

# Efficient Assignment of Identities in Anonymous Populations

Leszek Gąsieniec ✉ 🏠 

Department of Computer Science, University of Liverpool, UK

Jesper Jansson ✉ 🏠 

Graduate School of Informatics, Kyoto University, Japan

Christos Levcopoulos ✉ 🏠 

Department of Computer Science, Lund University, Sweden

Andrzej Lingas ✉ 🏠 

Department of Computer Science, Lund University, Sweden

---

## Abstract

We consider the fundamental problem of assigning distinct labels to agents in the probabilistic model of population protocols. Our protocols operate under the assumption that the size  $n$  of the population is embedded in the transition function. Their efficiency is expressed in terms of the number of states utilized by agents, the size of the range from which the labels are drawn, and the expected number of interactions required by our solutions. Our primary goal is to provide efficient protocols for this fundamental problem complemented with tight lower bounds in all the three aspects. W.h.p. (with high probability), our labeling protocols are silent, i.e., eventually each agent reaches its final state and remains in it forever, and they are safe, i.e., never update the label assigned to any single agent. We first present a silent w.h.p. and safe labeling protocol that draws labels from the range  $[1, 2n]$ . Both the number of interactions required and the number of states used by the protocol are asymptotically optimal, i.e.,  $O(n \log n)$  w.h.p. and  $O(n)$ , respectively. Next, we present a generalization of the protocol, where the range of assigned labels is  $[1, (1 + \varepsilon)n]$ . The generalized protocol requires  $O(n \log n / \varepsilon)$  interactions in order to complete the assignment of distinct labels from  $[1, (1 + \varepsilon)n]$  to the  $n$  agents, w.h.p. It is also silent w.h.p. and safe, and uses  $(2 + \varepsilon)n + O(n^c)$  states, for any positive  $c < 1$ . On the other hand, we consider the so-called pool labeling protocols that include our fast protocols. We show that the expected number of interactions required by any pool protocol is  $\geq \frac{n^2}{r+1}$ , when the labels range is  $1, \dots, n + r < 2n$ . Furthermore, we provide a protocol which uses only  $n + 5\sqrt{n} + O(n^c)$  states, for any  $c < 1$ , and draws labels from the range  $1, \dots, n$ . The expected number of interactions required by the protocol is  $O(n^3)$ . Once a unique leader is elected it produces a valid labeling and it is silent and safe. On the other hand, we show that (even if a unique leader is given in advance) any silent protocol that produces a valid labeling and is safe with probability  $> 1 - \frac{1}{n}$ , uses  $\geq n + \sqrt{\frac{n-1}{2}} - 1$  states. Hence, our protocol is almost state-optimal. We also present a generalization of the protocol to include a trade-off between the number of states and the expected number of interactions. Finally, we show that for any silent and safe labeling protocol utilizing  $n + t < 2n$  states, the expected number of interactions required to achieve a valid labeling is  $\geq \frac{n}{t+1}$ .

**2012 ACM Subject Classification** Theory of computation; Theory of computation → Design and analysis of algorithms; Theory of computation → Complexity theory and logic; Theory of computation → Distributed computing models

**Keywords and phrases** population protocol, state efficiency, time efficiency, one-way epidemics, leader election, agent identities

**Digital Object Identifier** 10.4230/LIPIcs.OPODIS.2021.12

**Funding** Research supported in part by VR grant 621-2017-03750 (Swedish Research Council).

**Acknowledgements** The authors are thankful to the anonymous referees for their valuable comments.



© Leszek Gąsieniec, Jesper Jansson, Christos Levcopoulos, and Andrzej Lingas; licensed under Creative Commons License CC-BY 4.0

25th International Conference on Principles of Distributed Systems (OPODIS 2021).

Editors: Quentin Bramas, Vincent Gramoli, and Alessia Milani; Article No. 12; pp. 12:1–12:21

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

## 1 Introduction

The problem of assigning and further maintaining unique identifiers for entities in distributed systems is one of the core problems related to network integrity. In addition, a solution to this problem is often an important preprocessing step for more complex distributed algorithms. The tighter the range that the identifiers are drawn from, the harder the assignment problem becomes.

In this paper we adopt the probabilistic population protocol model in which we study the problem of assigning distinct identifiers, which we refer to as *labels*, to all agents<sup>1</sup> The adopted model was originally intended to model large systems of agents with limited resources (state space) [4]. In this model the agents are prompted to interact with one another towards a solution of a shared task. The execution of a protocol in this model is a sequence of pairwise interactions between randomly chosen agents. During an interaction, each of the two agents: the *initiator* and the *responder* (the asymmetry assumed in [4]) updates its state in response to the observed state of the other agent according to the predefined (global) transition function. For more details about the population protocol model, see Appendix A.

Designing our population protocols for the problem of assigning unique labels to the agents (labeling problem), we make an assumption that the number  $n$  of agents is known in advance. Our protocols would also work if only an upper bound on the number of agents is known to agents. In fact, in such case the problem becomes easier as the range from which the labels are drawn is larger. In particular, if we do not have the limit on  $n$  we also do not have limit on the number of states to be used. More natural assumption is that such a limit is imposed. And indeed, there are plenty of population protocols which rely on the knowledge of  $n$  [12, 13].

Our labeling protocols include a preprocessing for electing a *leader*, i.e., an agent singled out from the population, which improves coordination of more complex tasks and processes. A good example is synchronization via phase clocks propelled by leaders. More examples of leader-based computation can be found in [5].

In the unique labeling problem adopted here, the number of utilized states needs to reflect the number of agents  $n$ . Also,  $\Omega(n \log n)$  is a natural lower bound on the expected number of interactions required to solve not only the labeling problem but any non-trivial problem by a population protocol. The main reason is that  $\Omega(n \log n)$  interactions are needed to achieve a positive constant probability that each agent is involved in at least one interaction [10].

Perhaps the simplest protocol for unique labeling in population networks is as follows [13] (cf. [11]). Initially, all agents hold label 1 which is equivalent with all agents being in state 1. In due course, whenever two agents with the same label  $i$  interact, the responder updates own label to  $i + 1$ . The advantage of this simple protocol is that it does not need any knowledge of the population size  $n$  and it utilizes only  $n$  states and assigns labels from the smallest possible range  $[1, n]^2$ . The severe disadvantage is that it needs at least a cubic in  $n$  number of interactions (getting rid of the last multiple label  $i$ , for all  $i = 1, \dots, n - 1$ , requires a quadratic number of interactions in expectation) to achieve the configuration in which the agents have distinct labels.

In the following two examples of protocols for unique labeling, we assume that the population size  $n$  is embedded in the transition function, such protocols are commonly used and known as non-uniform protocols [3], and one of the agents is distinguished as the leader, see leader based protocols [5].

<sup>1</sup> When the size of the label range is equal to the number of agents, the problem is also called *ranking* in the literature [12].

<sup>2</sup> We shall denote a range  $[p, \dots, q]$  by  $[p, q]$  from here on.

In the first of the two examples, we instruct the leader to pass labels  $n, n-1, \dots, 2$  to the encountered subsequently unlabeled yet agents and finally assign 1 to itself. The protocol uses only  $2n-1$  states ( $n$  states utilized by the leader and  $n-1$  states by other agents) and it assigns unique labels in the smallest possible range  $[1, n]$  to the  $n$  agents. Unfortunately, this simple protocol requires  $\Omega(n^2 \log n)$  interactions because as more agents get their labels, interactions between the leader and agents without labels become less likely. The probability of such an encounter drops from  $\frac{1}{n}$  at the beginning to  $\frac{1}{n(n-1)}$  at the end of the process.

By using randomization, we can obtain a much faster simple protocol as follows. We let the leader to broadcast the number  $n$  to all agents. It requires  $O(n \log n)$  interactions w.h.p.<sup>3</sup> [18]. When an agent gets the number  $n$ , it uniformly at random picks a number in  $[1, n^3]$  as its label. The probability that a given pair of agents gets the same label is only  $\frac{1}{n^3}$ . Hence, this protocol assigns unique labels to the agents with probability at least  $1 - \frac{1}{n}$ . It requires only  $O(n \log n)$  interactions w.h.p. The drawback is that it uses  $O(n^3)$  states and the large range  $[1, n^3]$ . This method also needs a large number of random bits independent for each agent.

Besides the efficiency and population size aspects, there are other deep differences between the three examples of labeling protocols. An agent in the first protocol never knows whether or not it shares its label with other agents. This deficiency cannot happen in the case of the second protocol but it takes place in the third protocol although with a small probability.

The labeling protocols presented in this paper are *silent* and *safe*. We say that a (non-necessarily labeling) protocol is silent if eventually each agent reaches its final state and remains in it forever. We say that a labeling protocol is safe if it never updates the label assigned to any single agent. While the concept of a silent population protocol is well established in the literature [12, 15], the concept of a safe labeling protocol is new. The latter property is useful in the situation when the protocol producing a valid labeling has to be terminated before completion due to some unexpected emergency or running out of time.

Observe that among the three examples of labeling protocols, only the second one is both silent and safe. The first example protocol is silent [12] but not safe. Finally, the third (probabilistic) one is silent and almost safe as it violates the definition only with small probability.

**Our contributions.** The primary objective of this paper is to provide efficient labeling protocols complemented with tight lower bounds in the aspects of the number of states utilized by agents, the size of the range from which the labels are drawn, and the expected number of interactions required by our solutions.

In particular, we provide positive answers to two following natural questions under the assumption that the number  $n$  of agents is known at the beginning.

1. Can one design a protocol for the labeling problem requiring an asymptotically optimal number of  $O(n \log n)$  interactions w.h.p., utilizing an asymptotically optimal number of  $O(n)$  states and an asymptotically minimal label range of size  $O(n)$  ?
2. Can one design a silent and safe protocol for the labeling problem utilizing substantially smaller number of states than  $2n$  and possibly the minimal label range  $[1, n]$  ?

We first present a population protocol that w.h.p. requires an asymptotically optimal number of  $O(n \log n)$  interactions to assign distinct labels from the range  $[1, 2n]$ . The protocol uses an asymptotically optimal number of  $O(n)$  states. We also present a more

---

<sup>3</sup> That is with the probability at least  $1 - \frac{1}{n^\alpha}$ , where  $\alpha \geq 1$  and  $n$  is the number of agents.

involved generalization of the protocol, where the range of assigned labels is  $[1, (1 + \varepsilon)n]$ . The generalized protocol requires  $O(n \log n / \varepsilon)$  interactions in order to complete the assignment of distinct labels from  $[1, (1 + \varepsilon)n]$  to the  $n$  agents, w.h.p. It uses  $(2 + \varepsilon)n + O(n^c)$  states, for any positive  $c < 1$ . Both protocols are silent w.h.p. and safe. Furthermore, we consider a natural class of population protocols for the unique labeling problem, the so-called *pool protocols*, including our fast labeling protocols. We show that for any protocol in this class that picks the labels from the range  $[1, n + r]$ , the expected number of interactions is  $\Omega(\frac{n^2}{r+1})$ .

Next, we provide a labeling protocol which uses only  $n + 5\sqrt{n} + O(n^c)$  states, for any positive  $c < 1$ , and the label range  $[1, n]$ . The expected number of interactions required by the protocol is  $O(n^3)$ . Once a unique leader is elected it produces a valid labeling and it is silent and safe. On the other hand, we show that (even if a unique leader is given in advance) any silent protocol that produces a valid labeling and is safe with probability larger than  $1 - \frac{1}{n}$ , uses at least  $n + \sqrt{\frac{n-1}{2}} - 1$  states. It follows that our protocol is almost state-optimal. In addition, we present a variant of this protocol which uses  $n(1 + \varepsilon) + O(n^c)$  states, for any positive  $c < 1$ . The expected number of interactions required by this variation is  $O(n^2/\varepsilon^2)$ , where  $\varepsilon = \Omega(n^{-1/2})$ . On the other hand, we show that for any silent and safe labeling protocol utilizing  $n + t < 2n$  states the expected number of interactions required to achieve a valid labeling is at least  $\frac{n^2}{t+1}$ .

All our labeling protocols include a preprocessing for electing a unique leader and assume the knowledge of the population size  $n$ . However, our almost state-optimal protocol (Single-Cycle protocol) can be made independent of  $n$  (see Section 4).

Our results are summarized in Tables 1 and 2.

**Main ideas of our protocols.** Our first fast labeling protocol roughly operates as follows. The leader initially has label 1 and a range of labels  $[2, n]$ . During the execution of the first phase, encountered unlabeled agents also get a label and an interval of labels that they can distribute among other agents. Upon a communication between a labeled agent that has a non-empty interval and an unlabeled agent, the latter agent gets a label from the interval and if the remaining part of the interval has length  $\geq 2$  then it is shared between the two agents. After  $O(n \log n)$  interactions, a sufficiently large fraction of agents is labeled and has no additional labels to distribute w.h.p. The leader counts its own interactions up to  $O(\log n)$  in order to trigger the second phase by broadcasting. In the latter phase, an agent with a label  $x$  and without a non-empty interval can distribute one additional label  $x + n$ . In the first phase of this protocol the labels from the range  $[1, n]$  are distributed rapidly among the agents. In the second phase the unlabeled agent still have a high chance of communicating with an agent that can distribute a label. Roughly, our second, generalized fast labeling protocol is obtained from the first one by constraining the set of agents that may distribute the labels  $x + n$  in the second phase to those having labels in the range  $[1, n\varepsilon]$ .

The main idea of the almost state-optimal labeling protocol (Single-Cycle protocol) is to use the leader and an auxiliary leader *nominated* by the leader to disperse the  $n$  labels jointly among the remaining free agents. The leader disperses the first and the auxiliary leader the second part of each individual label. When a free agent gets both partial labels, it combines them into its individual label and then informs the leaders about this. The two leaders operate in two embedded loops. For each of roughly  $\sqrt{n}$  partial labels of the leader, the auxiliary leader makes a full round of dispersing its roughly  $\sqrt{n}$  partial labels. In the generalized version of the protocol (k-Cycle protocol), the process is partially parallelized by letting the leader to form  $k$  pairs of dispensers, where each pair labels agents in a distinct range of size  $n/k$ .

■ **Table 1** Upper bounds on the number of states, the number of interactions and the range required by the labeling protocols presented in this paper. In Theorem 9,  $\varepsilon$  is  $\Omega(n^{-1})$  while in Theorem 12  $\Omega(n^{-0.5})$ .

Theorem	# states	# interactions	Range
Theorem 4	$O(n)$	$O(n \log n)$ w.h.p.	$[1, 2n]$
Theorem 9	$(2 + \varepsilon)n + O(n^c)$ , any $c < 1$	$O(n \log n / \varepsilon)$ w.h.p.	$[1, (1 + \varepsilon)n]$
Theorem 11	$n + 5 \cdot \sqrt{n} + O(n^c)$ , any $c < 1$	expected $O(n^3)$	$[1, n]$
Theorem 12	$(1 + \varepsilon)n + O(n^c)$ , any $c < 1$	expected $O(n^2 / \varepsilon^2)$	$[1, n]$

■ **Table 2** Lower bounds on the number of states or/and the number of interactions required by labeling protocols. (1) Any labeling protocol that is capable to produce a valid labeling. (2) The silent protocol in Theorem 14 (first part) is assumed to produce a valid labeling and be safe with probability greater than  $1 - \frac{1}{n}$ . (3) The silent protocol in Theorem 14 (2nd part) is assumed to produce a valid labeling and be safe with probability 1.

Protocol type	# states	# interactions	Theorem
any <sup>1</sup>	$n$	$\Omega(n \log n)$ w.h.p.	Theorem 13
silent, safe <sup>2</sup>	$n + \sqrt{\frac{n-1}{2}} - 1$	-	Theorem 14 (1st part)
silent, safe <sup>3</sup> , $n + t < 2n$ states	-	expected $\frac{n^2}{t+1}$	Theorem 14 (2nd part)
pool, range $[1, n + r]$	-	expected $\frac{n^2}{r+1}$	Theorem 17

**Related work.** There are several papers concerning labeling of processing units (also known as renaming or naming) in different communication models [14]. E.g., Berenbrink et al. [8] present efficient algorithms for the so-called loose and tight renaming in shared memory systems improving on or providing alternative algorithms to the earlier algorithms by Alistarh et al. [2, 1]. The loose renaming where the label space is larger than the number of units is shown to admit substantially faster algorithms than the tight renaming [1, 8].

The problem of assigning unique labels to agents has been studied in the model of population protocols by Beauquier et al. [7, 11]. In [11], the emphasis is on estimating the minimum number of states which are required by apparently non-safe protocols. In [7], the authors provide among other things a generalization of a leader election protocol to include a distribution of  $m$  labels among  $n$  agents, where  $m \leq n$ . In the special case of  $m = n$ , all agents will receive unique labels. No analysis on the number of interactions required by the protocol is provided in [7]. Their focus is on the feasibility of the solution, i.e., that the process eventually stabilizes in the final configuration. Their protocol seems inefficient in the state space aspect as it needs many states/bits to keep track of all the labels.

Doty et al. considered the labeling problem in [17] and presented a subroutine named “UniqueID” for it based on the technique of traversing a labeled binary tree and associating agents with nodes in the tree. The subroutine requires  $O(n \log n \log \log n)$  interactions.

The labeling problem has also been studied in the context of self-stabilizing protocols where the agents start in arbitrary (not predefined) states, see [12, 13]. In [13], Cai et al. propose a solution which coincides with our first example of labeling protocols presented in the introduction. In a very recent work [12], Burman et al. study both slow and fast labeling protocols, the latter utilizing an exponential number of states. The protocols in both papers require the exact knowledge of  $n$ . The work [12] focuses on self-stabilizing protocols which cannot be safe by definition. It is more proper to compare our protocols with the initialized version of the protocols in [12]. E.g., the leader-driven initialized (silent) ranking protocol

in [12] (see Lemma 4.1) requires  $O(n^2)$  interactions, uses  $O(n)$  states and it is safe. An analogous variant of the fast ranking protocol from [12] requiring  $O(n \log n)$  interactions and an exponential number of states is also safe but not silent. The known labeling protocols are summarized in Table 3 in Appendix B.

The most closely related problem more studied in the literature is that of counting the population size, i.e., the number of agents. It has been recently studied by Aspnes et al. in [6] and Berenbrink et al. in [10]. We assume that the population size is initially known. Alternatively, it can be computed by using the protocol counting the exact population size given in [10]. The aforementioned protocol computes the population size in  $O(n \log n)$  interactions w.h.p., using  $\tilde{O}(n)$  states. Another possibility is to use the protocol computing the approximate population size, presented in [10]. The latter protocol requires  $O(n \log^2 n)$  interactions to compute the approximate size w.h.p., using only a poly-logarithmic number of states. For references to earlier papers on protocols for counting or estimating the population size, in particular the papers that introduced the counting problem and that include the original algorithms on which the improved algorithms of Berenbrink et al. are based, see [10].

All our protocols include a preprocessing for electing a unique leader and its synchronization with the proper labeling protocol (e.g., see the proof of Theorem 4). There is a vast literature on population protocols for leader election [9, 16, 18, 19]. For our purposes, the most relevant is the protocol that elects a unique leader from a population of  $n$  agents using  $O(n \log n)$  interactions and  $O(n^c)$  many states, for any positive constant  $c < 1$ , w.h.p., described in [10, 16] (see also Fact 7). The newest results elaborate on state-optimal leader election protocols utilizing  $O(\log \log n)$  states. These include the fastest possible protocol [9] based on  $O(n \log n)$  interactions in expectation, and a slightly slower protocol [19] requiring  $O(n \log^2 n)$  interactions with high probability.

Our population protocols for unique labeling use also the known population protocol for (one-way) epidemics, or broadcasting. It completes spreading a message in  $\Theta(n \log n)$  interactions w.h.p. and it uses only two states [18] (see also Fact 4).

**Organization of the paper.** In the next section, we provide basic facts on probabilistic inequalities and population protocols for broadcasting, counting and leader election. In Section 3, we present our fast silent w.h.p. and safe protocol for unique labeling in the range  $[1, 2n]$  and its generalization to include the range  $[1, n(1 + \varepsilon)]$ . Section 4 is devoted to the almost state-optimal, roughly silent and safe protocol with the label range  $[1, n]$  and its variation. Section 5 presents lower bounds on the number of states or the number of interactions for silent, safe and the so-called pool protocols for unique labeling. We conclude with Final remarks.

## 2 Preliminaries

### Probabilistic bounds.

► **Fact 1** (The union bound). *For a sequence  $A_1, A_2, \dots, A_r$  of events,  $\text{Prob}(A_1 \cup A_2 \cup \dots \cup A_r) \leq \sum_{i=1}^r \text{Prob}(A_i)$ .*

► **Fact 2** (multiplicative Chernoff lower bound). *Suppose  $X_1, \dots, X_n$  are independent random variables taking values in  $\{0, 1\}$ . Let  $X$  denote their sum and let  $\mu = E[X]$  denote the sum's expected value. Then, for any  $\delta \in [0, 1]$ ,*

*$\text{Prob}(X \leq (1 - \delta)\mu) \leq e^{-\frac{\delta^2 \mu}{2}}$  holds. Similarly, for any  $\delta \geq 0$ ,  $\text{Prob}(X \leq (1 + \delta)\mu) \leq e^{-\frac{\delta^2 \mu}{2 + \delta}}$  holds.*

► **Fact 3** ([18]). *For all  $C > 0$  and  $0 < \delta < 1$ , during  $Cn \log n$  interactions, with probability at least  $1 - n^{-O(\delta^2 C)}$ , each agent participates in at least  $2C(1 - \delta) \log n$  and at most  $2C(1 + \delta) \log n$  interactions.*

**Broadcasting, counting and leader election.** We shall refer to the following broadcast process which can be completed during  $\Theta(n \log n)$  interactions w.h.p. Each agent is either in a state of M-type (got the message) or in a state of  $\neg$ M-type. Whenever an agent in a state of M-type interacts with an agent in a state of  $\neg$ M-type, the latter changes its state to a state of M-type (gets the message). The process starts when the first agent gets the message and completes when all agents have the message.

► **Fact 4.** *There is a constant  $c_0$ , such that for  $c \geq c_0$ , the broadcast process completes in  $cn \log n$  interactions with probability at least  $1 - n^{-\Theta(c)}$ .*

Berenbrink et al. [10] obtained among other things the following results on counting the population size, i.e., the number of agents.

► **Fact 5.** *There is a protocol for a population of an unknown number  $n$  of agents such that w.h.p., after  $O(n \log^2 n)$  interactions the protocol stabilizes and each agent holds the same estimation of the population size which is either  $\lceil \log n \rceil$  or  $\lfloor \log n \rfloor$ . The protocol uses  $O(\log^2 n \log \log n)$  states.*

► **Fact 6.** *There is a protocol for a population of an unknown number  $n$  of agents such that w.h.p., after  $O(n \log n)$  interactions the protocol stabilizes and each agent holds the exact population size. The protocol uses  $\tilde{O}(n)$  states.*

There is a vast literature on population protocols for leader election [18]. For our purposes, the following fact will be sufficient. Its idea is to start leader election with a subprotocol of [19] that elects a junta of substantially sublinear in  $n$  number of leaders. The junta is formed using  $O(n \log n)$  interactions. Then, when state space of size  $n^c$  is available,  $c < 1$ , only a constant number of rounds of leader elimination is needed, each requiring  $O(n \log n)$  interactions. For more details, see [10, 16].

► **Fact 7.** *There is a protocol that elects a unique leader from a population of  $n$  agents using  $O(n \log n)$  interactions and  $O(n^c)$  many states, for any positive constant  $c < 1$ , w.h.p. [10, 16].*

### 3 Labeling with asymptotically optimal number of interactions, nearly optimal number of states and range

In this section, we provide a silent w.h.p. and a safe labeling protocol that assigns unique labels from the range  $[1, 2n]$  to  $n$  agents in  $O(n \log n)$  interactions w.h.p. Then, we generalize the protocol to include the range  $[1, (1 + \varepsilon)n]$ , where  $\varepsilon$  does not have to be a constant; it can even be as small as  $O(n^{-1})$ . We show that the generalized protocol assigns unique labels from  $[1, (1 + \varepsilon)n]$  in  $O(n \log n / \varepsilon)$  interactions w.h.p. In the first protocol, the agents use  $O(n)$  states, in the second protocol only  $(2 + \varepsilon)n + O(n^c)$  states, for any positive  $c < 1$ .

**Range  $[1, 2n]$ .** The protocol runs in two main phases preceded by a leader election preprocessing. The idea of the first phase resembles that of load balancing [10], the difference is that tokens (in our case labels and interval sub-ranges) are distinct.

At the beginning of the first phase, the leader assigns the label 1 and also temporarily the interval  $[2, n]$  to itself. Next, whenever two agents interact, one with label and a temporarily assigned interval  $[q, r]$  where  $r > q$  and the other without label, the former agent shrinks its interval to  $[q, \lfloor \frac{q+r}{2} \rfloor]$  and it gives away the label  $\lfloor \frac{q+r}{2} \rfloor + 1$  and if  $\lfloor \frac{q+r}{2} \rfloor + 2 \leq r$  also the sub-interval  $[\lfloor \frac{q+r}{2} \rfloor + 2, r]$  to the latter agent. Furthermore, whenever an agent with label and a temporarily assigned singleton interval  $[q, q]$  interacts with an agent without label, the former agent cancels its interval and gives the label  $q$  to the latter agent. In the remaining cases, interactions have no effect. Note that during the first phase a sub-tree of the binary tree of the partition of the start interval  $[1, n]$  with  $n$  leaves determined by the protocol rules is formed; see Fig. 1 in Appendix C. Also observe that when an agent at an intermediate node of the tree interacts with an agent without label then the former agent migrates to the left child of the node while the latter agent lands at the right child of the node.

In the second phase, when an agent with a label  $i \in [1, n]$  at a leaf of the tree interacts with an agent without label for the first time then the latter agent gets the label  $i + n$ . Interactions between agents (if any) at intermediate nodes of the tree and agents without labels are defined as in the first phase.

The following lemmata are central in showing that  $O(n \log n)$  interactions are sufficient w.h.p. to implement our protocol.

► **Lemma 1.** *There is a constant  $c$  such that after  $cn \log n$  interactions in the first phase the number of agents without labels drops below  $n/4$  w.h.p.*

**Proof.** The proof is by contradiction. Suppose that a set  $F$  of at least  $n/4$  agents without labels survives at least  $cn \log n$  interactions, where the constant  $c$  will be specified later.

Consider first the leader agent starting with the interval  $[2, n]$  during the aforementioned interactions. When the agent interacts with an agent without label its interval is roughly halved. We shall call such an interaction a success. The probability of success is at least  $\frac{1}{4n}$ . The expected number of successes is at least  $\frac{c}{4} \log n$ . By using Chernoff multiplicative bound given in Fact 2, we can set  $c$  to enough large constant so the probability of at least  $\log_2 n + 1$  successes will be at least  $1 - \frac{1}{n^2}$ . This means that the leader will end up without any interval with so high probability during the  $cn \log n$  interactions. The leader chooses the leftmost path in the binary partition tree of the start interval  $[1, n]$ . Consider an arbitrary path  $P$  from the root to a leaf in the tree. Note that several agents during distinct interactions can appear on the path. Define as a success an interaction in which an agent currently on  $P$  interacts with an agent without label. The expected number of successes is again at least  $\frac{c}{4} \log n$  and again we can conclude that there are at least  $\log_2 n + 1$  successes with probability at least  $1 - \frac{1}{n^2}$ . Simply, the probabilities of interacting with an agent without label are the same for all agents with labels, i.e., on some paths in the tree. Another way to argue is that the leader could make other decisions as to which roughly half of interval to preserve and the path choice. By the union bound (Fact 1), we conclude that all the  $n$  paths from the root to the leaves in the tree could be developed during the  $cn \log n$  interactions, so all agents would get a label, with probability at least  $1 - \frac{1}{n}$ . We obtain a contradiction with the so long existence of the set  $F$ . ◀

► **Lemma 2.** *If the second phase starts after  $cn \log n$  interactions, where  $c$  is the constant from Lemma 1, then only  $O(n \log n)$  interactions are needed to assign labels in  $[1, 2n]$  to the remaining agents without labels, w.h.p.*

**Proof.** The number of agents without labels at the beginning of the second phase is at most  $n/4$  w.h.p. Hence, at the beginning of this phase the number of agents with labels is at least  $\frac{3}{4}n$  w.h.p. An agent with label  $i \leq n$  at a leaf of the tree can give the label  $i + n$  to an agent

without label only once. Since this can happen at most  $\frac{n}{4}$  times, the number of agents with labels in  $[1, n]$  that can give a label is always at least  $\frac{n}{2}$  w.h.p. We conclude that for an agent without label the probability of an interaction with an agent that can give a label is at least almost  $\frac{1}{2n}$ . Hence, after each  $O(n)$  interactions the expected number of agents without label halves. It follows that the expected number of such interactions rounds is  $O(\log n)$ . Consequently, the number of the rounds is also  $O(\log n)$  w.h.p. by Chernoff bound (Fact 2).

An alternative way to obtain the  $O(n \log n)$  bound on the number of interactions w.h.p. is to use Fact 3 with  $C = O(\frac{1}{1/2})$  and  $\delta = \frac{1}{2}$ . Then, each agent will interact with at least  $C \log n$  agents w.h.p. during  $Cn \log n$  interactions. Consequently, the probability that a given agent does not interact with any agent that can give a label during the aforementioned interactions is  $(1 - \frac{1}{2})^{O(2 \log n)}$ . Hence, by picking enough large  $C$ , we conclude that each agent (in particular without label) will interact with at least one agent that can give a label during the  $Cn \log n$  interactions w.h.p. ◀

► **Lemma 3.** *During both phases, no pair of agents gets the same label.*

**Proof.** The uniqueness of the label assignments in the first phase follows from the disjointness of the labels and intervals assigned to agents before and after each interaction. This argument also works for the labels not exceeding  $n$  assigned later in the second phase. Finally, the uniqueness of the labels of the form  $i + n$  follows from the uniqueness of the labels of the agents passing these labels. ◀

► **Theorem 4.** *There is a safe protocol for population of  $n$  agents that w.h.p. assigns unique labels in the range  $[1, 2n]$  to the agents equipped with  $O(n)$  states in  $O(n \log n)$  interactions. The protocol is also silent w.h.p.*

**Proof.** Under the assumption that the leader election preprocessing provides a unique leader, the correctness of label assignment in both phases w.h.p. and the fulfilling of the definition of a silent and safe protocol follows from Lemmata 1, 2, and 3 and the specification of the protocol, respectively.

For the purpose of the leader election preprocessing, we use the simple leader election protocol using  $O(n \log n)$  interactions and  $O(n^c)$  states, for any positive constant  $c$ , described in [10, 16] (Fact 7). The phase clock (based on junta of leaders) from [19] is also formed in  $O(n \log n)$  interactions, using  $O(\log \log n)$  states and we use this clock to count the required (by the simple leader election protocol) time  $\Omega(n \log n)$ . When this time is reached on the clock we switch from leader election to our proper labeling protocol. The two aforementioned processes can be run simultaneously, resulting in additional state usage  $O(n^c \log \log n)$  (still fine for our needs). Thus, the leader election preprocessing and its synchronization with the proper labeling protocol in two phases add  $O(n \log n)$  interactions and  $o(n)$  states w.h.p. It provides a unique leader w.h.p. It follows that w.h.p. the whole protocol provides a correct labeling, it is silent and safe. In fact, we can make it safe (with probability 1) by prohibiting agents to change or get rid of an assigned label. Note that this constraint does not affect the operation of the protocol when a unique leader is provided by the preprocessing.

Both phases require  $O(n \log n)$  interactions w.h.p. by Lemmata 1, 2.

To put the two phases described in Lemmata 1, 2 together, we let the leader agent to count its interactions. When the number of interactions of the leader in the first phase exceeds an appropriate multiplicity of  $\log n$ , the total number of interactions in the first phase achieves the required lower bound from Lemma 1 w.h.p. by Fact 3. Therefore, then the leader starts broadcasting the message on the transition to the second phase to the other agents. By Fact 4, the broadcasting increases the number of interactions only by  $O(n \log n)$  w.h.p. (The leader can also stop the second phase in a similar fashion.)

## 12:10 Efficient Assignment of Identities in Anonymous Populations

To save on the number of states, instead of having states corresponding to all possible sub-intervals of  $[1, n]$ , we consider states corresponding to the nodes of the interval partition tree (see Fig. 1 in Appendix C) whose sub-tree is formed in the first phase. More precisely, we associate two states with each intermediate node of the binary tree on  $n$  leaves and  $n - 1$  intermediate nodes. They indicate whether or not the agent at the intermediate node has already received the message about the transition to the second phase. Next, we associate four states to each leaf of the tree. They indicate similarly whether or not the agent at the leaf has already received the phase transition message and whether or not the agent has already passed a label to an agent without label in the second phase, respectively. With each label in the range  $[n + 1, 2n]$ , we associate only a single state. Additionally, there are  $O(\log n)$  states used by the leader to count interactions in order to start the second phase. Recall also that the leader election preprocessing requires  $o(n)$  additional states. Thus the total number of states does not exceed  $2n + 4n + n + o(n)$ . ◀

By combining the protocol of Theorem 4 with that of Berenbrink et al. for exact counting the population size (Fact 6), we obtain the following corollary on unique labeling when the population size is unknown to agents initially.

► **Corollary 5.** *There is a protocol for a population of  $n$  agents that assigns unique labels in the range  $[1, 2n]$  to the agents initially not knowing the number  $n$ , equipped with  $\tilde{O}(n)$  states, in  $O(n \log n)$  interactions w.h.p.*

**Proof.** We run first the protocol for exact counting (Fact 6) and then our protocol for unique labeling (Theorem 4) using the leader elected by the counting protocol. We can synchronize the three protocols in a similar fashion as we synchronized the two phases of our protocol additionally using  $O(n \log n)$  interactions and  $O(\log n)$  states. ◀

By using the method of approximate counting from [10] (Fact 5) instead of that for exact counting (Fact 6), we can decrease the number of states to  $O(n)$  at the cost of increasing the label range to  $[1, 8n]$  and the number of interactions required to  $O(n \log^2 n)$ .

**Range  $[1, (1 + \varepsilon)n]$ .** The new protocol is obtained by the following modifications in the previous one. The leader which counts the number of own interactions starts broadcasting the phase transition message when the number of agents without labels drops below  $n\varepsilon/4$  w.h.p. (see Lemma 6). The information about the transition to the second phase affects only the agents at the leaves of the interval partition tree, corresponding to labels in  $[1, n\varepsilon]$ . When they get the message about the phase transition, they know that they can pass a label which is the sum of their own label and  $n$  to the first agent without label they interact with. For this reason, only the agents at the leaves corresponding to labels in  $[1, n\varepsilon]$  as well as the agents that are at the nodes that are ancestors of the aforementioned leaves participate in the broadcasting of the phase transition message. (Observe that the number of agents at these ancestors is  $O(n\varepsilon)$  and an agent at such an ancestor also has a label in  $[1, n\varepsilon]$ .) In the second phase, besides the agents at the leaves corresponding to labels in  $[1, n\varepsilon]$  and the agents without labels, also the agents at the intermediate nodes of the tree (if any) can really interact, in fact as in the first phase.

The following generalization of Lemma 1 is straightforward; see Appendix D for the proof.

► **Lemma 6.** *Let  $c$  be the constant from the statement of Lemma 1. During  $cn \log n/\varepsilon$  interactions in the first phase the number of agents without label drops below  $n\varepsilon/4$  w.h.p.*

Having Lemma 6, we can easily generalize Lemma 2 to the following one; see Appendix D for the proof.

► **Lemma 7.** *If the second phase starts after  $cn \log n/\varepsilon$  interactions, where  $c$  is the constant from Lemmata 1, 6, then only  $O(n \log n/\varepsilon)$  interactions are needed to assign labels in  $[1, (1 + \varepsilon)n]$  to the remaining agents without labels, w.h.p.*

We also need the following auxiliary lemma on broadcasting constrained to a subset of agents; see Appendix D for the proof.

► **Lemma 8.** *The leader can inform  $\Theta(n\varepsilon)$  agents with labels not exceeding  $O(n\varepsilon)$  about the phase transition using only these agents in  $O(n \log n/\varepsilon)$  interactions.*

The proof of the following theorem is analogous to that of Theorem 4 with Lemmata 1, 2 replaced by Lemmata 6, 7.

► **Theorem 9.** *Let  $\varepsilon > 0$ . There is a silent w.h.p. and safe protocol for a population of  $n$  agents that assigns unique labels in the range  $[1, (1 + \varepsilon)n]$  to  $n$  agents equipped with  $(2 + \varepsilon)n + O(n^c)$  states, for any positive  $c < 1$ , in  $O(n \log n/\varepsilon)$  interactions w.h.p.*

**Proof.** Under the assumption that the leader election preprocessing provides a unique leader, the correctness of the label assignment in both phases w.h.p. and the fulfillment of the definition of a silent and safe protocol follow from Lemmata 6, 7, and 8 by the same arguments as in the proof of Theorem 4.

The leader election preprocessing and its synchronization with the proper labeling protocol require  $O(n \log n)$  interactions and  $o(n^c)$  states, for  $c < 1$ , w.h.p. as described in the proof of Theorem 4. Analogously, it follows that w.h.p. the whole protocol provides a valid labeling, it is silent and safe. Again, it can be transformed to a safe protocol by prohibiting agents to change or get rid of an assigned label.

By Lemmata 6, 7, both phases require  $O(n \log n/\varepsilon)$  interactions w.h.p. The broadcasting about the phase transition starts when the number of agents without labels in the first phase drops below  $n\varepsilon/4$  w.h.p. By Lemma 8, it requires  $O(n \log n/\varepsilon)$  interactions w.h.p. since only the  $\Theta(n\varepsilon)$  agents in states corresponding to labels in  $[1, n\varepsilon]$  are involved in it.

The estimation of the number of needed states is more subtle than in Theorem 4. With each intermediate node of the interval partition tree that does not correspond to a label in  $[1, n\varepsilon]$  (equivalently, that is not an ancestor of a leaf corresponding to a label in  $[1, n\varepsilon]$ ), we associate a single state. (Recall here that if an agent at an intermediate node of the tree encounters an agent without label then the former agent moves to the left child of the node.) With each intermediate node corresponding to a label in  $[1, n\varepsilon]$ , we associate two states. They indicate whether or not the agent at the node has already got the message about phase transition. Next, with each leaf of the tree corresponding to a label  $i$  in  $[1, n\varepsilon]$ , we associate four states. They indicate whether or not the agent at the leaf has already got the message about the phase transition, and whether or not the agent has already passed the label  $i + n$  to some agent without label, respectively. To each of the remaining leaves, we associate only a single state.

We also need  $O(\log n/\varepsilon)$  additional states for the leader to count the number of own interactions in order to start broadcasting the message on transition to phase two at a right time step. In fact, we can get rid of the  $O(\frac{1}{\varepsilon})$  factor here by letting the leader to count approximately each  $\Theta(1/\varepsilon)$  interaction. Simply, the leader can count only interactions with agents which have got labels not exceeding  $O(\varepsilon n)$ .

Finally, we have  $n\varepsilon$  states corresponding to the labels in  $[n + 1, (1 + \varepsilon)n]$ . Thus, totally only  $(2 + O(\varepsilon))n + O(n^c)$  states, for any positive  $c < 1$ , are sufficient. To get rid of the constant factor at  $\varepsilon$ , it is sufficient to run the protocol for a smaller  $\varepsilon' = \Omega(\varepsilon)$ . It does not change the asymptotic upper bound on the number of required interactions w.h.p. and even it decreases the range of the labels. ◀

## 12:12 Efficient Assignment of Identities in Anonymous Populations

Note that  $\varepsilon$  in Theorem 9 does not have to be a constant; it can even be as small as  $O(n^{-1})$ .

By combining the protocol of Theorem 9 with that of Berenbrink et al. for exact counting the population size (Fact 6), we obtain the following corollary on unique labeling when the population size is unknown to agents initially. The proof is analogous to that of Corollary 5.

► **Corollary 10.** *Let  $\varepsilon > 0$ . There is a protocol for a population of  $n$  agents that assigns unique labels in the range  $[1, (1 + \varepsilon)n]$  to the agents initially not knowing the number  $n$ , equipped with  $\tilde{O}(n)$  states in  $O(n \log n / \varepsilon)$  interactions w.h.p.*

### 4 State- and range-optimal labeling

In this section we propose and analyze state-optimal protocols, which are silent and safe once a unique leader is elected, and utilize labels from the smallest possible range  $[1, n]$ . We assume the number of agents  $n$  to be known. We propose such a labeling protocol *Single-Cycle* which utilizes  $n + 5\sqrt{n} + O(n^c)$  states, for any positive  $c < 1$ , and the expected number of interactions required by the protocol is  $O(n^3)$ . We show in Section 5 that any silent and safe labeling protocol requires  $n + \sqrt{\frac{n-1}{2}} - 1$  states, see Theorem 14. Thus, our protocol is almost state-optimal. Finally, we propose a partial parallelization of *Single-Cycle* protocol called *k-Cycle* protocol which utilizes  $(1 + \varepsilon)n$  states and  $O((n/\varepsilon)^2)$  interactions for  $\varepsilon = \Omega(n^{-1/2})$ .

**Labeling protocol.** The state efficient labeling protocol starts from a preprocessing electing a unique leader. Its main idea is to use two agents: the initial leader  $A$  and a nominated (by  $A$ ) agent  $B$ , as partial *label dispensers*. These two agents jointly dispense unique labels for the remaining *free* (non-labeled yet) agents in the population where agent  $A$  dispenses the first and agent  $B$  the second part of each individual label. For the simplicity of presentation, we assume that  $n$  is a square of some integer. During execution of the protocol agent  $A$  uses partial labels  $\text{label}(a) \in \{0, \dots, \sqrt{n} - 1\}$  and  $B$  uses partial labels  $\text{label}(b) \in \{1, \dots, \sqrt{n}\}$ . The two dispensers label every agent by a unique pair of partial labels  $(\text{label}(a), \text{label}(b))$  where the combination  $(i, j)$  is interpreted as the integer label  $i \cdot \sqrt{n} + j$ . The protocol first labels all *free* (different to dispensers unlabeled) agents and eventually gives labels  $(0, 2)$  to agent  $B$  and  $(0, 1)$  to agent  $A$ .

In a nutshell, the labeling process is based on a single cycle of interactions between dispensers  $A$  and  $B$  and the free agents. Agent  $A$  awaits an interaction with a free agent  $F$  when  $A$  dispenses to  $F$  its current partial label  $\text{label}(a)$ . Now  $F$  awaits an interaction with  $B$  in order to receive the second part of its label. And when this happens agent  $F$  concludes with the combined label and agent  $B$  awaits an interaction with  $A$  to inform that the next free agent needs to be labeled. On the conclusion of this interaction if  $\text{label}(b) > 1$  agent  $B$  adopts new partial label  $\text{label}(b) - 1$ , otherwise  $B$  adopts  $\text{label}(b) = \sqrt{n}$  and agent  $A$  adopts new label  $\text{label}(a) - 1$ . The only exception is when  $\text{label}(a) = 0$  and  $\text{label}(b) = 2$  when agent  $B$  adopts label  $(0, 2)$  and agent  $A$  adopts label  $(0, 1)$  and both agents conclude the labeling process. For more details, see the definition of the transition function in Appendix E.

► **Theorem 11.** *Single-cycle utilizes  $n + 5 \cdot \sqrt{n} + O(n^c)$  states, for any positive  $c < 1$ , and the minimal label range  $[1, n]$ . The expected number of interactions required by the protocol is  $O(n^3)$ . Once a unique leader is elected, it produces a valid labeling of the  $n$  agents and it is silent and safe.*

**Proof.** Assume that the leader election preprocessing provides a unique leader. Then, the protocol is silent and safe by its definition. All  $1l$  labels are dispensed in the sequential manner and the labeling process concludes when the two dispensers finalize their own labels.

In particular, as soon as the two dispensers  $A$  and  $B$  are established they operate in a short cycle formed of steps  $C1, C2$  and  $C3$  labeling one by one all free agents in the population. One can observe that the sequence of cycles mimics the structure of two nested loops where the external loop iterates along the partial labels of  $A$  and the internal one along partial labels of  $B$ . In total, we have  $n - 2$  iterations where the expected number of interactions required by each iteration is  $O(n^2)$ . Thus one can conclude that the expected number of interactions required by the whole labeling process but for the leader election preprocessing is  $O(n^3)$ . By the definition of the protocol the range of assigned labels is  $[1, n]$ . Finally, as indicated earlier in this section the number of states utilized by the protocol but for the leader election preprocessing is equal to  $n + 5 \cdot \sqrt{n} + 4$ .

The leader election preprocessing and its synchronization with the proper labeling protocol require additional  $O(n \log n)$  interactions and additional  $o(n^c)$  states, for  $c < 1$ , w.h.p. as described in the proof of Theorem 4. ◀

Observe that when the exact value of  $n$  is embedded in the transition function on the conclusion all agents become dormant, i.e., they stop participating in the labeling process. One could redesign the protocol such that the labels are dispensed by  $A$  and  $B$  in the increasing order using a diagonal method, e.g.,  $(0, 0), (0, 1), (1, 0), (0, 2), (1, 1), (2, 0), (0, 3), (1, 2), (2, 1), (3, 0)$  etc., where agent  $A$  gets label  $(0, 0)$ , agent  $B$  gets label  $(0, 1)$ , the first labeled free agent gets  $(1, 0)$ , the second  $(0, 2)$ , then  $(1, 1)$  and  $(2, 0)$ , when  $A$  and  $B$  start using the next diagonal, etc. Each pair  $(i, j)$  is interpreted as  $(i + j)(i + j + 1)/2 + i$ , e.g.,  $(0, 1) = 1, (0, 2) = 3, (0, 3) = 6$  and in general  $(0, j) = j(j + 1)/2, (1, j - 1) = j(j + 1)/2 + 1, (1, j - 2) = j(j + 1)/2 + 2, \dots, (j, 0) = j(j + 1)/2 + j = (j + 1)(j + 2)/2 - 1 = (0, j + 1) - 1$ . In this case the size of the population does not need to be known in advance, however, the two dispensers will never stop searching for free agents yet to be labeled.

**Faster Labeling.** We observe that one can partially parallelize *Single-Cycle* protocol by instructing leader  $A$  to form  $k$  pairs of dispensers where each pair labels agents in a distinct range of size  $n/k$ . In such case the new  $k$ -cycle protocol requires extra  $2k$  states to allow leader  $A$  initialize the labeling process (create two dispensers) in all  $k$  cycles. Thus the total number of states is bounded by  $n + 2k + k \cdot (5\sqrt{n/k} + 4) = n + 6k + 5k \cdot \sqrt{n/k} < n + 6(k + \sqrt{nk}) < n + 12\sqrt{nk}$ , as  $k < \sqrt{nk}$ , plus the number of states required by the leader election preprocessing. We use the same method for the leader election preprocessing and its synchronization with the proper labeling protocol described in the proof of Theorem 4. Analogously, it adds  $O(n \log n)$  interactions and  $O(n^c)$  states, for any positive  $c < 1$ . As we need to pick  $k$  for which  $n + 12\sqrt{nk} \leq n + n\varepsilon$  we conclude that  $k \leq n\varepsilon^2/144$ .

One can show that for  $k = n\varepsilon^2/144$ , the expected number of interactions required by the  $k$ -cycle protocol is  $O(n^2/\varepsilon^2)$ . Note that in order to initialize  $k$  cycles the leader  $A$  has to communicate with  $2k - 1$  free agents. As  $k$  is at most a small fraction of  $n$  during the search for dispensers for each cycle the number of free agents is always greater than  $n/2$  (in fact it is very close to  $n$ ). Thus the probability of forming a new dispenser during any interaction is greater than  $1/2n$ , i.e., the product of the probability  $1/n$  that the random scheduler selects leader  $A$  as the initiator, times the probability greater than  $1/2$  that the responder is a free agent. In order to finish the initialization, we need to create new dispensers  $2k - 1$  times. Using Chernoff bound, we observe that after  $O(kn) = O(n^2/\varepsilon^2)$  interactions all  $k$  cycles have their two dispensers formed. As each cycle dispenses  $n/k = 144/\varepsilon^2$  labels and the expected number of interactions required to dispense a single label is  $O(n^2)$  with high probability, the expected number of interactions required by a specific cycle to generate all labels is  $O(n^2/\varepsilon^2)$  also with high probability. As observed earlier, the leader election preprocessing adds

## 12:14 Efficient Assignment of Identities in Anonymous Populations

only  $O(n \log n)$  interactions w.h.p. Hence, the expected number of interactions required to conclude the labeling process is  $O(n^2/\varepsilon^2)$ . Finally, note that for small values of  $\varepsilon$  approaching  $n^{-1/2}$   $k$ -cycle protocol reduces to *Single-cycle* protocol and for constant  $\varepsilon$  the number of interactions required by the protocol is  $O(n^2)$ .

► **Theorem 12.** *For  $k = n\varepsilon^2/144$ , where  $\varepsilon = \Omega(n^{-1/2})$ , and the minimal label range  $[1, n]$ , the proposed  $k$ -cycle labeling protocol provides a space-time trade-off in which utilization of  $(1 + \varepsilon)n + O(\log \log n)$  states permits the expected number of interactions  $O(n^2/\varepsilon^2)$ .*

### 5 Lower bounds

In this chapter, we derive several lower bounds on the number of states or the number of interactions required by silent, safe or the so-called pool protocols for unique labeling. Importantly, these lower bounds also hold in our model assuming that the population size is known to the agents initially and also when a unique leader is available initially.

The following general lower bound valid for any range of labels follows immediately from the definitions of a population protocol and the problem of unique labeling, respectively.

► **Theorem 13.** *The problem of assigning unique labels to  $n$  agents requires  $\Omega(n \log n)$  interactions w.h.p. and the agents have to be equipped with at least  $n$  states.*

**Proof.**  $\Omega(n \log n)$  interactions are needed w.h.p. since each agent has to interact at least once, see, e.g., the introduction in [10]. The lower bound on the number of states follows from the symmetry of agents, so any agent has to be prepared to be assigned an arbitrary label with at least a logarithmic bit representation. ◀

**A sharper lower bound on the number of states.** We obtain the following lower bound on the number of states required by a silent protocol which produces a valid labeling of the  $n$  agents and is safe w.h.p. The lower bound holds even if the protocol is provided with a unique leader and the knowledge of the number of agents. It almost matches the upper bound established in the previous section.

► **Theorem 14.** *A silent protocol which produces a valid labeling of the  $n$  agents and is safe with probability larger than  $1 - \frac{1}{n}$  requires at least  $n + \sqrt{\frac{n-1}{2}} - 1$  states. Also, if a silent protocol, which produces a valid labeling of the  $n$  agents and is safe with probability 1, uses  $n + t$  states, where  $t < n$ , then the expected number of interactions required by the protocol to provide a valid labeling is  $\frac{n^2}{t+1}$ .*

**Proof.** Let  $I$  be the set of ordered pairs of the  $n$  agents.  $I$  can be interpreted as the set of possible pairwise interactions between the agents.

Let  $Z$  be a finite run of the protocol, i.e., a finite sequence of pairs in  $I$ . Suppose that the execution of  $Z$  is successful, i.e., each agent reaches a final state with a distinct label, and no agent gets assigned two or more distinct labels during the run.

Let  $F_Z$  be the set of final states achieved by the agents after the execution of the run  $Z$ . We have  $|F_Z| = n$ . Also, let  $R_Z$  stand for the set of remaining states used in this run. Observe that if an agent is in a state in  $F_Z$  then it has a label.

For an agent  $x$ , let  $f_Z(x) \in F_Z$  be the last state achieved by the agent in the run  $Z$ , and let  $pred_Z(x)$  be the next to the last state achieved by the agent  $x$  in the run. Since for at most one agent the common initial state can be the final one,  $pred_Z(\cdot)$  is defined for at least  $n - 1$  agents. If  $pred_Z(x) \in F_Z$  and  $pred_Z(x)$  assigns a distinct label from that

assigned by  $f_Z(x)$  to  $x$  then we have a contradiction with our assumptions on  $Z$ . In turn, if  $pred_Z(x) \in F_Z$  and  $pred_Z(x)$  assign the same label as that assigned by  $f_Z(x)$  to  $x$  then we have a contradiction with the validity of the final labeling resulting from  $Z$ . We conclude that if  $pred_Z(x)$  is defined then  $pred_Z(x) \in R_Z$ .

Next, let  $A_Z$  be the set of agents  $x$  that achieved their final state in the run  $Z$  by an interaction of  $x$  in the state  $pred_Z(x)$  with an agent in a state in  $F_Z$ . For the proof of the following claim under the assumptions of the first statement in the theorem, see Appendix F.

▷ **Claim 15.** There is a finite run  $Z$  of the protocol such that after the execution of  $Z$ , each agent is in a final state with a distinct label, no single agent is assigned distinct labels during  $Z$ , and for any pair of distinct agents  $x, y \in A_Z$ ,  $pred_Z(x) \neq pred_Z(y)$ .

From here on, we assume that the run  $Z$  satisfies the claim. Consequently,  $|R_Z| \geq |A_Z|$ .

Let  $B_Z$  be the set of remaining agents that got their final state in  $F_Z$  in an interaction where both agents were in states outside  $F_Z$ , i.e., in  $R_Z$ . Since the agents in  $B$  achieved distinct final states with distinct labels in the aforementioned interactions, we infer that  $2|R_Z|^2 \geq |B_Z|$  and thus  $|R_Z| \geq \sqrt{|B_Z|/2}$ . Simply, there are  $|R_Z|^2$  ordered pairs of states in  $R_Z$ , and when agents in the states forming such a pair interact they can achieve at most two distinct states in  $F_Z$ . (Consequently, if  $2|R_Z|^2 < |B_Z|$  then there would be a pair of agents in  $B_Z$  that would achieve the same final state in the run and hence it would have the same label at the end of the considered run.)

Thus, we obtain  $|R_Z| \geq \max\{|A_Z|, \sqrt{\frac{n-1-|A_Z|}{2}}\} \geq \sqrt{\frac{n-1}{2}} - 1$  by straightforward calculations. This completes the proof of the first statement of the theorem.

To prove the second statement of the theorem, we need  $|R_Z| \geq |A_Z|$  to hold for any run  $Z$  resulting in a valid labeling of the agents without updating the label of any single agent. The existence of such a run  $Z$  implied by Claim 15 is not sufficient to obtain a lower bound on the expected number of required interactions. The stronger assumptions on the silent protocol in the second statement of the theorem requiring the protocol to provide always a valid labeling without updating the label of any single agent solves the problem. Namely, if  $pred_Z(x) = pred_Z(y)$  for  $x, y \in A_Z$  then following the notation and argumentation from the proof of Claim 15 neither  $Z_1i_1Z_2i_3$  nor any of its lengthening can provide a valid labeling without updating the label of any single agent. We obtain a contradiction with the aforementioned assumptions. Thus, the inequality  $|R_Z| \geq |A_Z|$  holds for arbitrary run  $Z$  ending with a valid labeling without updating the label of any single agent.

To prove the second statement, we may also assume w.l.o.g. that  $|A_Z| < n$  since otherwise  $t \geq |R_Z| \geq |A_Z| \geq n$ . Hence, the set  $B_Z$  of agents is non-empty. Let  $x$  be a last agent in  $B_Z$  that being in the state  $pred(x)$  gets its final state  $f(x)$  by an interaction with another agent  $y$  in a state  $s$ . If  $y$  belongs to  $B_Z$  then both  $x$  and  $y$  are the two last agents in  $B_Z$  that simultaneously get their final states in  $F_Z$  in the same interaction. The probability of the interaction between them is only  $\frac{1}{n^2}$ . Suppose in turn that  $y$  belongs to  $A_Z$ . We know that  $t \geq |R_Z| \geq |A_Z|$  from the previous part. Thus, there are at most  $t$  agents in  $B_Z$  in the state  $s$  with which the agent  $x$  in the state  $pred_Z(x)$  could interact. The probability of such an interaction is at most  $\frac{t}{n^2}$ . We conclude that the probability of an interaction between the agent  $x$  and the agent  $y$  after which  $x$  gets its final state  $f(x)$  is at most  $\frac{t+1}{n^2}$ , which proves the second statement. ◀

▶ **Corollary 16.** *If for  $\varepsilon > 0$ , a silent protocol that produces a valid labeling of the  $n$  agents and is safe with probability 1 uses only  $n + O(n^{1-\varepsilon})$  states then the expected number of interactions required by the protocol to achieve a valid labeling is  $\Omega(n^{1+\varepsilon})$ .*

**A lower bound for the range  $[1, n + r]$ .** Our fast protocols presented in Section 3 are examples of a class of natural protocols for the unique labeling problem that we term *pool protocols*.

In each step of a pool protocol, a subset of agents owns explicit or implicit pools of labels which are pairwise disjoint and whose union is included in the assumed range of labels. When two agents interact, they can repartition the union of their pools among themselves. Before the start of a pool protocol, only a single agent (the leader) owns a pool of labels. This initial pool corresponds to the assumed range of labels. An agent can be assigned a label from its own pool only. After that, the label is removed from the pool and cannot be changed. Finally, an agent without an assigned label cannot give away the whole own pool during an interaction with another agent without getting some part of the pool belonging to the other agent.

► **Theorem 17.** *The expected number of interactions required by a pool protocol to assign unique labels in the range  $[1, n + r]$ , where  $r \geq 0$ , to the population of  $n$  agents is at least  $\frac{n^2}{r+1}$ .*

**Proof.** We shall say that an agent has the  $P$  property if the agent owns a non-empty pool or a label has been assigned to the agent. Observe that if an agent accomplishes the  $P$  property during running a pool protocol then it never loses it. Also, all agents have to accomplish the  $P$  property sooner or later in order to complete the assignment task. During each interaction of a pool protocol at most one more agent can get the  $P$  property. Since at the beginning only one agent has the  $P$  property, there must exist an interaction after which only one agent lacks this property. By the disjointedness of the pools and labels, the assumed label range, and the definition of a pool protocol, there are at most  $r + 1$  agents among the remaining ones that could donate a sub-pool or label from own pool to the agent missing the  $P$  property. The expected number of interactions leading to an interaction between the agent missing the  $P$  property and one of the at most  $r + 1$  agents is  $\frac{n^2}{r+1}$ . ◀

## 6 Final remarks

Our upper bound of  $n + 5 \cdot \sqrt{n} + O(n^c)$ , for any positive  $c < 1$ , on the number of states achieved by a protocol for unique labeling that is silent and safe once a unique leader is elected almost matches our lower bound of  $n + \sqrt{\frac{n-1}{2}} - 1$ . We can combine our protocols for unique labeling with the recent protocols for counting or approximating the population size due to Berenbrink et al. [10] in order to get rid of the assumption that the population size is known to one of the agents initially. Since the aforementioned protocols from [10] either require  $\tilde{O}(n)$  states or  $O(n \log^2 n)$  interactions, the resulting combinations lose some of the near-optimality or optimality properties of our protocols (cf. Corollaries 5, 10). The related question if one can design a protocol for counting or closely approximating the population size simultaneously requiring  $O(n \log n)$  interactions w.h.p. and at most  $cn$  states, where  $c$  is a low constant, is of interest in its own right.

---

## References

- 1 D. Alistarh, H. Attiya, S. Gilbert, A. Giurgiu, and R. Guerraoui. Fast randomized test-and-set and renaming. In *Proc. of the International Symposium on Distributed Computing, DISC*, volume 6343 of *Lect. Notes in Comput. Sci.*, pages 94–108. Springer, 2010.
- 2 D. Alistarh, O. Denysyuk, L. Rodrigues, and N. Shavit. Balls-into-leaves: Sub-logarithmic renaming in synchronous message-passing systems. In *Proc. of the 2014 ACM Symposium on Principles of Distributed Computing, PODC*, pages 232–241. ACM, 2014.

- 3 D. Alistarh and R. Gelashvili. Recent algorithmic advances in population protocols. *ACM SIGACT News*, 49(3):63–73, 2018.
- 4 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.
- 5 D. Angluin, J. Aspnes, and D. Eisenstat. Fast computation by population protocols with a leader. In *Proc. of the International Symposium on Distributed Computing , DISC*, volume 4167 of *Lect. Notes in Comput. Sci.*, pages 61–75. Springer, 2006.
- 6 J. Aspnes, J. Beauquier, J. Burman, and D. Sohler. Time and space optimal counting in population protocols. In *Proc. of the 20th International Conference on Principles of Distributed Systems, OPODIS*, volume 70 of *LIPICs*, pages 13:1–13:17. Dagstuhl - LZI, 2016.
- 7 J. Beauquier, J. Burman, L. Rosaz, and B. Rozoy. Non-deterministic population protocols. In *Proc. of the 20th International Conference on Principles of Distributed Systems, OPODIS*, *Lect. Notes in Comput. Sci.*, pages 61–75. Springer, 2012.
- 8 P. Berenbrink, A. Brinkmann, R. Elsässer, T. Friedetzky, and L. Nagel. Randomized renaming in shared memory systems. In *Proc. of the IEEE Int. Parallel and Distributed Processing Symp., IPDPS*, pages 542–549. IEEE Computer Society, 2015.
- 9 P. Berenbrink, G. Giakkoupis, and P. Kling. Optimal time and space leader election in population protocols. In *Proc. of the 52nd ACM-SIGACT Symposium on Theory of Computing, STOC*, pages 119–129. ACM, 2020.
- 10 P. Berenbrink, D. Kaaser, and T. Radzik. On counting the population size. In *Proc. of the ACM Symposium on Principles of Distributed Computing, PODC*, pages 43–52. ACM, 2019.
- 11 J. Burman, J. Beauquier, and D. Sohler. Space-optimal naming in population protocols. In *Proc. of the 33rd International Symposium on Distributed Computing, DISC*, volume 146 of *LIPICs*, pages 9:1–9:16. Dagstuhl - LZI, 2019.
- 12 J. Burman, H. Chen, H. Chen, D. Doty, T. Nowak, E. Severson, and C. Xu. Time-optimal self-stabilizing leader election in population protocols. In *Proc. of the ACM Symposium on Principles of Distributed Computing, PODC*, pages 33–44. ACM, 2021.
- 13 S. Cai, T. Izumi, and K. Wada. How to prove impossibility under global fairness: On space complexity of self-stabilizing leader election on a population protocol model. *Theory of Computing Systems*, 50:433–445, 2012.
- 14 A. Castañeda, S. Rajsbaum, and M. Raynal. The renaming problem in shared memory systems: An introduction. *Comput. Sci. Rev.*, 5(3):229–251, 2011.
- 15 S. Dolev, M.G. Gouda, and M. Schneider. Memory requirements for silent stabilization. *Acta Inf.*, 36(6):447–462, 1999.
- 16 D. Doty and M. Eftekhari. A survey of size counting in population protocols. *Theoretical Computer Science*, 894:91–102, 2021.
- 17 D. Doty, M. Eftekhari, O. Michail, P. G. Spirakis, and M. Theofilatos. Exact size counting in uniform population protocols in nearly logarithmic time. *ArXiv*, 2018. preprint. [arXiv: 1805.04832](https://arxiv.org/abs/1805.04832).
- 18 R. Elsässer and T. Radzik. Recent results in population protocols for exact majority and leader election. *Bull. EATCS*, 126, 2018.
- 19 L. Gąsieniec and G. Stachowiak. Enhanced phase clocks, population protocols, and fast space optimal leader election. *J. ACM*, 68(1):2:1–2:21, 2021.

## **A** The computational model of population protocols

There is given a population of  $n$  agents that can pairwise interact in order to change their states and in this way perform a computation. A population protocol can be formally specified by providing a set  $Q$  of possible states, a set  $O$  of possible outputs, a transition function  $\delta : Q \times Q \rightarrow Q \times Q$ , and an output function  $o : Q \rightarrow O$ . The current state  $q \in Q$  of an agent is updated during interactions. Consequently, the current output  $o(q)$  of the agent also becomes updated during interactions. The current state of the set of  $n$  agents is given by a vector in  $Q^n$  with the current states of the agents. A computation of a population protocol

## 12:18 Efficient Assignment of Identities in Anonymous Populations

is specified by a sequence of pairwise interactions between agents. In every time step, an ordered pair of agents is selected for interaction by a probabilistic scheduler independently and uniformly at random. The first agent in the selected pair is called the initiator while the second one is called the responder. The states of the two agents are updated during the interaction according to the transition function  $\delta$ .

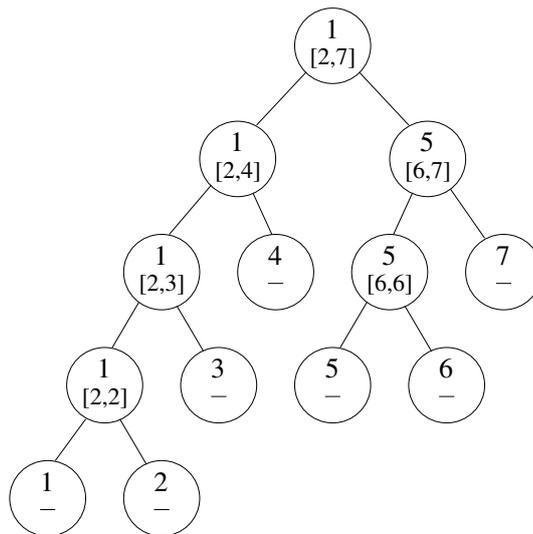
We can specify a problem to solve by a population protocol by providing the set of input configurations, the set  $O$  of possible outputs, and the desired output configurations for given input configurations. For the unique labeling problem, all agents are initially in the same state  $q_0$ . The set  $O$  is just the set of positive integers. A desired configuration is when all agents output their distinct labels. The *stabilization time* of an execution of a protocol is the number of interactions until the states of agents form a desired configuration from which no sequence of pairwise interactions can lead to a configuration outside the set of desired configurations.

### B Related work (Table 3)

■ **Table 3** Upper bounds on the number of interactions, the number of states and the range used by the known labeling protocols. In case of the self-stabilizing labeling protocols in [12], the “safe” property can eventually hold only for their initialized versions.

$n$	# interactions	# states	Range	Properties	Paper
unknown	$O(n^3)$ w.h.p.	$n$	$[1, n]$	silent	[13]
unknown	$O(n \log n \log \log n)$ w.h.p.	$n^{O(1)}$	$[1, n^{O(1)}]$	silent	[17]
known	$O(n^2)$ expected	$O(n)$	$[1, n]$	silent, safe	[12]
known	$O(n \log n)$ w.h.p.	$\exp(O(n^{\log n} \log n))$	$[1, n]$	safe	[12]

### C Figure 1



■ **Figure 1** An example of the partition tree of the start interval  $[1, 7]$ .

## D Labeling with asymptotically optimal number of interactions within $[1, (1 + \epsilon)n]$ (proofs)

**Proof of Lemma 6.** The proof is a generalization of that for Lemma 1. Define  $F_\epsilon$  as a set of at least  $\epsilon n/4$  agents without labels that survive at least  $cn \log / \epsilon$  interactions in the first phase. Note that for an arbitrary agent, the probability of interaction with a member in  $F_\epsilon$  is at least  $\frac{\epsilon}{4n}$ . The rest of the proof is analogous to that of Lemma 1. It is sufficient to replace  $F$  by  $F_\epsilon$  and the probability  $\frac{1}{4n}$  of an interaction with a member in  $F$  with that  $\frac{\epsilon}{4n}$  of an interaction with a member in  $F_\epsilon$ . ◀

**Proof of Lemma 7.** The number of agents without labels at the beginning of the second phase is smaller than  $\epsilon n/4$  w.h.p. Hence, at the beginning of the second phase the number of agents with labels in the range  $[1, \epsilon n]$  is at least  $\frac{3\epsilon n}{4}$  w.h.p. Recall that such an agent at a leaf of the tree can give a label to an agent without label only once. It follows that the number of agents with labels in  $[1, \epsilon n]$  that can give a label to an agent without label is always at least  $\frac{\epsilon n}{2}$  w.h.p. We conclude that for an agent without label the probability of an interaction with an agent that can give a label is at least almost  $\frac{\epsilon}{2n}$ . Hence, after each  $O(n/\epsilon)$  interactions the expected number of agents without labels halves. It follows that the expected number of such interactions rounds is  $O(\log n)$ . Consequently, the number of the rounds is also  $O(\log n)$  w.h.p. by Fact 2.

An alternative way to obtain the  $O(n \log n/\epsilon)$  bound on the number of interactions w.h.p. is to use Fact 3 analogously as in the proof of Lemma 2. The difference is that  $C$  is set to  $O(\frac{2}{\epsilon})$  instead of  $O(2)$  since the set of agents that can give a label is of size at least  $\frac{n\epsilon}{2}$  now. ◀

**Proof of Lemma 8.** During the initial part of the broadcasting process, after every  $O(n/\epsilon)$  interactions, the expected number of agents participating in the broadcasting process doubles. Hence, after  $O(n \log / \epsilon)$  interactions, the expected number of informed agents will be  $\Omega(n\epsilon)$ . Then, the expected number of uninformed agents will be halved for every  $O(n/\epsilon)$  interactions. So the expected number of rounds, each consisting of  $O(n/\epsilon)$  interactions, needed to complete the broadcasting is  $O(\log n)$ . It remains to turn the latter bound to a w.h.p. one. This can be done by using the Chernoff bounds (Fact 2).

Alternatively, we can define for the purpose of the analysis of the doubling part, a binary broadcast tree. An informed agent at an intermediate node of the tree after an interaction with an uninformed agent moves to a child of the node while the other agent now informed places at the other child (cf. the partition tree in the proofs of Lemmata 1, 6). Then, we can use the technique from the proofs of Lemmata 1, 6 to show that only  $O(n \log n/\epsilon)$  interactions are required w.h.p. to achieve a configuration where only a constant fraction of the agents participating in the broadcasting is uninformed. To derive the same asymptotic upper bound on the number of interactions required by the halving part w.h.p., we can use Fact 3 with  $C = O(\epsilon^{-1})$  analogously as in the proofs of Lemmata 2, 7. ◀

## E The transition function of the state optimal protocol

**State utilization in *Single-Cycle* protocol.**

- [Agent  $A$ ] Since  $\text{label}(a) \in \{0, \dots, \sqrt{n}-1\}$  dispenser  $A$  utilizes  $2 \cdot \sqrt{n} + 2$  states including:
- $A.\text{init} = (1)$  # the initial (leadership) state of dispenser  $A$ ,
  - $A[\text{label}(a), \text{await}(F)]$  # dispenser  $A$  carrying partial label  $\text{label}(a)$  awaits interaction with a free agent  $F$ ,

## 12:20 Efficient Assignment of Identities in Anonymous Populations

- $A[\text{label}(a), \text{await}(B)]$  # dispenser  $A$  carrying partial label  $\text{label}(a)$  awaits interaction with dispenser  $B$ ,
- $A.\text{final} = (0, 1)$  # the final state of  $A$ .

[**Agent  $B$** ] Since  $\text{label}(b) \in \{0, \dots, \sqrt{n}\}$  dispenser  $B$  utilizes  $2 \cdot \sqrt{n} + 3$  states including:

- $B[\text{label}(b), \text{await}(F)]$  # dispenser  $B$  carrying partial label  $\text{label}(b)$  awaits interaction with a free agent  $F$ ,
- $B[\text{label}(b), \text{await}(A)]$  # dispenser  $B$  carrying partial label  $\text{label}(b)$  awaits interaction with dispenser  $A$
- $B.\text{final} = (0, 2)$  # the final state of  $B$ .

[**Agent  $F$** ] Since free agents carry partial labels  $\text{label}(a) \in \{0, \dots, \sqrt{n} - 1\}$  and eventually adopt one of the  $n - 2$  destination labels (excluding dispensers) they utilize  $n + \sqrt{n} - 1$  states including:

- $F.\text{init} = (0)$  # the initial (non-leader) state of  $F$
- $F[\text{label}(a), \text{await}(B)]$  # free agent  $F$  carrying partial label  $\text{label}(a)$  awaits interaction with dispenser  $B$ ,
- $F.\text{final} = (\text{label}(a), \text{label}(b))$  # the final state of  $F$ .

In total *Single-Cycle* protocol requires  $n + 5 \cdot \sqrt{n} + 4$  states.

### Transition function in *Single-Cycle* protocol.

**Step 0: Initialization.** During the first interaction of  $A$  with a free agent the second dispenser  $B$  is nominated. Both dispensers adopt their largest labels. Agent  $A$  awaits a free agent in the initial state while agent  $B$  awaits a free agent carrying a partial label obtained from  $A$ .

- $(A.\text{init}, F.\text{init})$   
 $\rightarrow (A[\text{label}(a) = \sqrt{n} - 1, \text{await}(F)], B[\text{label}(b) = \sqrt{n}, \text{await}(F)]),$

The three steps  $C_1, C_2$ , and  $C_3$  of the labeling cycle are given below.

**Step  $C_1$ : Agent  $A$  dispenses partial label.** During an interaction of agent  $A$  with a free agent  $F$  the current partial label  $\text{label}(a)$  is dispensed to  $F$ . Both agents await interactions with dispenser  $B$  which is ready to interact with partially labeled  $F$  but not  $A$ .

- $(A[\text{label}(a), \text{await}(F)], F.\text{init})$   
 $\rightarrow (A[\text{label}(a), \text{await}(B)], F[\text{label}(a), \text{await}(B)])$  # Go to Step  $C_2$

**Step  $C_2$ : Agent  $B$  dispenses partial label.** During an interaction of agent  $B$  with a free agent  $F$  which carries partial label  $\text{label}(a)$ , the complementary current partial label  $\text{label}(b)$  is dispensed to  $F$ . Agent  $F$  concludes in the final state with the combined label  $(\text{label}(a), \text{label}(b))$ . Agent  $B$  is now ready for interaction with  $A$ .

- $(B[\text{label}(b), \text{await}(F)], F[\text{label}(a), \text{await}(B)])$   
 $\rightarrow (B[\text{label}(b), \text{await}(A)], F.\text{final} = (\text{label}(a), \text{label}(b)))$  # Go to Step  $C_3$

**Step  $C_3$ : Agent  $A$  and  $B$  negotiate a new label or conclude.** In the case when  $\text{label}(a) = 0$  and  $\text{label}(b) = 2$  the dispensers  $A$  and  $B$  conclude in states  $(0, 1)$  and  $(0, 2)$  respectively, see the first transition. Otherwise a new combination of partial labels is agreed and the protocol goes back to Step  $C_1$ .

- $(A[\text{label}(a) = 0, \text{await}(B)], B[\text{label}(b) = 2, \text{await}(A)])$   
 $\rightarrow (A.\text{final} = (0, 1), B.\text{final} = (0, 2))$  # Conclude the labeling process

- $(A[\text{label}(a) = 0, \text{await}(B)], B[\text{label}(b) > 2, \text{await}(A)])$  or  
 $(A[\text{label}(a) > 0, \text{await}(B)], B[\text{label}(b) > 1, \text{await}(A)])$   
 $\rightarrow (A[\text{label}(a), \text{await}(F)], B[\text{label}(b) - 1, \text{await}(F)])$  # Go to Step  $C_1$
- $(A[\text{label}(a) > 0, \text{await}(B)], B[\text{label}(b) = 1, \text{await}(A)])$   
 $\rightarrow (A[\text{label}(a) - 1, \text{await}(F)], B[\text{label}(b) = \sqrt{n}, \text{await}(F)])$  # Go to Step  $C_1$

## F Lower bounds (proofs)

**Proof of Claim 15.** The proof of the claim is by a contradiction with the assumptions on the labeling protocol. The general intuition is that if  $\text{pred}_Z(x) = \text{pred}_Z(y)$  for two agents  $x, y \in A_Z$  then we can associate with a prefix of  $Z$  a slightly modified equally likely run  $Z'$  which assigns the same label to a pair of agents. Hence, the modified run does not produce a valid labeling or it has to assign at least two different labels to some agent.

To obtain the contradiction, we assume that for each finite run  $Z$  in which the agents achieve final states with distinct labels without assigning distinct labels to any single agent during the run, there is a pair of agents  $x, y \in A_Z$ , where  $\text{pred}_Z(x) = \text{pred}_Z(y)$ . Let us consider such a pair of agents  $x, y \in A_Z$  that minimizes the length of the prefix of  $Z$  in which both agents achieve their final states in  $F_Z$ . We may assume w.l.o.g. that  $x$  gets its final state  $f_Z(x)$  in an interaction  $i_1$  with an agent  $x'$  that is in a state in  $F_Z$ , and in a later interaction  $i_2$ ,  $y$  gets its final state  $f_Z(y)$ , in the run  $Z$ . (Note that  $x'$  cannot be in a final state different from its own, i.e., in  $F_Z \setminus \{f_Z(x')\}$  since this would require updating its label contradicting the assumption on  $Z$ .) Thus, the shortest prefix of  $Z$  in which both  $x$  and  $y$  get their final states has the form  $Z_1 i_1 Z_2 i_2$ . Then, if we replace the latter interaction  $i_2$  by the interaction  $i_3$  between  $y$  and the agent  $x'$  in the state  $f_Z(x')$  analogous to  $i_1$ , it will result in achieving by  $y$  the state  $f_Z(x)$  since  $\text{pred}_Z(x) = \text{pred}_Z(y)$ . Thus, neither the run  $Z_1 i_1 Z_2 i_3$  nor any of its extensions can yield a valid labeling of the agents without updating labels for some of them. Importantly, the runs  $Z_1 i_1 Z_2 i_2$  and  $Z_1 i_1 Z_2 i_3$  are equally likely (\*).

We initialize two sets  $S_{\text{valid}}$  and  $S_{\text{invalid}}$  of strings (sequences) over the alphabet  $I$ . Then, for each run  $Z$  in which the agents achieve final states with distinct labels without updating the label of any single agent, we insert the prefix  $Z_1 i_1 Z_2 i_2$  into  $S_{\text{valid}}$  and the corresponding sequence  $Z_1 i_1 Z_2 i_3$  into  $S_{\text{invalid}}$ . Note that by the choice of  $i_1, i_2$ , no string in  $S_{\text{valid}}$  is a prefix of another string in  $S_{\text{valid}}$ . The analogous property holds for  $S_{\text{invalid}}$ . By the construction of the sets, each run  $Z$  in which the agents achieve final states with distinct labels without updating the label of any single agent has to overlap with or be a lengthening of a string in  $S_{\text{valid}}$ . Furthermore, no run of the protocol that overlaps with a string in  $S_{\text{invalid}}$  or it is a lengthening of a string in  $S_{\text{invalid}}$  results in a valid labeling without updating the label of any single agent. Define the function  $g : S_{\text{valid}} \rightarrow S_{\text{invalid}}$  by  $g(Z_1 i_1 Z_2 i_2) = Z_1 i_1 Z_2 i_3$ . By the property (\*), the probability that a string over  $I$  is equal to  $Z_1 i_1 Z_2 i_2$  or it is a lengthening of  $Z_1 i_1 Z_2 i_2$  is not greater than the probability that a string over  $I$  is equal to  $g(Z_1 i_1 Z_2 i_2)$  or it is a lengthening of  $g(Z_1 i_1 Z_2 i_2)$ . The function  $g$  is not necessarily a bijection. Suppose that  $g(Z_1 i_1 Z_2 i_2) = g(Z'_1 i'_1 Z'_2 i'_2)$ . Then, we have  $Z_1 i_1 Z_2 i_3 = Z'_1 i'_1 Z'_2 i'_3$ . Consequently, the strings  $Z_1 i_1 Z_2 i_2$  and  $Z'_1 i'_1 Z'_2 i'_2$  may only differ in the last interaction, i.e.,  $i_2$  may be different from  $i'_2$ . However,  $i_2$  and  $i'_2$  have to include the same agent ( $y$  in the earlier construction) that appears in  $i_3$ . We conclude that the aforementioned two strings in  $S_{\text{valid}}$  can differ by at most one agent in the last interaction. It follows that  $g$  maps at most  $n - 1$  strings in  $S_{\text{valid}}$  to the same string in  $S_{\text{invalid}}$ . Hence, the event that the agents eventually achieve their final states yielding a valid labeling without updating the label of any single agent is at most  $n - 1$  times more likely than the complement event, contradicting our assumptions. ◀