

Published in: *Fundamenta Informaticae*, IOS Press.

© 2017 IOS Press. Personal use of this material is permitted. Permission from IOS Press must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

The Version of Record is available online at: <http://dx.doi.org/10.3233/FI-2017-1537>

Naming a Channel with Beeps ^{*}

Bogdan S. Chlebus[†] Gianluca De Marco[‡] Muhammed Talo[†]

Abstract

We consider a communication channel in which the only possible communication mode is transmitting beeps. A beep is a signal which, when transmitted by a station, reaches all the other stations instantaneously. Stations are anonymous, in that they do not have any individual identifiers. The algorithmic goal is to assign names to the stations in such a manner that the names make a contiguous segment of positive integers starting from 1. We develop a Las Vegas naming algorithm, for the case when the number of stations n is known, and a Monte Carlo algorithm, for the case when the number of stations n is not known. The given randomized algorithms are provably optimal with respect to the expected time $\mathcal{O}(n \log n)$, the expected number of random bits used $\mathcal{O}(n \log n)$, and the probability of error.

Key words: beeping channel, anonymous network, randomized algorithm, Las Vegas algorithm, Monte Carlo algorithm, lower bound.

^{*}The work of the first author was supported by the National Science Foundation under Grant 1016847.

[†]Department of Computer Science and Engineering, University of Colorado Denver, Denver, Colorado 80217, USA.

[‡]Dipartimento di Informatica, Università degli Studi di Salerno, Fisciano, 84084 Salerno, Italy.

1 Introduction

The considered model of communication is anonymous channel with beeping. There are a number of stations attached to the channel that are devoid of any identifying features they can refer to during communication. The goal is to have the stations assign names to themselves by executing a distributed communication algorithm.

The stations communicate among themselves in synchronous rounds. All of them start together in the same round. The channel provides a binary feedback to all the attached stations: when no stations transmit then nothing is sensed on the communication medium, and when some station does transmit then every station detects a beep.

A beeping channel resembles multiple-access channels, in that it can be interpreted as a single-hop radio network. The difference between the two models is in the feedback they provide. Multiple access channels come in two variants: with and without collision detection. A multiple access channel without collision detection provides the following binary feedback: when exactly one station transmits then the transmitted message is heard by every station, and otherwise, when either no station transmits or multiple stations do, then this results in each station sensing silence. A multiple access channel with collision detection provides the following ternary feedback: silence occurs when no station transmits, a message is heard when exactly one station transmits, and collision occurs when multiple stations transmit simultaneously, which prevents any of the transmitted messages to be heard and is sensed as distinct from silence.

A channel with beeping has its communication capabilities restricted only to carrier sensing, without even the functionality of transmitting specific bits as messages. The only apparent mode of exchanging information on such a synchronous channel with beeping is to suitably encode it by sequences of beeps and silences.

Modeling communication by a mechanism as limited as beeping has been motivated by diverse aspects of communication and distributed computing. Beeping provides a detection of collision on a transmitting medium by sensing it. Communication by only carrier sensing can be placed in a general context of investigating wireless communication on the physical level and modeling interference of concurrent transmissions, of which the signal-to-interference-plus-noise ratio (SINR) model is among the most popular and well studied; see [32, 37, 47]. Beeping is then a very limited mode of wireless communication, with feedback in the form of either a simplest signal, which is identical to interference, or lack thereof.

Another motivation comes from biological systems, in which agents exchange information in a distributed manner, while the environment severely restricts how such agents communicate; see [2, 44, 48].

Finally, communication with beeps belongs to the area of distributed computing by weak devices, where the involved agents have restricted computational and communication capabilities. In this context, the devices are modeled as finite-state machines that communicate asynchronously by exchanging states or messages from a finite alphabet. Examples of this approach include the “population-protocols” model introduced by Angluin et al. [5], (see also [6, 8, 41]), and the “stone-age” distributed computing model proposed by Emek and Wattenhoffer [24].

Previous work. The model of communication by discrete beeping was introduced by Cornejo and Kuhn [17], who considered a general-topology wireless network in which nodes use only carrier

sensing to communicate, and developed algorithms for node coloring. They were inspired by “continuous” beeping studied by Degeysys et al. [22] and Motskin et al. [42], and by the implementation of coordination by carrier sensing given by Flury and Wattenhofer [27].

Afek et al. [1] considered the problem to find a maximal independent set of nodes in a distributed manner when the nodes can only beep, under additional assumptions regarding the knowledge of the size of the network, waking up the network by beeps, collision detection among concurrent beeps, and synchrony. Brandes et al. [11] studied the problem of randomly estimating the number of nodes attached to a single-hop beeping network. Czumaj and Davies [19] approached systematically the tasks of deterministic broadcasting, gossiping, and multi-broadcasting on the bit level in general-topology symmetric beeping networks. In a related work, Hounkanli and Pelc [35] studied deterministic broadcasting in asynchronous beeping networks of general topology with various levels of knowledge about the network. Förster et al. [28] considered leader election by deterministic algorithms in general multi-hop networks with beeping. Gilbert and Newport [31] studied the efficiency of leader election in a beeping single-hop channel when nodes are state machines of constant size with a specific precision of randomized state transitions. Huang and Moscibroda [36] considered the problems of identifying subsets of stations connected to a beeping channel and compared their complexity to those on multiple-access channels. Yu et al. [51] considered the problem of constructing a minimum dominating set in networks with beeping.

Related work. A beeping channel is related to multiple access channels and single-hop radio networks [14]. The similar task to the one considered in this paper, to have all stations transmit successfully, once by each station among the n stations, can be accomplished in expected time $\mathcal{O}(n)$ on a multiple access channel. This variant of the “contention resolution problem” for such channels has been the topic of continuous interest, since multiple access channels started to be investigated [7, 13, 16, 20, 21, 33, 34, 49]. In view of the lower bound $\Omega(n \log n)$ on time of naming for beeping channels, given in this paper, resolving contention in the model of a beeping channel is demonstrably more costly, in that it involves the additional logarithmic multiplicative overhead. This can be attributed to the key difference between the two models, in that a multiple access channel differentiates between single and multiple transmissions by its feedback, while a beeping channel returns a beep in both cases.

Networks of nodes communicating by beeps share common features with radio networks with collision detection. Ghaffari and Haeupler [29] gave efficient leader election algorithm by treating collision detection as “beeping” and transmitting messages as bit strings. Their approach by way of “beep waves” was adopted to broadcasting in networks with beeping by Czumaj and Davies [19]. In a related work, Ghaffari et al. [30] developed randomized broadcasting and multi-broadcasting in radio networks with collision detection.

Communication in anonymous networks was first considered by Angluin [4], who showed that randomization was necessary to name nodes in suitably symmetric communication settings. The work that followed was either on specific network topologies, like rings or hypercubes, or on problems in general message-passing systems. The early work on anonymous rings was done by Attiya et al. [10], while anonymous hypercubes were studied first by Kranakis and Krizanc [38]. Algorithmic problems in anonymous networks of general topologies were considered by Afek and Matias [3] and Schieber and Snir [46]. Angluin et al. [6] considered self-stabilizing protocols for anonymous asynchronous agents.

Anonymous systems of asynchronous processes communicating by shared memory were studied

by Lipton and Park [40], Egecioglu and Singh [23], Panconesi et al. [45], Kutten et al. [39], and Buhrman et al. [12], among others. The paper by Chlebus et al. [15] is an example of work on anonymous distributed computing with shared read-write memory, in the case of synchrony.

Impossibility results and lower bounds are surveyed by Attiya and Ellen [9] and Fich and Ruppert [26]. Yao’s minimax principle was formulated by Yao [50], with an alternative direct proof given by Fich et al. [25]; see also the book by Motwani and Raghavan [43] for a presentation of applications.

2 Technical preliminaries

A beeping channel is a network consisting of some n stations connected to a communication medium that provides a limited feedback. When no stations transmit at a point in time then nothing is sensed on the communication medium, and when some station does transmit then every station detects a beep.

We consider synchronous beeping channels, in the sense that an execution of a communication algorithm is partitioned into consecutive rounds. All the stations start execution together. In each round, a station may either beep or pause. When some station beeps in a round, then each station hears the beep in this round, otherwise all the stations receive silence as feedback in this round. When multiple stations beep together in a round then we call this a *collision*.

We say that a parameter of a communication network is *known* when it can be used in codes of algorithms. The relevant parameter used in this paper is the number of stations n . We consider two cases, in which either n is known or it is not.

The stations attached to the channel are anonymous in that they do not have access to any predetermined identifiers that can be referred to when executing a communication algorithm. The stations are distinguishable by an external observer, and we refer to such an observer’s “name” when referring to a station to distinguish it from the other stations.

Randomized algorithms use random bits, understood as outcomes of tosses of a fair coin. All different random bits used by our algorithms are considered stochastically independent from one another.

Randomized algorithms are often categorized as Las Vegas or Monte Carlo, see [43]. We use these terms in the following meaning. An algorithm is *Las Vegas* when it terminates almost surely and it returns a correct output upon termination. An algorithm is *Monte Carlo* when it terminates almost surely but an error may occur when the algorithm terminates; the probability of error converges to zero with the number of communicating agents converging to infinity. An “error” occurs when the output generated by the algorithm does not satisfy all the properties, as specified in the correctness requirements. Our naming algorithms have as their goal to assign unique identifiers to the stations, moreover we want names to be integers in the contiguous range $\{1, 2, \dots, n\}$; we denote this range as $[n]$. The Monte Carlo naming algorithm that we develop assigns names that completely fill in an interval of integers of the form $[k]$, for $k \leq n$. If it happens that $k < n$ then this means that there are duplicate integers assigned as names, which is the only form of error that can possibly occur.

We will use a procedure to detect collisions, called DETECT-COLLISION, whose pseudocode is in Figure 1. The procedure is executed by a group of stations, and they all start their executions simultaneously. The procedure takes two rounds. Each of the participating stations simulates the

Procedure DETECT-COLLISION

```
tossv ← outcome of a random coin toss

if tossv = heads                               /* first round */
    then beep else pause

if tossv = tails                               /* second round */
    then beep else pause

return (a beep was heard in each of the two rounds)
```

Figure 1: A pseudocode for a station v . The procedure takes two rounds to execute. It detect a collision and returns “true” when a beep is heard in each of the rounds, otherwise it does not detect a collision and returns “false.”

toss of a fair coin, with the outcomes independent among the participating stations. Depending on the outcome of a toss, a station beeps either in the first round or the second one of the allocated pair of rounds. A collision is detected only when two consecutive beeps are heard.

Lemma 1 *If k stations perform m time-disjoint calls of procedure DETECT-COLLISION, each station participating in exactly one call, then collision is not detected in any of these calls with probability 2^{-k+m} .*

Proof: Consider a call of DETECT-COLLISION performed concurrently by i stations, for $i \geq 1$. We argue by “deferred decisions.” One of these stations tosses a coin and determines its outcome X . The other $i - 1$ stations participating concurrently in this call also toss their coins. Here we have $i - 1 \geq 0$, so it could be the case that there is no such a station; if $i = 1$ then collision cannot be detected. The only possibility not to detect a collision in case of $i - 1 > 0$ is for each among these $i - 1$ stations to also produce X . This happens with probability 2^{-i+1} in one call. The probability of producing only **false** during the m calls is the product of these probabilities. When we multiply them out over m instances of the procedure being performed, then the outcome is 2^{-k+m} , because numbers i sum up to k and the number of factors is m . \square

Pseudocode conventions and notations. We present algorithms formally using pseudocodes, like the one in Figure 1. The convention we apply is that all the stations stay synchronized to execute the same line of code in any given round. This means that if a conditional statement is executed, like an if-then instruction or a while-loop, then a station that pauses through it still counts rounds to stay synchronized, as if it were actively executing the code.

Each station has its private copy of every among the variables used in the pseudocode. When the values of these copies may vary across the stations, then we add the station’s external-observer “name” as a subscript of the variable’s name to emphasize that, and otherwise, when all the copies of a variable stay equal at all times across all the stations then no subscript is used.

We use the notation $\lg x$ to denote the logarithm to the base 2 of x . The base of logarithms is not specified when this is not necessary, like in an asymptotic notation $\mathcal{O}(\log x)$.

3 Lower bounds and impossibilities

We will show impossibility results to justify methodological approach to naming algorithms we apply, and use lower bounds on performance metrics for such algorithms to argue about the optimality of the algorithms, as developed in subsequent sections. We begin with an observation, formulated as Proposition 1, that if the system is sufficiently symmetric then randomness is necessary to break symmetry. The given argument is standard and is given for completeness sake; see [4, 9, 26].

Proposition 1 *There is no deterministic naming algorithm for a synchronous channel with beeping with at least two stations, in which all stations are anonymous, such that it eventually terminates and assigns proper names.*

Proof: We argue by contradiction. Suppose that there exists a deterministic algorithm that eventually terminates with proper names assigned to the anonymous stations. Let all the stations start initialized to the same initial state. The following invariant is maintained in each round: the internal states of the stations are all equal.

We proceed to show this invariant by induction on the round numbers. The base of induction is satisfied by the assumption about the initialization. For the inductive step, we assume that the stations are in the same state, by the inductive hypothesis. Then either all of them pause or all of them beep in the next round, as determined by the state. The effect is that either all of them hear their own beep or all of them pause and hear silence. This results in the same internal state transition, which completes the inductive step.

When the algorithm eventually terminates, then each station assigns to itself the identifier determined by its state. The identifier is the same in all stations because their states are the same, by the invariant. This violates the desired property of names to be distinct, because there are at least two stations with the same name. \square

Proposition 1 justifies developing randomized naming algorithms. We continue with “entropy” arguments; see the book by Cover and Thomas [18] for a systematic exposition of information theory.

An execution of a randomized naming algorithm coordinates and translates random bits into names. This same amount of entropy needs to be processed/communicated on the channel, by the Shannon’s noiseless coding theorem. An analogue of the following Proposition 2 was stated in [15] for the model of synchronized processors communicating by reading and writing to shared memory. The proof is similar, we include it here for completeness sake.

Proposition 2 *If a randomized naming algorithm for a channel with beeping is executed by n anonymous stations and is correct with probability p_n then it requires $\Omega(n \log n)$ random bits in total to be generated with probability at least p_n . In particular, a Las Vegas naming algorithm uses $\Omega(n \log n)$ random bits almost surely.*

Proof: Let us refer to the n anonymous stations by the identifiers known only to an external observer; we call them the *unknown names*. The names given by an execution of a randomized naming algorithms are referred to as the *given names*.

A permutation of the unknown names is produced by ordering the stations by the given names, with the respective probability. The algorithm determines a probability distribution on these permutations, conditional on successful naming. The conditional distribution is uniform because of the symmetry: all the processors execute the same code, without explicit private identifiers, and all use random bits produced by tossing fair coins. This determines a probability space of $n!$ elementary events, each identified by a permutation of the unknown identifiers. Each such an elementary event occurs with the conditional probability $1/n!$. The Shannon entropy of this random variable equals $\lg(n!) = \Theta(n \log n)$. Setting on one permutation, by way of assigning names, requires the amount $\lg(n!) = \Theta(n \log n)$ of entropy, by the Shannon's noiseless coding theorem. The source of this entropy are random bits generated by the stations. \square

One round of an execution of a naming algorithm allows the stations that do not transmit to learn at most one bit, because, from the perspective of these stations, a round is either silent or there is a beep. Intuitively, the running time is proportional to the amount of entropy that is needed to assign names. This intuition leads to Proposition 3. In its proof, we combine the Shannon's entropy [18] with the Yao's principle [50].

Proposition 3 *A randomized naming algorithm for a beeping channel with n stations operates in $\Omega(n \log n)$ expected time, when it is either a Las Vegas algorithm or a Monte Carlo algorithm with the probability of error smaller than $1/2$.*

Proof: We apply the Yao's minimax principle to bound the expected time of a randomized algorithm by the distributional complexity of naming. We consider Las Vegas algorithms first.

A randomized algorithm using strings of random bits generated by stations can be considered as a deterministic algorithm \mathcal{D} on all possible assignments of such (sufficiently long) strings of bits to stations as their inputs. We consider assignments of strings of bits of an equal length with the uniform distribution among all such assignments of strings of the same length. On a given assignment of input strings of bits to stations, the deterministic algorithms either assigns proper names or fails to do so. A failure to assign proper names with some input is interpreted as the randomized algorithm continuing to work with additional random bits, which comes at an extra time cost. This is justified by a combination of two factors. One is that the algorithm is Las Vegas, so it halts almost surely and with a correct output. Another is that the probability to assign a specific finite sequence as a prefix of a used sequence of random bits is positive. So if starting with a specific string of bits, as a prefix of a possibly longer needed string, would mean inability to terminate with a positive probability, then the naming algorithm would not be Las Vegas.

The common length of these input strings is a parameter. We consider all sufficiently large positive integer values for this parameter such that their exist strings of random bits of this length resulting in assignments of proper names. For a given length of input strings, we remove input assignments that do not result in assignment proper names and consider a uniform distribution of the remaining inputs. This is the same as the uniform distribution conditional on the algorithm terminating with input strings of bits of a given length.

Let us consider such a deterministic algorithm \mathcal{D} assigning names, and using strings of bits at

stations as inputs, these strings being of a fixed length, assigned under a uniform distribution for this length, and such that they result in termination. An execution of this algorithm produces a “feedback sequence of bits,” where we translate the feedback from the channel round by round, say, with symbol 1 representing a beep and symbol 0 representing silence. Each such a feedback sequence translates into a specific assignment of names, because this is the only information shared by all the stations. These feedback sequences also have a uniform distribution, by the same symmetry argument as used in the proof of Proposition 2. The expected length of such a feedback sequence is the expected time of algorithm \mathcal{D} . The expected time of algorithm \mathcal{D} is therefore at least $\lg n! = \Omega(n \log n)$, by the Shannon’s noiseless coding theorem. We conclude that, by the Yao’s principle, the original randomized Las Vegas algorithm has the expected time that is $\Omega(n \log n)$.

A similar argument, by the Yao’s principle, applies to a Monte Carlo algorithm that is incorrect with a constant probability smaller than $1/2$. The only difference in the argument is that when a given assignment of input string bits does not result in a proper assignment of names, then either the algorithm continues to work with more bits for an extra time, or terminates with error. \square

Next, we consider the case when the number of stations n is unknown. The following Proposition 4 is about the inevitability of error in such situations, which is due to the fact that we cannot compare the number of randomly assigned names to the (unknown) number of computing/communicating agents. Intuitively, when two computing/communicating agents generate the same string of bits, then their actions are the same, and so they get the same name assigned. In other words, we cannot distinguish the case when there is only one such an agent present from cases when at least two of them work in unison.

Proposition 4 was formulated in [15] for the model of synchronous parallel random-access machine, in which processors communicate among themselves by reading from and writing to shared read-write registers. It applies to a synchronous beeping channel, when we understand actions of stations as either beeping or pausing in a round; we include the proof to make the exposition self-contained.

Proposition 4 ([15]) *For an unknown number of stations n , if a randomized naming algorithm is executed by n anonymous stations, then an execution is incorrect, in that duplicate names are assigned to different stations, with probability that is at least $n^{-\Omega(1)}$, assuming that the algorithm uses $\mathcal{O}(n \log n)$ random bits with probability $1 - n^{-\Omega(1)}$.*

Proof: Suppose the algorithm uses at most $cn \lg n$ random bits with probability p_n when executed by a channel with n stations, for some constant $c > 0$. Then one of these stations uses at most $c \lg n$ bits with probability p_n , by the pigeonhole principle.

Consider an execution of the same algorithm for $n + 1$ stations. Let us distinguish a station v . One of the remaining n stations, say w , uses at most $c \lg n$ bits with the probability p_n . Station v generates the same string of bits with probability $2^{-c \lg n} = n^{-c}$. The random bits generated by the station w and v are independent. Therefore, duplicate names occur with probability at least $n^{-c} \cdot p_n$. When the bound $p_n = 1 - n^{-\Omega(1)}$ holds, then the probability of duplicate names is at least $n^{-c}(1 - n^{-\Omega(1)}) = n^{-\Omega(1)}$. \square

We conclude this section with a fact about impossibility of developing a Las Vegas naming algorithm when the number of stations n is unknown.

Proposition 5 ([15, 39]) *There is no Las Vegas naming algorithm for a channel with beeping with at least two stations such that it does not refer to the number of stations.*

Proposition 5 was shown by Kutten et al. [39] for the model of processes communicating by asynchronous processes reading from and writing to shared registers. This fact depends not that much on asynchrony as on n being unknown, and it was shown by Chlebus et al. [15] to hold for anonymous synchronous shared-memory systems of read-write registers. The proof given in [15] is general enough to be directly applicable here as well, as both models are synchronous. Proposition 5 justifies developing Monte Carlo algorithm for unknown n , which we do in Section 5.

4 A Las Vegas algorithm

We give a Las Vegas naming algorithm for the case when n is known. The idea is to have stations choose rounds to beep from a segment of integers. As a convenient probabilistic interpretation, these integers are interpreted as bins, and, after selecting a bin, a ball is placed in the bin.

An execution proceeds by considering the consecutive bins one by one. First, a bin is verified to be nonempty by making the owners of the balls in the bin beep. When no beep is heard then the next bin is considered, otherwise the nonempty bin is verified for collisions. Such a verification is performed by $\mathcal{O}(\log n)$ consecutive calls of procedure DETECT-COLLISION given in Section 2. When a collision is not detected then the stations that placed their balls in this bin assign themselves the next available name, otherwise the stations whose balls are in this bin place their balls in a new set of bins. When each station has a name assigned, we verify if the maximum assigned name is n . If this is the case then the algorithm terminates, otherwise we repeat. The algorithm is called BEEP-NAMING-LV, its pseudocode is in Figure 2.

Algorithm BEEP-NAMING-LV is analyzed by modeling its executions by a process of throwing balls into bins, which we call the *ball process*. The process proceeds through stages. There are n balls in the first stage. When a stage begins and there are some i balls eligible for the stage then the number of used bins is $i \lg n$. Each ball is thrown into a randomly selected bin. Next, balls that are single in their bins are removed and the remaining balls that participated in collisions advance to the next stage. The process terminates when no eligible balls remain.

Lemma 2 *The number of times a ball is thrown into a bin during an execution of the ball process that starts with n balls is at most $3n$ with probability at least $1 - e^{-n/4}$.*

Proof: In each stage, we throw some k balls into at least $k \lg n$ bins. The probability that a given ball ends up single in a bin is at least

$$1 - \frac{k}{k \lg n} = 1 - \frac{1}{\lg n},$$

which we denote as p . A ball is thrown repeatedly in consecutive iterations until it lands single in a bin. Our immediate concern is the number of trials to have all balls end up as singletons in their bins.

Suppose that we perform some m independent Bernoulli trials, each with probability p of success, and let X be the number of successes. Here a success means that a new ball ends up single in its bin. We show next that $m = \Theta(n)$ suffices with large probability to have the inequality $X \geq n$.

Algorithm BEEP-NAMING-LV

```
repeat
  counter  $\leftarrow$  0 ; left  $\leftarrow$  1 ; right  $\leftarrow$   $n \lg n$  ; namev  $\leftarrow$  null
  repeat
    slotv  $\leftarrow$  random number in the interval [left, right]
    for  $i \leftarrow$  left to right do
      if  $i = \text{slot}_v$  then
        beep
        if a beep was just heard then
          collision  $\leftarrow$  false
          for  $j \leftarrow$  1 to  $\beta \lg n$  do
            if DETECT-COLLISION then collision  $\leftarrow$  true
          if not collision then
            counter  $\leftarrow$  counter + 1
            namev  $\leftarrow$  counter
    if namev = null then beep /* station  $v$  participated in a collision */
    if a beep was just heard then
      left  $\leftarrow$  counter
      right  $\leftarrow$   $(n - \text{counter}) \lg n$ 
  until no beep was heard in the previous round
until counter =  $n$ 
```

Figure 2: A pseudocode for a station v . The number of stations n is known. The constant $\beta > 1$ is a parameter determined in the analysis. Procedure DETECT-COLLISION has its pseudocode in Figure 1. The variable **name** is to store the assigned identifier.

The expected number of successes is $\mathbb{E}[X] = \mu = pm$. We use the Chernoff bound in the form

$$\Pr(X < (1 - \varepsilon)\mu) < e^{-\varepsilon^2\mu/2} , \quad (1)$$

for any $0 < \varepsilon < 1$; see [43]. It suffices to have the inequality $(1 - \varepsilon)\mu \geq n$. Let us set $\varepsilon = \frac{1}{2}$, so that it suffices to have $\mu \geq 2n$ or $pm \geq 2n$, which is extended into the following form:

$$\left(1 - \frac{1}{\lg n}\right) \cdot m \geq 2n . \quad (2)$$

If we choose $m = 3n$ then inequality (2) holds for sufficiently large n .

The probability that this inequality (2) does not hold is estimated from above by (1). Here we have that $\mu \geq 2n$ so the right-hand side of (1) is $e^{-n/4}$. \square

We proceed to Theorem 1, which summarizes all the good properties of algorithm BEEP-NAMING-LV, in particular that the algorithm is Las Vegas. In the proof, we model executions

of the algorithm as the ball process starting with n balls. The main difference between the ball process and the algorithm is that collisions of balls in bins are detected with certainty, by the specification of the process, while in the algorithm collisions between tentative names might be overlooked with some positive probability.

Theorem 1 *Algorithm BEEP-NAMING-LV, for each $\beta > 0$, terminates almost surely and there is no error when it terminates. For each $a > 0$, there exists $\beta > 1$ and $c > 0$ such that the algorithm assigns unique names, works in time at most $cn \lg n$, and uses at most $cn \lg n$ random bits, all these properties holding with the probability that is at least $1 - n^{-a}$, for sufficiently large n .*

Proof: Consider an iteration of the main repeat-loop. An error can occur in this iteration only when there is a collision that is not detected by procedure DETECT-COLLISION in none of its $\beta \lg n$ calls. Such an error results in duplicate names, so that the number of assigned different names is smaller than n . The maximum name assigned in an iteration is the value of the variable `counter`, which has the same value at each station. The algorithm terminates by having an iteration that produces `counter` = n , but then there are no repetitions among the names, and so there is no error.

Next we show that termination is a sure event. Consider an iteration of the main repeat-loop. There are n balls and each of them is being thrown repeatedly until either it is not involved in a collision or there is a collision but it is not detected. Eventually each ball is left to reside in its bin with probability 1. This means that each iteration ends almost surely.

Next, we introduce the notations for two events in an iteration of the main repeat-loop. We denote by $\neg E$ the complements of an event E .

Let A be the event that there is a collision that passes undetected. The iteration fails to assign proper names if and only if the event A holds. Let B be the event that the total number of throws of balls into bins is at most $3n$. We have that $\Pr(\neg B) \leq e^{-n/4}$, by Lemma 2.

When a ball lands in a bin then it is verified for a collision $\beta \lg n$ times. If there is a collision then it passes undetected with probability at most $n^{-\beta}$. This is because one call of procedure DETECT-COLLISION detects a collision with probability at least $\frac{1}{2}$, by Lemma 1, in which $m = 1$ and $k \geq 2$.

We estimate the probability of the event that an iteration fails to assign proper names, which is the same as of event A . This is accomplished as follows:

$$\begin{aligned}
\Pr(A) &= \Pr(A \cap B) + \Pr(A \cap \neg B) \\
&= \Pr(A | B) \cdot \Pr(B) + \Pr(A | \neg B) \cdot \Pr(\neg B) \\
&\leq \Pr(A | B) + \Pr(\neg B) \\
&\leq 3n \cdot n^{-\beta} + e^{-n/4}, \tag{3}
\end{aligned}$$

where we used the union bound to obtain the last line (3). It follows that at least i iterations are needed with the probability that is at most $(e^{-n/4} + 3n^{1-\beta})^i$, which converges to 0 as i grows unbounded, assuming only that $\beta > 1$ and n is sufficiently large.

Let us consider the event $\neg A \cap B$, which occurs when balls are thrown at most $3n$ times and all collisions are detected, when modeling an iteration of the main repeat loop. The probability that

the event $\neg A \cap B$ holds can be estimated from below as follows:

$$\begin{aligned}
\Pr(\neg A \cap B) &= \Pr(\neg A \mid B) \cdot \Pr(B) \\
&\geq (1 - 3n^{1-\beta}) \cdot (1 - e^{-n/4}) \\
&\geq 1 - 3n^{1-\beta} - e^{-n/4}(1 - 3n^{1-\beta}) .
\end{aligned} \tag{4}$$

This bound (4) is at least $1 - n^{-a}$ for sufficiently large $\beta > 1$, when also n is large enough.

Bound (4) holds for the first iteration of the main repeat loop. So with probability at least $1 - n^{-a}$ the first iteration assigns proper names with at most $3n$ balls thrown in total. Let us assume that this event occurs. Then the whole execution takes time at most $cn \lg n$, for a suitably large $c > 0$. This is because procedure DETECT-COLLISION is executed at most $3\beta n \lg n$ times, and each of its calls takes two rounds. One assignment of a value to variable `slot` requires $\lg(n \lg n) < 2 \lg n$ bits, for sufficiently large n . There are at most $3n$ such assignments, for a total of at most $cn \lg n$ random bits, for a suitably large $c > 0$. \square

Algorithm BEEP-NAMING-LV runs in the optimal expected time $\mathcal{O}(n \log n)$, by Proposition 3, and it uses the optimum expected number of random bits $\mathcal{O}(n \log n)$, by Proposition 2, these propositions given in Section 3.

5 A Monte Carlo algorithm

We give a randomized naming algorithm for the case when n is unknown. In view of Proposition 5, no Las Vegas algorithm exists in this case, so we develop a Monte Carlo one.

The algorithm again can be interpreted as repeatedly throwing balls into bins and verifying for collisions. A bin is determined by a string of some k bits. Each station chooses one such a string randomly. The algorithm proceeds to repeatedly identify the smallest lexicographically string among those not considered yet. This is accomplished by procedure NEXT-STRING which operates as a search implemented by using beeps. Having identified a nonempty bin, all the stations that placed their balls into this bin verify if there is a collision in this bin by calling DETECT-COLLISION a suitably large number of times. In case no collision has been detected, the stations whose balls are in the bin assign themselves the consecutive available name as a temporary one. This continues until all the balls have been considered. If no collision has ever been detected in the current stage, then the algorithm terminates and the temporary names are considered as the final assigned names, otherwise the algorithm proceeds to the next stage.

We specify procedure NEXT-STRING first. It operates as a radix search to identify the smallest string of bits by considering consecutive bit positions. It uses two variables `my-string` and k , where k is the length of the bit strings considered and `my-stringv` is the string of k bits generated by station v . The procedure begins by setting to 1 all bit positions in variable `string`, which has k such bit positions. Then the consecutive bit positions $i = 1, 2, \dots, k$ are considered one by one. For a given bit position i , all the stations v , that still can possibly have the smallest string and whose bit on position i in `my-stringv` is 0, do beep. This determines the first i bits of the smallest string, because if a beep is heard then the i th bit of the smallest string is 0 and otherwise it is 1. This is recorded by setting the i th bit position in the variable `string` to the determined bit. The stations eligible for beeping, if their i th bit is 0, are those whose strings agree on the first $i - 1$ positions with the smallest string. After all k bit positions have been considered, the variable `string` is returned.

Procedure NEXT-STRING

```
string ← a string of  $k$  bit positions, with all of them set to 1
for  $i \leftarrow 1$  to  $k$  do
    if (my-string $v$  matches string on the first  $i - 1$  bit positions)
        and (the  $i$ th bit of my-string $v$  is 0)
            then beep
    if a beep was heard in the previous round then
        set the  $i$ th bit of string to 0
return (string)
```

Figure 3: A pseudocode for a station v . This procedure is used by algorithm BEEP-NAMING-MC. The variables `my-string` and k are the same as those in the pseudocode in Figure 4.

Procedure NEXT-STRING has its pseudocode in Figure 3. Its relevant property is summarized in Lemma 3.

Lemma 3 *Procedure NEXT-STRING returns the smallest lexicographically string among the non-null string values of the private copies of the variable `my-string`.*

Proof: The string that is output is obtained by processing all the input strings `my-string` through consecutive bit positions. We show the invariant that after i bits have been considered, for $0 \leq i \leq k$, then the bits on these positions make the prefix of the first i bits of the smallest string.

The invariant is shown by induction on i . When $i = 1$ then the bits on previously considered positions make an empty string, as no positions have been considered yet, and the empty string is a prefix of the smallest string. Suppose that the invariant holds for all i such that $0 \leq i < k$, and consider the stations whose variable `my-string` has the same bits on these first i positions as variable `string`. This set includes the station v with the smallest `my-string` by the inductive hypothesis. If the bit on the $(i + 1)$ st position of `my-string v` is 0 then v beeps and `string` has its bit on position $i + 1$ set to 0. Otherwise there is no station with 0 on the $(i + 1)$ st position of `my-string v` , because `my-string v` is smallest. Then there is no beep and 1 at position $i + 1$ in `string` is not modified. This completes the proof of the invariant.

The procedure NEXT-STRING terminates after k bit positions have been processed. The proved invariant for $i = k$ means that the smallest `my-string v` and the final value of the variable `string` are identical. \square

The naming algorithm we develop is called BEEP-NAMING-MC. Its pseudocode is in Figure 4. The algorithm proceeds through stages, where a stage is implemented by an iteration of the main repeat-loop in the pseudocode. The number of bins in the i th stage is 2^k where $k = 2^i$. The variable k is doubled in the beginning of each iteration of the main loop. During a stage, first

Algorithm BEEP-NAMING-MC

```
k ← 1
repeat
  k ← 2k
  collision ← false
  counter ← 0
  my-stringv ← a random string of k bits
  repeat
    if my-stringv ≠ null then smallest-string ← NEXT-STRING
    if my-stringv = smallest-string then
      for i ← 1 to β · k do
        if DETECT-COLLISION then collision ← true
      if not collision then
        counter ← counter + 1
        namev ← counter
        my-stringv ← null
    if my-stringv ≠ null then beep
  until no beep was heard in the previous round
until not collision
```

Figure 4: A pseudocode for a station v . Constant $\beta > 0$ is an integer parameter determined in the analysis. Procedure DETECT-COLLISION has its pseudocode in Figure 1 and procedure NEXT-STRING has its pseudocode in Figure 3. The variable **name** is to store the assigned identifier.

the next bin with a ball is identified by calling procedure NEXT-STRING. Next, this bin is verified for collisions by calling procedure DETECT-COLLISION βk times, for a constant $\beta > 0$, which is a parameter to be settled in analysis. During such a verification, the stations whose balls are in this bin participate only.

Theorem 2 summarizes the good properties of algorithm BEEP-NAMING-MC. In particular, that it is a Monte Carlo algorithm with a suitably small probability of error.

Theorem 2 *Algorithm BEEP-NAMING-MC, for each $\beta > 0$, terminates almost surely. For each $a > 0$, there exists $\beta > 0$ and $c > 0$ such that the algorithm assigns unique names, works in time at most $cn \lg n$, and uses at most $cn \lg n$ random bits, all these properties holding with a probability that is at least $1 - n^{-a}$.*

Proof: We interpret an iteration of the outer repeat-loop as a stage in a process of throwing n balls into 2^k bins and verifying βk times for collisions. The string selected by a station is the name of

the bin. When at least one collision is detected then k gets incremented and another iteration is performed. An error occurs when there is a collision but it is not detected.

Next we estimate from above the probability of not detecting a collision. To this end, we consider two cases, depending on which of the inequalities $2^k < n$ or $2^k \geq n$ hold, for a given k .

In the first case, when $2^k < n$, collisions occur with certainty, by the pigeonhole principle. Let m be the number of occupied bins. This results in $m \leq 2^k$ verifications performed, one for each bin, where procedure DETECT-COLLISION is called βk times per verification. By Lemma 1, the probability of not detecting a collision, with just one call of DETECT-COLLISION occurring in each of these verifications, is at most

$$2^{-n+m} \leq 2^{-n+2^k} .$$

When βk calls of DETECT-COLLISION occur in each verification, as is the case by the design of the algorithm, the probability of not detecting a collision is at most

$$2^{(-n+2^k)\beta k} .$$

Intuitively at this point, since $2^k < n$, this probability is maximized for $n = 2^k + 1$ and it is about $2^{-\beta k} \approx n^{-\beta}$, as $k \approx \lg n$.

A precise argument to obtain an estimate is by considering two sub-cases, and is as follows. If it is the sub-case that $2^k < n/2$, then

$$2^{(-n+2^k)\beta k} = 2^{-\Omega(n)} .$$

Otherwise, when it is the sub-case that $n > 2^k \geq n/2$, then $n - 2^k \geq 1$ and $k \geq \lg n - 1$, so that we obtain the following estimate:

$$2^{(-n+2^k)\beta k} \geq 2^{-\beta(\lg n - 1)} \geq 2^\beta n^{-\beta} .$$

We can conclude this sub-case with the following two estimates: $2^{-\Omega(n)}$ and $2^\beta n^{-\beta}$, of which the latter is larger, for sufficiently large n . It is sufficient to take $\beta > a$, as then the inequality $2^\beta n^{-\beta} < n^{-a}$ holds for sufficiently large n .

The second case occurs when $2^k \geq n$. This implies that $k \geq \lg n$. When a collision occurs in a bin, then it is verified by at least $\beta \lg n$ calls of procedure DETECT-COLLISION. This gives the probability of not detecting one such a collision that is at most $n^{-\beta}$. Multiple bins with collisions make the probability of not detecting any of them even smaller. Now it is enough to take $\beta > a$, as then $n^{-\beta} < n^{-a}$ holds.

This completes estimating the probability of error by at most n^{-a} , for sufficiently large β , and all correspondingly large n .

Next, we estimate the probability that the running time is $\mathcal{O}(n \log n)$. Let us consider a stage with sufficiently many bins, say, when $k = d \lg n$ for $d > 2$. Then the number of bins is $2^k = n^d$. The probability that there is no collision at all in this stage is at least

$$\left(1 - \frac{n}{n^d}\right)^n \geq 1 - \frac{n}{n^{d-1}} = 1 - n^{-d+2} . \quad (5)$$

Choosing $d = a + 2$ we obtain that the algorithm terminates by the iteration of the outer repeat-loop when $k = d \lg n$ with the probability that is at least $1 - n^{-a}$. One iteration of the outer repeat

loop, for some k , is proportional to $k \cdot n$. The total time spent up to and including $k = d \lg n$ is proportional to

$$\sum_{i=1}^{\lg((a+2) \lg n)} 2^i \cdot n \leq n \cdot 2(a+2) \lg n = \mathcal{O}(n \log n) \quad (6)$$

with the probability at least $1 - n^{-a}$.

The number of bits generated up to and including the iteration for $k = d \lg n$ is also proportional to (6). This is because the number of bits generated in one iteration of the main repeat-loop is proportional to $k \cdot n$, similarly as the running time.

To show that the algorithm terminates almost surely, it is sufficient to demonstrate that the probability of a collision converges to zero with k increasing. The probability of no collision for $k = d \lg n$ is at most n^{-d+2} , by (6). If k grows to infinity then $d = k / \lg n$ increases to infinity as well, and then n^{-d+2} converges to 0 as a function of d . \square

Algorithm BEEP-NAMING-MC is optimal with respect to the following performance measures: the expected running time $\mathcal{O}(n \log n)$, by Proposition 3, the expected number of random bits used $\mathcal{O}(n \log n)$, by Proposition 2, and the probability of error, as determined by the number of used bits, by Proposition 4.

6 Conclusion

We considered a channel in which a synchronized beeping is the only means of communication. We showed that names can be assigned to the anonymous stations by randomized algorithms. The algorithms are either Las Vegas or Monte Carlo, depending on whether the number of stations n is known or not, respectively. The performance characteristics of the two algorithms, such as the running time, the number of random bits used, and the probability of error, are proved to be asymptotically optimal.

The algorithms we developed rely in an essential manner on synchronization of the channel. It would be interesting to consider an anonymous asynchronous beeping channel and investigate how to assign names to stations in such a communication environment.

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] Y. Afek, N. Alon, O. Barad, E. Hornstein, N. Barkai, and Z. Bar-Joseph. A biological solution to a fundamental distributed computing problem. *Science*, 331(6014):183–185, 2011.
- [3] Y. Afek and Y. Matias. Elections in anonymous networks. *Information and Computation*, 113(2):312–330, 1994.
- [4] D. Angluin. Local and global properties in networks of processors. In *Proceedings of the 12th ACM Symposium on Theory of Computing (STOC)*, pages 82–93, 1980.

- [5] D. Angluin, J. Aspnes, Z. Diamadi, M. J. Fischer, and R. Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, 18(4):235–253, 2006.
- [6] D. Angluin, J. Aspnes, M. J. Fischer, and H. Jiang. Self-stabilizing population protocols. *ACM Transactions on Autonomous and Adaptive Systems*, 3(4), 2008.
- [7] A. F. Anta, M. A. Mosteiro, and J. R. Muñoz. Unbounded contention resolution in multiple-access channels. *Algorithmica*, 67(3):295–314, 2013.
- [8] J. Aspnes and E. Ruppert. An introduction to population protocols. *Bulletin of the EATCS*, 93:98–117, 2007.
- [9] H. Attiya and F. Ellen. *Impossibility Results for Distributed Computing*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, 2014.
- [10] H. Attiya, M. Snir, and M. K. Warmuth. Computing on an anonymous ring. *Journal of the ACM*, 35(4):845–875, 1988.
- [11] P. Brandes, M. Kardas, M. Klonowski, D. Pająk, and R. Wattenhofer. Approximating the size of a radio network in beeping model. In *Proceedings of the 23rd International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, volume 9988 of *Lecture Notes in Computer Science*, pages 358–373. Springer, 2016.
- [12] H. Buhrman, A. Panconesi, R. Silvestri, and P. Vitanyi. On the importance of having an identity or, is consensus really universal? *Distributed Computing*, 18(3):167–176, 2006.
- [13] J. Capetanakis. Tree algorithms for packet broadcast channels. *IEEE Transactions on Information Theory*, 25(5):505–515, 1979.
- [14] 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*, volume I, pages 401–456. Kluwer Academic Publishers, 2001.
- [15] B. S. Chlebus, G. D. Marco, and M. Talo. Anonymous processors with synchronous shared memory. *CoRR*, abs/1507.02272, 2015.
- [16] I. Cidon and M. Sidi. Conflict multiplicity estimation and batch resolution algorithms. *IEEE Transactions on Information Theory*, 34(1):101–110, 1988.
- [17] A. Cornejo and F. Kuhn. Deploying wireless networks with beeps. In *Proceedings of the 24th International Symposium on Distributed Computing (DISC)*, volume 6343 of *Lecture Notes in Computer Science*, pages 148–162. Springer, 2010.
- [18] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley, 2nd edition, 2006.
- [19] A. Czumaj and P. Davies. Communicating with beeps. In *Proceedings of the 19th International Conference on Principles of Distributed Systems (OPODIS)*, volume 46 of *LIPICs*, pages 30:1–30:16. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.
- [20] G. De Marco and D. R. Kowalski. Contention resolution in a non-synchronized multiple access channel. In *Proceedings of the 27th IEEE International Symposium on Parallel and Distributed Processing (IPDPS)*, pages 525–533, 2013.

- [21] G. De Marco and D. R. Kowalski. Fast nonadaptive deterministic algorithm for conflict resolution in a dynamic multiple-access channel. *SIAM Journal on Computing*, 44(3):868–888, 2015.
- [22] J. Degeysys, I. Rose, A. Patel, and R. Nagpal. DESYNC: self-organizing desynchronization and TDMA on wireless sensor networks. In *Proceedings of the 6th International Conference on Information Processing in Sensor Networks (IPSN)*, pages 11–20. ACM, 2007.
- [23] Ö. Eğecioğlu and A. K. Singh. Naming symmetric processes using shared variables. *Distributed Computing*, 8(1):19–38, 1994.
- [24] Y. Emek and R. Wattenhofer. Stone age distributed computing. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 137–146, 2013.
- [25] F. E. Fich, F. Meyer auf der Heide, P. Ragde, and A. Wigderson. One, two, three...infinity: Lower bounds for parallel computation. In *Proceedings of the 17th ACM Symposium on Theory of Computing (STOC)*, pages 48–58, 1985.
- [26] F. E. Fich and E. Ruppert. Hundreds of impossibility results for distributed computing. *Distributed Computing*, 16(2-3):121–163, 2003.
- [27] R. Flury and R. Wattenhofer. Slotted programming for sensor networks. In *Proceedings of the 9th International Conference on Information Processing in Sensor Networks (IPSN)*, pages 24–34. ACM, 2010.
- [28] K. Förster, J. Seidel, and R. Wattenhofer. Deterministic leader election in multi-hop beeping networks. In *Proceedings of the 28th International Symposium on Distributed Computing (DISC)*, volume 8784 of *Lecture Notes in Computer Science*, pages 212–226. Springer, 2014.
- [29] M. Ghaffari and B. Haeupler. Near optimal leader election in multi-hop radio networks. In *Proceedings of the 24th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 748–766. SIAM, 2013.
- [30] M. Ghaffari, B. Haeupler, and M. Khabbazian. Randomized broadcast in radio networks with collision detection. *Distributed Computing*, 28(6):407–422, 2015.
- [31] S. Gilbert and C. C. Newport. The computational power of beeps. In *Proceedings of the 29th International Symposium on Distributed Computing (DISC)*, volume 9363 of *Lecture Notes in Computer Science*, pages 31–46. Springer, 2015.
- [32] O. Goussevskaia, Y. A. Pignolet, and R. Wattenhofer. Efficiency of wireless networks: Approximation algorithms for the physical interference model. *Foundations and Trends in Networking*, 4(3):313–420, 2010.
- [33] A. G. Greenberg, P. Flajolet, and R. E. Ladner. Estimating the multiplicities of conflicts to speed their resolution in multiple access channels. *Journal of the ACM*, 34(2):289–325, 1987.
- [34] J. F. Hayes. An adaptive technique for local distribution. *IEEE Transactions on Communications*, 26(8):1178–1186, 1978.

- [35] K. Hounkanli and A. Pelc. Asynchronous broadcasting with bivalent beeps. In *Proceedings of the 23rd International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, volume 9988 of *Lecture Notes in Computer Science*, pages 291–306. Springer, 2016.
- [36] B. Huang and T. Moscibroda. Conflict resolution and membership problem in beeping channels. In *Proceedings of the 27th International Symposium on Distributed Computing (DISC)*, volume 8205 of *Lecture Notes in Computer Science*, pages 314–328. Springer, 2013.
- [37] T. Jurdziński and D. R. Kowalski. Distributed randomized broadcasting in wireless networks under the SINR model. In M.-Y. Kao, editor, *Encyclopedia of Algorithms*. Springer US, 2014.
- [38] E. Kranakis and D. Krizanc. Distributed computing on anonymous hypercube networks. *Journal of Algorithms*, 23(1):32–50, 1997.
- [39] S. Kutten, R. Ostrovsky, and B. Patt-Shamir. The Las-Vegas processor identity problem (How and when to be unique). *Journal of Algorithms*, 37(2):468–494, 2000.
- [40] R. J. Lipton and A. Park. The processor identity problem. *Information Processing Letters*, 36(2):91–94, 1990.
- [41] O. Michail, I. Chatzigiannakis, and P. G. Spirakis. *New Models for Population Protocols*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, 2011.
- [42] A. Motskin, T. Roughgarden, P. Skraba, and L. J. Guibas. Lightweight coloring and desynchronization for networks. In *Proceedings of the 28th IEEE International Conference on Computer Communications (INFOCOM)*, pages 2383–2391, 2009.
- [43] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
- [44] S. Navlakha and Z. Bar-Joseph. Distributed information processing in biological and computational systems. *Communications of the ACM*, 58(1):94–102, 2014.
- [45] A. Panconesi, M. Papatriantafyllou, P. Tsigas, and P. M. B. Vitányi. Randomized naming using wait-free shared variables. *Distributed Computing*, 11(3):113–124, 1998.
- [46] B. Schieber and M. Snir. Calling names on nameless networks. *Information and Computation*, 113(1):80–101, 1994.
- [47] S. Schmid and R. Wattenhofer. Algorithmic models for sensor networks. In *Proceedings of the 20th International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 2006.
- [48] A. Scott, P. Jeavons, and L. Xu. Feedback from nature: An optimal distributed algorithm for maximal independent set selection. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 147–156, 2013.
- [49] B. S. Tsybakov and V. A. Mikhailov. Free synchronous packet access in a broadcast channel with feedback. *Problems of Information Transmission*, 14(4):32–59, 1978.
- [50] A. C. Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proceedings of the 18th Symposium on Foundations of Computer Science (FOCS)*, pages 222–227. IEEE Computer Society, 1977.

- [51] J. Yu, L. Jia, D. Yu, G. Li, and X. Cheng. Minimum connected dominating set construction in wireless networks under the beeping model. In *Proceedings of the 2015 IEEE Conference on Computer Communications (INFOCOM)*, pages 972–980, 2015.