randomized Computing*, vol. 1, pp. 401–456. Kluwer Academic Publishers, Dordrecht (2001)

5. Clementi, A., Monti, A., Silvestri, R.: Selective families, superimposed codes, and broadcasting on unknown radio networks. In: Proc. of the 12th Ann. ACM-SIAM Symp. on Discrete Algorithms, pp. 709–718 (2001)
6. Farach-Colton, M., Mosteiro, M.A.: Sensor network gossiping or how to break the broadcast lower bound. In: Tokuyama, T. (ed.) ISAAC 2007. LNCS, vol. 4835, pp. 232–243. Springer, Heidelberg (2007)
7. Fernández Anta, A., Mosteiro, M.A.: Contention resolution in multiple-access channels: k-selection in radio networks. *Discrete Mathematics, Algorithms and Applications* 2(4), 445–456 (2010)
8. Fernández Anta, A., Mosteiro, M.A., Muñoz, J.R.: Unbounded Contention Resolution in Multiple-Access Channels. arXiv:1107.0234v1 [cs.DC] (July 2011)
9. Gerèb-Graus, M., Tsantilas, T.: Efficient optical communication in parallel computers. In: 4th Ann. ACM Symp. on Parallel Algorithms and Architectures, pp. 41–48 (1992)
10. Greenberg, A., Winograd, S.: 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)
11. Greenberg, R.I., Leiserson, C.E.: Randomized routing on fat-trees. *Advances in Computing Research* 5, 345–374 (1989)
12. Gusella, R.: A measurement study of diskless workstation traffic on an ethernet. *IEEE Transactions on Communications* 38(9), 1557–1568 (1990)
13. Hayes, J.F.: An adaptive technique for local distribution. *IEEE Trans. Comm.* COM-26, 1178–1186 (1978)
14. Indyk, P.: Explicit constructions of selectors and related combinatorial structures, with applications. In: Proc. of the 13th Ann. ACM-SIAM Symp. on Discrete Algorithms, pp. 697–704 (2002)
15. Komlòs, J., Greenberg, A.: An asymptotically nonadaptive algorithm for conflict resolution in multiple-access channels. *IEEE Trans. Inf. Theory* 31, 303–306 (1985)
16. Kowalski, D.R.: On selection problem in radio networks. In: Proc. 24th Ann. ACM Symp. on Principles of Distributed Computing, pp. 158–166 (2005)
17. Kushilevitz, E., Mansour, Y.: An $\Omega(D \log(N/D))$ lower bound for broadcast in radio networks. *SIAM Journal on Computing* 27(3), 702–712 (1998)
18. Leland, W.E., Taqqu, M.S., Willinger, W., Wilson, D.V.: On the self-similar nature of ethernet traffic (extended version). *IEEE/ACM Transactions on Networking* 2, 1–15 (1994)
19. Martel, C.U.: Maximum finding on a multiple access broadcast network. *Inf. Process. Lett.* 52, 7–13 (1994)
20. Mikhailov, V., Tsybakov, B.S.: Free synchronous packet access in a broadcast channel with feedback. *Problemy Peredachi Inform* 14(4), 32–59 (1978)
21. Mitzenmacher, M., Upfal, E.: *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, Cambridge (2005)
22. Nakano, K., Olariu, S.: A survey on leader election protocols for radio networks. In: Proc. of the 6th Intl. Symp. on Parallel Architectures, Algorithms and Networks (2002)
23. Willard, D.E.: Log-logarithmic selection resolution protocols in a multiple access channel. *SIAM Journal on Computing* 15, 468–477 (1986)