



Counting in one-hop beeping networks

Arnaud Casteigts, Yves Métivier, John Michael Robson, Akka Zemmari

► To cite this version:

Arnaud Casteigts, Yves Métivier, John Michael Robson, Akka Zemmari. Counting in one-hop beeping networks. Theoretical Computer Science, 2019, 780, pp.20-28. [⟨10.1016/j.tcs.2019.02.009⟩](https://doi.org/10.1016/j.tcs.2019.02.009). [⟨hal-02388912⟩](https://hal.science/hal-02388912)

HAL Id: hal-02388912

<https://hal.science/hal-02388912v1>

Submitted on 25 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC 4.0 - Attribution - Non-commercial use - International License

Counting in One-Hop Beeping Networks

A. Casteigts, Y. Métivier, J.M. Robson and A. Zemmari

*Université de Bordeaux - Bordeaux INP
LaBRI UMR CNRS 5800
351 cours de la Libération, 33405 Talence, France
{acasteig, metivier, robson, zemmari}@labri.fr*

Abstract

We consider networks of processes which interact with beeps. In the basic model defined by Cornejo and Kuhn [8], which we refer to as the *BL* variant, processes can choose in each round either to beep or to listen. Those who beep are unable to detect simultaneous beeps. Those who listen can only distinguish between silence and the presence of at least one beep. Beeping models are weak in essence and even simple tasks may become difficult or unfeasible with them.

In this paper, we address the problem of computing how many participants there are in a one-hop network: the *counting* problem. We first observe that no algorithm can compute this number with certainty in *BL*, even if it is randomised (Las Vegas). We thus consider the stronger variant where beeping nodes are able to detect simultaneous beeps, referred to as $B_{cd}L$ (*cd* stands for *collision detection*). We prove that at least n rounds are necessary in $B_{cd}L$, and we present an algorithm whose running time is $O(n)$ rounds with high probability. Further experimental results show that its expected running time is less than $10n$. Finally, we discuss how this algorithm can be adapted in other beeping models. In particular, we show that the algorithm can be emulated in *BL*, at the cost of a logarithmic slowdown and of trading its Las Vegas nature (certain result, uncertain time) against a Monte Carlo one (certain time, uncertain result).

Keywords: Beeping model, Size computation, Counting problem, Las Vegas algorithms.

1. Introduction

Distributed algorithms are concerned with assumptions relating to various aspects like the structure of the network (e.g. trees, rings, planar graphs, complete graphs) or knowledge available to the nodes (e.g. a bound on the network size, unique identifiers, or port numbering). Another important aspect is the size of messages, which may range from unbounded (the *LOCAL* model) to logarithmic size (*CONGEST* model), to constant size (e.g. simple bits) [23].

A natural approach in distributed computing is to reduce the assumptions as much as possible, in order to make positive results more general. Hence, when a problem is solved in some strong model, one naturally strives to solve it in a weaker model. In a recent series of works [8, 24, 1, 12, 25, 10], the community has started to explore new models of communications that are even weaker than constant size messages in anonymous networks, namely *beeping models*.

In these models, the only communication capabilities offered to the nodes are to *beep* or to *listen* for the beeps of others. Several variants exist. In [8], a node that beeps is unable to detect whether other nodes have beeped at the same time. When listening, it can distinguish between silence or the presence of at least one beep, but it cannot distinguish between one and several beeps. In Section 6 of [1], a stronger variant is considered where beeping nodes can detect whether other nodes are beeping simultaneously (sender side collision detection). In [24] and Section 4 of [1], yet another variant is considered where the nodes can tell the difference between silence, one beep, and several beeps. The ability to detect internal collision is denoted by B_{cd} (B otherwise) and that of detecting peripheral collisions is denoted by L_{cd}

(L otherwise). The various models in literature can be reformulated in these terms. The basic model introduced by Cornejo and Kuhn in [8] is BL ; the model considered by Afek et al. in [1] (Section 6) and Scott et al. in [25] is $B_{cd}L$; and the model considered in [24] and in Section 4 of [1] is BL_{cd} . To the best of our knowledge, $B_{cd}L_{cd}$ has only been considered in [4].

Studying weak models of computation is interesting in its own right. In addition, beeping models turn out to be relevant to model real-world applications or phenomena. For instance, they reflect the features of a network at the lowest level (physical layer), where a node can probe or emit signals, with or without collision detection. At a higher level of abstraction, beeping models also reflect some communication patterns in biology, such as *lateral inhibition* among neighboring cells [7].

Contributions. In this paper, we consider the counting problem in which nodes must determine the size of the network. We focus on the case that the communication graph is complete (*one-hop* networks), implying that each node can hear and be heard by all the others. As it turns out, even this version of the problem cannot be solved deterministically, due to the inherent lack of symmetry-breaking mechanisms in the beeping model. We thus consider randomised algorithms and start by providing a *Las Vegas* (LV) algorithm – correct result but uncertain time, though finite with probability 1 – for solving the counting problem in $B_{cd}L$. The expected running time of this algorithm is linear in the number of nodes (with high probability). On the negative side, we show that n is a lower bound, which makes the algorithm optimal up to a constant factor (experimentally estimated to less than 10). Unfortunately, we also observe that no LV algorithm exists in the weaker BL model, leading us to consider a *Monte Carlo* (MC) variant in BL where the result is correct only with some threshold probability. This algorithm relies on a technique developed in [4], which enables the emulation of B_{cd} beeps with B beeps, at the cost of a logarithmic slowdown. We also discuss how to adapt the algorithm in the stronger $B_{cd}L_{cd}$ model, this time with a (constant factor) gain in time complexity.

Related Work. As explained by Chlebus [5], detecting a collision in a radio network is to be able to distinguish between 0 messages and at least 2 messages while detecting a collision in the beeping model is to be able to distinguish between 1 message and at least 2 messages. Thus results on the counting problem in the context of radio networks cannot be applied directly in the context of beeping models. Another problem which has been well-studied, in this context, is the k -Selection problem (also known as all-broadcast); as explained by Anta and Mosteiro [2] (see also [14, 3]): it 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”. Finally, selection problems, in general, differ from the counting problem studied in this paper

1. by the fact that the collision detection does not have the same meaning in the context of radio networks and in the context of the beeping models, and
2. by the fact that in our case we look for the exact number of nodes and the k -selection problem asks that nodes access the unique shared channel at least once.

Afek et al. [1], from considerations concerning the development of certain cells, studied the Maximal Independent Set (MIS) problem in the discrete beeping model BL as presented in [8]. They consider, in particular, the wake-on-beep model (sleeping nodes wake up upon receiving a beep) and sender-side collision detection $B_{cd}L$: they give a $O(\log^2 n)$ rounds MIS algorithm. After this work, Scott et al. [25] present in the model $B_{cd}L$ a randomised algorithm with feedback mechanism whose expected time to compute a MIS is $O(\log n)$.

More generally, Navlakha and Bar-Joseph present in [22] a general survey on similarities and differences between distributed computations in biological and computational systems and, in this framework, the importance of the beeping model.

In [8], Cornejo and Kuhn study the interval colouring problem: an interval colouring assigns to each node an interval (contiguous fraction) of resources such that neighbouring nodes do not share resources (it

is a variant of graph colouring). They assume that each node knows its degree and an upper bound of the maximum degree Δ of the graph. They present in the beeping model BL a probabilistic algorithm which never stops and stabilises with a correct $O(\Delta)$ -interval coloring in $O(\log n)$ periods with high probability, where: n is the size of the graph, Δ its maximum degree and a period is Q time slots with $Q \geq \Delta$, thus it stabilises in $O(Q \log n)$ slots.

Emek and Wattenhofer introduce in [9] a model for distributed computations which resembles the beeping model: networked finite state machines (nFSM for short). This model enables the sending of the same message to all neighbours of a node; however it is asynchronous, the states of nodes belong to a finite set, and the set of messages is also finite. In the nFSM model they give a 2-MIS algorithm for graphs of size n using a set of messages of size 3 with a time complexity equal to $O(\log n^2)$.

Another similar model was considered in [15]. In this model, in every step, each node acts either as a transmitter or as a receiver. When a node is a transmitter, its message reaches all its neighbours. When it acts as a receiver, it receives a message if, and only if, exactly one of its neighbours transmits. The authors devise algorithms to broadcast efficiently in this model.

In [19], De Marco and Stachowiak consider a pretty close model: a number of stations is able to communicate by transmitting and listening to a shared channel, and a message is successfully delivered to all stations if, and only if, its source station is the only transmitter at a time. They discuss the impact of asynchrony in this model. They also consider and discuss how important the knowledge of the number of stations is, with a focus on the problem of contention resolution.

Concerning the counting problem, in the context of the classical message passing model, counting Monte Carlo algorithms for anonymous rings are presented in [13, 20]. The time complexity is $O(n)$. It is also investigated in dynamic networks [16, 17, 21]. In [16] nodes have unique identifiers, the size of messages is $O(\log n)$. Nodes know the size of the graph in $O(n^2)$ rounds. Networks are dynamic and anonymous in [17]; there is a leader in the network and the termination of the counting algorithm is detected by a heuristic. In [21] communication is asynchronous message passing, nodes have no identifiers and there exists a leader; counting algorithm which is presented produces a correct stabilizing solution that do not guarantee termination.

Summary. The paper is organised as follows. In Section 2 we provide definitions and basic observations regarding the beeping model. Section 3 presents our main results on the counting problem, including the lower bound, our optimal algorithm, and its time complexity analysis. Finally, Section 4 discusses the adaptations of our algorithm in other variants of the beeping model, including its emulation as a Monte Carlo algorithm in the weakest BL variant.

2. Network Model and Definitions

We consider a wireless network model and we follow definitions given in [1] and [8]. The network is anonymous: unique identifiers are not available to distinguish the processes. Communications are synchronous and encoded by a graph $G = (V, E)$ where the nodes V represent processes and the edges E represent pairs of processes that can hear each other. Since we focus here on one-hop networks, G is a complete graph. Time is divided into discrete synchronised time intervals called *slots* (following the usual terminology in wireless networks). All processes wake up and start computation at the same slot. In each slot, all processors act in parallel and either beep or listen. In addition, processors can perform an unrestricted amount of local computation in-between two slots (in effect, our algorithms require little computation).

In this paper, we consider several variants of beeping models:

- if a process beeps, there are two cases:

1. it cannot know whether another process beeps simultaneously, this case is denoted by B ;

2. it can distinguish whether it beeped alone or if at least one neighbour beeped concurrently, it is an internal collision; this case is called sender side collision detection, and it is denoted in this paper B_{cd} ;
- if a process listens, there are also two cases:
 1. it can distinguish between silence or the presence of at least one beep, this model is denoted L ;
 2. it can distinguish between silence or the presence of one beep or the presence of at least two beeps; in this case it is a peripheral collision, this model is denoted L_{cd} in this paper.

Finally, a beeping model is defined by choosing between B or B_{cd} and between L and L_{cd} .

Remark 1. In general, nodes are active or passive. When they are active they beep or listen; in the description of algorithms we say explicitly when a node beeps meaning that a non beeping active node listens.

The time complexity, also called *slot complexity*, is the maximum number of slots needed until every node has completed its computation. Our algorithms are typically structured into *phases*, each of which corresponds to a small (constant or logarithmic) number of slots. In the algorithm, we specify which one is the current slot by means of a **switch** instruction with as many **case** statements as there are slots in the phase. Phases repeat until some condition holds for termination.

Remark 2. An algorithm given in a beeping model induces an algorithm in the (synchronous) message passing model. Thus, given a problem, any lower bound on the round complexity in the message passing model also holds for slot complexity in the beeping model.

In this paper, results on graphs having n nodes are expressed with high probability (w.h.p. for short), meaning with probability $1 - o(n^{-1})$. We write $\log n$ for the binary logarithm of n .

Distributed Randomised Algorithm. A randomised (or probabilistic) algorithm is an algorithm which makes choices based on given probability distributions. A *distributed* randomised algorithm is a collection of local randomised algorithms (in our case, all identical).

A *Las Vegas* algorithm is a randomised algorithm which terminates with probability one, and always produces a correct result. A *Monte Carlo* algorithm is a randomised algorithm which terminates deterministically, but whose result may be incorrect with a certain probability. Intuitively, Las Vegas algorithms have uncertain execution time but certain result, while Monte Carlo does the reverse. Classical considerations on symmetry breaking in anonymous beeping networks (see for instance Lemma 4.1 in [1]), imply that:

Remark 3. There is no Las Vegas (and a fortiori no deterministic) algorithm in BL nor in BL_{cd} which allows a node to distinguish between an execution where it is isolated and one where it has exactly one neighbour.

From this remark we deduce that there is no Las Vegas counting algorithm in BL , which advocates the use of stronger models. In what follows, we consider the $B_{cd}L$ model, where a beeping node can detect simultaneous beeps by others. We give a Las Vegas algorithm in this model, which is then turned into a Monte Carlo algorithm in BL using emulation techniques.

The counting problem. We say that an algorithm solves the counting problem if and only if by the end of its execution, *every* node knows the size of the network.

3. One-Hop Network Size Computation in $B_{cd}L$

In the $B_{cd}L$ model, a node that beeps can detect simultaneous beeps by others, while a node that listens cannot distinguish between one and several beeps. Hence, only beeping nodes can detect collisions, which proves sufficient to solve the counting problem. In this section, we first prove that at least n slots are required to learn the number of nodes n ; then we present a Las Vegas algorithm which takes $O(n)$ slots to terminate (w.h.p.). Simulation results refine this complexity to less than $10n$ slots (*i.e.* less than 3.32 phases of 3 slots each).

3.1. A Lower Bound

The following lemma establishes that n slots are required for computing n in a complete graph.

Lemma 1. *Any Las Vegas beeping algorithm that counts the number of nodes in a complete graph K_n needs n slots.*

PROOF (BY CONTRADICTION). Let \mathcal{A} be such an algorithm and let $\mathcal{E}_{\mathcal{A}}$ be an execution of \mathcal{A} that terminates in less than n slots in the complete graph K_n . Then it holds that at least one node, say v , never beeped alone. Let $\mathcal{E}'_{\mathcal{A}}$ be another execution of \mathcal{A} , but this time in the complete graph K_{n+1} composed of the same nodes as before, plus v' . Let all the nodes behave as they did over $\mathcal{E}_{\mathcal{A}}$ and let v' act exactly like v . Since v never beeped alone in $\mathcal{E}_{\mathcal{A}}$, it does the same in $\mathcal{E}'_{\mathcal{A}}$ and so does v' as well, making both executions indistinguishable (two beeps are indistinguishable from three). Hence, the nodes in $\mathcal{E}'_{\mathcal{A}}$ terminate as in $\mathcal{E}_{\mathcal{A}}$, having counted n instead of $n + 1$, which is a contradiction. \square

Remark 4. The bound holds even when collision detection is also available for listening nodes (thus, in $B_{cd}L_{cd}$ and a fortiori in BL_{cd}), due to the same argument that two beeps are indistinguishable from three on the listener side.

3.2. An Algorithm for Computing One-Hop Network Size in $B_{cd}L$

We now propose a beeping algorithm that solves the counting problem in one-hop networks in $B_{cd}L$. Then we characterise its slot complexity both analytically and experimentally.

3.2.1. Informal description:

The basic idea is to have nodes beep alone as often as possible, so that they can detect it and become counted. This is achieved by means of three slots: In slot 1, the non-counted nodes beep with some probability. If they beeped alone, they re-beep in slot 2 to inform the other nodes and they can be counted. In slot 3, all non-counted nodes beep. If no one beeped in the third slot, then it means that every node is counted and the algorithm can terminate globally. The key ingredient is to use an adaptive probability $1/k$ to increase the chances to beep alone in the first slot, which is done by increasing the probability (decreasing k) whenever no one beeped, and decrease it (increasing k) whenever a collision is detected. As shown in the analysis section further below, this technique leads to a linear running time in the number of nodes.

The details are given on Algorithm 1. It describes the phase that repeats until termination is detected. Each **case** statement corresponds to a slot, and an extra **case** statement groups together various computation to be performed at the end of the phase: updating the count variable (line 19), detecting termination (line 21), adjusting the probability (lines 22 to 27).

Algorithm 1: Computing the size of a one-hop network in $B_{cd}L$.

```
1 Boolean counted  $\leftarrow$  false
2 Boolean terminated  $\leftarrow$  false
3 Integer k  $\leftarrow$  2
4 Integer size  $\leftarrow$  1
5 repeat
6   switch slot do
7     case 1 do
8       if counted = false then
9         | beep with probability  $\frac{1}{k}$ ; // contends
10      case 2 do
11        if I am the only one that beeped in slot 1 then
12          | beep; // wins the contest
13          | counted  $\leftarrow$  true
14      case 3 do
15        if counted = false then
16          | beep; // will contend in the next phase
17      end of phase:
18        if a neighbor beeped in slot 2 then
19          | size  $\leftarrow$  size + 1
20        if no one beeped in slot 3 then
21          | terminate  $\leftarrow$  true
22        if no one beeped in slot 1 then
23          | if ( $k > 2$ ) then
24            |  $k \leftarrow k - 1$ 
25        else
26          | if there was a collision in slot 1 then
27            |  $k \leftarrow k + 1$ 
28 until terminated = true;
```

3.2.2. Analysis of the Algorithm:

For the sake of analysis, it is easier to think in terms of phases rather than slots. The algorithm progresses every time a node beeps alone in the first slot of a phase. The probability that this happens in a phase depends upon the value of k and the number of nodes n' still contending in this phase. Intuitively, if k is too large, then it is likely that no node will beep in the first slot, while if k is too small, they will likely be several to do so. In fact, the probability of success is maximum when $k = n'$, which is what the algorithm attempts to maintain. The following fact is important:

Fact 3.1. *Uncounted nodes all have the same k .*

In the analysis below, we distinguish between the case that k is within “good bounds” and the case that it is not. We call a *bad* phase one in which k drifts out of these bounds or keeps drifting away (if it is already out). More formally, let $I = [n', 3n']$. A phase ϕ is a *bad phase* if:

- (i) at the start of ϕ , $k \leq n'$ or $k \geq 3n'$, and
- (ii) the phase moves k farther away from I (or k remains unchanged at 2).

We prove the following lemma:

Lemma 2. *The probability of a phase to be a bad one is upper bounded by 0.4.*

PROOF. Let ϕ be a phase. Assume that k verifies the conditions in (i), i.e., $k \leq n'$ or $k \geq 3n'$. To prove the lemma, we study both cases separately.

- $k \leq n'$, then k will decrease iff no node beeps. This happens with probability $p = (1 - \frac{1}{k})^{n'}$, which is maximum when $k = n'$ and always less than e^{-1} , which is less than 0.4.
- $k \geq 3n'$, then k will increase iff at least two nodes beep. This is a particular case of (thus is less likely than) having at least *one* node beep, which happens with probability $p = 1 - (1 - \frac{1}{k})^{n'}$. This formula is maximum when $k = 3n'$ and it is always less than $\frac{1}{3}$, which is less than 0.4. \square

We also have the following lemma:

Lemma 3. *Let ϕ be a phase. If $k \in I$ at the start of ϕ , then with probability at least $1/5$, a node will be counted.*

PROOF. The probability that exactly one node beeps (and thus is counted) in phase ϕ is given by $p = \frac{n'}{k} (1 - \frac{1}{k})^{n'-1}$. Indeed, let v be a vertex that is not yet counted. The probability for v to beep is $\frac{1}{k}$, and the probability that *all* the other $n' - 1$ nodes do not beep is given by $(1 - \frac{1}{k})^{n'-1}$. The probability p is a decreasing function of k for $k > n'$ and thus it is lower bounded when taking $k = 3n'$.

Using the lower bound $(1 - x)^n \geq 1 - nx$, we obtain:

$$p \geq \frac{1}{3} \left(1 - \frac{n' - 1}{3n'} \right) \geq \frac{2}{9}.$$

\square

We now give the main result of this section:

Theorem 4. *Let G be a one-hop network of size n . The execution of Algorithm 1 in G ends in at most $55n$ phases ($165n$ slots), with high probability.*

PROOF. We denote by T a number of phases, and by T_I the number of phases where k is in the interval I , and by B_T the number of bad phases (both within the first T phases). We have the following fact:

Fact 3.2. *$T_I \geq T - 4n - 2 \times B_T$, where $4n$ is an upper bound on the number of “external” phases, that is, $n - 2$ initial phases from $k = 2$ to $k = n$ and $3n - 3$ phases where k decreases from $3n$ to 3 . The factor 2 for B_T corresponds to the fact that bad phases have symmetrical “good phases” where k comes closer to the good interval (and their number is at most B_T).*

Note that B_T is dominated by a binomial random variable $BIN(T, 0.4)$ with parameters T and 0.4. On the other hand, we have the following formulation of the Chernoff bound (see [11]):

If $BIN(n, p)$ is a binomial random variable with parameters n and p then

$$\Pr(|BIN(n, p) - np| > a) < 2e^{-a^2/3np},$$

for any $0 < a < np$.

Then, taking $T = 55n$, we obtain:

$$\begin{aligned} \Pr(B_T \geq 23n) &= \Pr(B_T - 0.4 \times T \geq n) \\ &\leq \Pr(|B_T - 0.4 \times T| \geq n) \\ &\leq 2 \times e^{-n^2/3 \times 0.4 \times T} = 2 \times e^{-n/66}. \end{aligned} \tag{1}$$

This shows that if $T = 55n$ then, w.h.p., the number of bad phases is upper bounded by $23n$. Thus, using Fact 3.2, we obtain that, w.h.p., the number of phases inside interval I is at least $5n$. Since each phase inside I counts a node with probability at least $1/5$, we can use the same arguments to prove that, w.h.p., the number of nodes counted during $T = 55n$ phases is at least n . Which ends the proof. \square

From Lemma 1:

Corollary 5. *Algorithm 1 is optimal up to a constant factor.*

3.2.3. Experimental Results:

The algorithm was implemented and run on complete graphs with size n ranging over all powers of 2 from 8 to 512 (inclusive). For each value of n , we performed 10000 runs and measured the number of phases before the algorithm terminates.¹ A linear regression with Gnuplot gives us an average number of phases of $3.3197n$ (with very small regression error of $\pm 0.03074\%$). Since there are 3 slots per phases, the expected slot complexity in practice is less than $10n$.

4. Adaptation of the algorithm in other beeping models

This section explores the possible adaptations of Algorithm 1 in other variants of the beeping model, namely in $B_{cd}L_{cd}$ (the strongest), in BL (the weakest), and in BL_{cd} . We show that the extra power available in $B_{cd}L_{cd}$ allows us to reduce the slot complexity by a constant factor. Adaptation to BL comes at the price of sacrificing certainty (the adaptation is a Monte Carlo algorithm), it consists in emulating B_{cd} beeps using techniques from [4]. Finally, we briefly discuss the case of BL_{cd} , in which naive adaptations do not seem to work.

4.1. Adaptation of the algorithm in $B_{cd}L_{cd}$

In the $B_{cd}L_{cd}$ variant, both beeping and listening nodes can detect if several beeps occur simultaneously. In our algorithm, the purpose of the second slot is for a node to inform others that it was the only one to beep. This step is no longer necessary since listening nodes can detect it from the first slot. As a result, the phases of the algorithm can be simplified from 3 to 2 slots and the expected running time of the algorithm decreases by one third. (According to our experimentations, this would thus decrease the expected number of slots from less than $10n$ to less than $6.67n$.)

4.2. Adaptation of the algorithm in BL

As discussed in Section 2, no algorithm can solve the counting problem with guaranteed result in BL , whether it be deterministic or Las Vegas. However, we can adapt it into a Monte Carlo algorithm (uncertain result) using emulation techniques from [4] and [1]. We describe this adaptation here and analyse the resulting uncertainty. We show that this uncertainty can be bounded by any constant threshold if the nodes know an upper bound on the size of the network. Otherwise, we must settle for a certainty bound that depends on the size of the network (which is unknown).

The adaptation consists in replacing the **beep** instruction in Line 9 of Algorithm 1 by a call to the procedure **EmulateB_{cd}inBL**. The other **beep** instructions in Algorithm 1 need not be replaced. Procedure **EmulateB_{cd}inBL** emulates a B_{cd} beep (beep with collision detection) by means of several basic B beeps. It is itself organised into phases of two slots. In each of these phases, the node beeps in one of the two slots, chosen uniformly at random, and it listens in the other. The intuitive idea is that if several nodes are to execute this procedure simultaneously, then they will eventually beep in different slots and thus detect it. Unfortunately, as long as the number of phases is finite, there is a non-zero probability that

¹Source code available upon request.

these nodes always beep simultaneously in the procedure, which is why the resulting algorithm is Monte Carlo.

The detail of the procedure is given in Algorithm 2. For technical reasons, we have the nodes generate their random bits (called signature) beforehand and use the same every time the emulation procedure is called. Besides using less random bits, this improves the probability of success by avoiding the extra union bound due to repeating a successful emulation several times. Let r be the number of round in each emulation (its value is discussed later on), we denote by s the signature of a node, that is, the word formed by r bits generated uniformly at random.

Algorithm 2: A Procedure to emulate a B_{cd} in the BL model.

```

1 Procedure Emulate $B_{cd}$ in $BL$ (IN:  $s$  : word of bits associated to the vertex; OUT:
   collision : boolean)
2   Boolean collision  $\leftarrow$  false;
3   Integer  $i \leftarrow 0$ ;
4   repeat
5     switch slot do
6       case 1 do
7         if  $s[i] = 0$  then beep;
8         else listen;
9       case 2 do
10        if  $s[i] = 1$  then beep;
11        else listen;
12      end of phase:
13      if someone beeped in a different slot than I did then
14        collision  $\leftarrow$  true
15     $i \leftarrow i + 1$ 
16 until  $i = r$ ;
17 End Procedure

```

4.2.1. Analysis of the Algorithm:

We say that the procedure is *correct* if it detects when several nodes execute it simultaneously. The probability that it fails is maximum when two nodes only execute it, which corresponds to $1/2^r$. Hence, the probability that the procedure is not correct is upper bounded by $1/2^r$. The results in Lemma 6 are obtained by replacing r with different values.

Lemma 6. *For any $\varepsilon > 0$ and for any $n > 0$:*

1. *if $r = \lceil \log(\frac{n}{\varepsilon}) \rceil$, then the procedure is correct with probability at least $1 - \frac{\varepsilon}{n} > 1 - \varepsilon$,*
2. *if $r = \lceil \log(\frac{1}{\varepsilon}) \rceil$, then, for any node v , the procedure is correct on v with probability at least $1 - \varepsilon$,*
3. *if $r = \lceil 2 \log(n) \rceil$, then, the procedure is correct w.h.p.*

Since n is not known a priori, only the second result can be used effectively by the algorithm. From Lemma 6 we obtain:

Theorem 7. *For any graph G of size n and any $0 < \varepsilon < 1$: If $r = \lceil \log(\frac{1}{\varepsilon}) \rceil$, each node v computes the size of the network in $O(n \log(\frac{1}{\varepsilon}))$ slots, and the result is correct with probability at least $1 - \varepsilon$.*

Now, if the nodes know a polynomial upper bound $N = O(n^{\Theta(1)})$ on the size of the network, then the other two results from Lemma 6 can be used effectively by the algorithm as follows.

Proposition 8. *For any graph G of size n and any $0 < \varepsilon < 1$:*

- *if $r = \lceil \log(\frac{N}{\varepsilon}) \rceil$, the running time of the main algorithm is $O(n \log(\frac{n}{\varepsilon}))$, and the result is correct with probability at least $1 - \varepsilon$.*
- *if $r = \lceil 2 \log(N) \rceil$, the running time of the main algorithm is $O(n \log n)$, and the result is correct with probability $1 - o(\frac{1}{n})$.*

The latter is particularly relevant, as it implies that the correct result can be obtained with high probability. Furthermore, assuming that an upper bound is known is reasonable in practical scenarios. For instance, when a sensor network is deployed, one might want to learn afterward how many sensors (among the initial amount) successfully started to operate. Similarly, when sensors use sleeping patterns (some are active, others sleep to save energy), this makes it possible to learn the number of active ones.

4.3. Adaptation of the algorithm in BL_{cd}

Now that we have seen how to adapt the algorithm in the strongest ($B_{cd}L_{cd}$) and the weakest (BL) beeping models, we turn our attention to the symmetrical BL_{cd} model. In this variant, only listening nodes can detect simultaneous beeps. Algorithm 1 can be adapted into this model. The basic idea is that instead of having a node inform the others when it detects that it beeped alone, we could do the reverse and have listeners inform the beeping node when it beeped alone. Hence, the first and last slot of each phase remain as in Algorithm 1. The second slot is split into two slots. In the first of these (new slot 2), any node that was listening in slot 1 beeps. This is to detect pathetic cases where *all* the nodes beep at the same time. In the second of these slots (new slot 3), the listening node(s) beep if they detected a collision in slot 1. Based on this beep, if a node beeped in slot 1, it detects that it was alone (other nodes also know it, since they were all listening in slot 1). The only exception, which new slot 2 can't avoid, is the special case that $n = 1$ from the start. In this case, the node cannot decide whether it is alone to beep (and thus must count itself) or several nodes are beeping. This problem does not exist if $n \geq 2$, even when a single node remains to be counted in the end (since counted nodes keep on listening and participating until termination).

As for the slot complexity, the number of phases still depends solely on the drawings of slot 1 and thus remains unchanged. The fact that every phase is now composed of 4 slots makes the running time increase by one third. (According to our experimentations, this would thus increase the expected number of slots from less than $10n$ to less than $13.34n$.)

5. Concluding remarks

In this paper, we presented an algorithm for counting the number of processes in one-hop beeping networks. Due to natural limitations of the beeping model, we observed that no algorithm can solve the problem with guaranteed result unless some form of collision detection is available (either for beeping nodes or for listening nodes). We presented an algorithm in the case that beeping nodes can detect collision ($B_{cd}L$), and prove that it runs in linear time in the number of nodes. We proved that this is optimal up to a constant factor and estimated the expected value of this factor experimentally. Interestingly, the algorithm can be adapted in $B_{cd}L_{cd}$ with a small gain in slot complexity (1/3 less slots) and in BL_{cd} with a small overhead (4/3 more slots). Its adaptation in BL is not as direct, but we proved that the algorithm can be emulated in this variant by a Monte Carlo algorithm which runs in a larger, though reasonable amount of time ($\log n$ more slots). Whether this is optimal for BL is an open question. Another open question is whether, and to which extent, the number of nodes in the network could be efficiently approximated by a randomised algorithm.

References

- [1] Y. Afek, N. Alon, Z. Bar-Joseph, A. Cornejo, B. Haeupler, and F. Kuhn. Beeping a maximal independent set. *Distributed Computing*, 26(4):195–208, 2013.
- [2] A. F. Anta and M. A. Mosteiro. Contention resolution in multiple-access channels: k -selection in radio networks. In *Computing and Combinatorics, 16th Annual International Conference, COCOON 2010, Nha Trang, Vietnam, July 19-21, 2010. Proceedings*, pages 378–388, 2010.
- [3] M. Bienkowski, M. Klonowski, M. Korzeniowski, and D. R Kowalski. Randomized mutual exclusion on a multiple access channel. *Distributed Computing*, (DOI 10.1007/s00446-016-0265-z), 2016.
- [4] A. Casteigts, Y. Métivier, J. M. Robson, and A. Zemmari. Design patterns in beeping algorithms: Examples, emulation, and analysis. *Inf. Comput.*, 264:32–51, 2019.
- [5] B. S. Chlebus. Randomized communication in radio networks. In P.M. Pardalos, S. Rajasekaran, J.H. Reif, and J.D.P. Rolim, editors, *Handbook of Randomized Computing*, pages 401–456. Kluwer Academic, 2001.
- [6] B. S. Chlebus, G. De Marco, and D. R. Kowalski. Scalable wake-up of multi-channel single-hop radio networks. *Theor. Comput. Sci.*, 615:23–44, 2016.
- [7] J. Collier, N. Monk, P. Maini, and J. Lewis. Pattern formation by lateral inhibition with feedback: a mathematical model of delta-notch intercellular signalling. *Journal of Theoretical Biology*, 183(4):429–446, 1996.
- [8] A. Cornejo and F. Kuhn. Deploying wireless networks with beeps. In *Proc. of 24th International Symposium on Distributed Computing (DISC)*, pages 148–162, 2010.
- [9] Y. Emek and R. Wattenhofer. Stone age distributed computing. In *Proc. of 33rd Symposium on Principles of Distributed Computing (PODC)*, pages 137–146, 2013.
- [10] S. Gilbert and C. Newport. The computational power of beeps. In *Proc. of 29th International Symposium on Distributed Computing (DISC)*, 2015.
- [11] M. Habib, C. McDiarmid, J. Ramirez-Alfonsin, and B. Reed, editors. *Probabilistic methods for algorithmic discrete mathematics*, volume 16 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1998.
- [12] B. Huang and Th. Moscibroda. Conflict resolution and membership problem in beeping channels. In *Proc. of 27th International Symposium on Distributed Computing (DISC)*, pages 314–328, 2013.
- [13] A. Itai and M. Rodeh. Symmetry breaking in distributed networks. *Inf. Comput.*, 88(1):60–87, 1990.
- [14] D. R. Kowalski. On selection problem in radio networks. In *Proceedings of the Twenty-Fourth Annual ACM Symposium on Principles of Distributed Computing, PODC 2005, Las Vegas, NV, USA, July 17-20, 2005*, pages 158–166, 2005.
- [15] D. R. Kowalski and A. Pelc. Time of deterministic broadcasting in radio networks with local knowledge. *SIAM J. Comput.*, 33(4):870–891, 2004.
- [16] F. Kuhn, N. A. Lynch, and R. Oshman. Distributed computation in dynamic networks. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 513–522, 2010.

- [17] G. A. Di Luna, S. Bonomi, I. Chatzigiannakis, and R. Baldoni. Counting in anonymous dynamic networks: An experimental perspective. In *Algorithms for Sensor Systems - 9th International Symposium on Algorithms and Experiments for Sensor Systems, Wireless Networks and Distributed Robotics, ALGOSENSORS 2013, Sophia Antipolis, France, September 5-6, 2013, Revised Selected Papers*, pages 139–154, 2013.
- [18] G. De Marco, T. Jurdzinski, M. Rozanski, and G. Stachowiak. Subquadratic non-adaptive threshold group testing. In Ralf Klasing and Marc Zeitoun, editors, *Fundamentals of Computation Theory - 21st International Symposium, FCT 2017, Bordeaux, France, September 11-13, 2017, Proceedings*, volume 10472 of *Lecture Notes in Computer Science*, pages 177–189. Springer, 2017.
- [19] G. De Marco and G. Stachowiak. Asynchronous shared channel. In *Proc. of 37rd Symposium on Principles of Distributed Computing (PODC), Washington, DC, USA, July 25-27, 2017*, pages 391–400, 2017.
- [20] Y. Métivier, J. M. Robson, and A. Zemmari. Analysis of fully distributed splitting and naming probabilistic procedures and applications. *Theor. Comput. Sci.*, 584:115–130, 2015.
- [21] O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Naming and counting in anonymous unknown dynamic networks. In *Stabilization, Safety, and Security of Distributed Systems - 15th International Symposium, SSS 2013, Osaka, Japan, November 13-16, 2013. Proceedings*, pages 281–295, 2013.
- [22] S. Navlakha and Z. Bar-Joseph. Distributed information processing in biological and computational systems. *Commun. ACM*, 58(1):94–102, 2015.
- [23] D. Peleg. *Distributed computing - A Locality-sensitive approach*. SIAM Monographs on discrete mathematics and applications, 2000.
- [24] J. Schneider and R. Wattenhofer. What is the use of collision detection (in wireless networks)? In *Proc. of 24th International Symposium on Distributed Computing (DISC)*, pages 133–147, 2010.
- [25] A. Scott, P. Jeavons, and L. Xu. Feedback from nature: an optimal distributed algorithm for maximal independent set selection. In *Proc. of 33rd Symposium on Principles of Distributed Computing (PODC)*, pages 147–156, 2013.