

Almost logarithmic-time space optimal leader election in population protocols

Leszek Gąsieniec¹, Grzegorz Stachowiak², and Przemysław Uznański³

¹University of Liverpool, UK

²University of Wrocław, Poland

³ETH Zürich, Switzerland

Abstract

The model of population protocols refers to a large collection of simple indistinguishable entities, frequently called *agents*. The agents communicate and perform computation through pairwise interactions. We study fast and space efficient leader election in population of cardinality n governed by a random scheduler, where during each time step the scheduler uniformly at random selects for interaction exactly one pair of agents.

We present the first $o(\log^2 n)$ -time leader election protocol. It operates in expected parallel time $\mathcal{O}(\log n \log \log n)$ which is equivalent to $\mathcal{O}(n \log n \log \log n)$ pairwise interactions. This is the fastest currently known leader election algorithm in which each agent utilises asymptotically optimal number of $\mathcal{O}(\log \log n)$ states. The new protocol incorporates and amalgamates successfully the power of assorted *synthetic coins* with variable rate *phase clocks*.

Regular SPAA submission

Contact author: Leszek Gąsieniec, l.a.gasieniec@liverpool.ac.uk

1 Introduction

The computational model of *population protocols* was introduced in the seminal paper by Angluin *et al.* [AAD⁺04]. Their model provides a universal platform for the formal analysis of pairwise interactions within a large collection of indistinguishable entities, frequently referred to as *agents*. In this model the agents rely on very limited communication and computation power. The actions of agents are prompted by their pairwise interactions with the outcome determined by a finite state machine \mathcal{F} . When two agents engage in an interaction they cross-examine the content of their local states, and on the conclusion of this encounter their states change according to the transition function forming an integral part of \mathcal{F} . A population protocol terminates with success when eventually all agents stabilise w.r.t. the output (which depends only on their states).

The number of states utilised by the finite state machine \mathcal{F} constitutes the *space complexity* of the protocol. In the *probabilistic variant* of population protocols, introduced in [AAD⁺04] and used in this paper, in each step the interacting pair of agents is chosen uniformly at random by the *random scheduler*. In this variant one is also interested in the *time complexity*, i.e., the time needed to stabilise (converge) the protocol. More recently the studies on population protocols focus on performance in terms of *parallel time* defined as the total number of pairwise interactions (leading to stabilisation) divided by the size of the population. The parallel time can be also interpreted as the local time observed by agents proportional to the number of interactions it participates in.

Populations protocols attracted studies on several central problems in distributed computing. This includes the *majority* problem, a special instance of *consensus* [Fis83], where the final configuration of states must indicate the larger fraction of the population. The first attempt to computing majority with population protocols can be found in [AAD⁺04]. Later, a neat 3-state one-way protocol for approximate majority was given in [AAE08a]. In more recent work [AGV15] Alistarh *et al.* consider time-precision trade-offs in exact majority population protocols. Further studies on time-space trade-offs can be found in [AAE⁺17, BCER17, BEF⁺18] and [AAG18], where in the latter an asymptotically space-optimal protocol is given. The convergence (stabilisation) of majority protocols was also studied in more specific network topologies [DV12, GHMS15, MNRS14], as well as in the deterministic setting [GHM⁺16, MNRS14]. A useful survey [MCS11] discusses a range of combinatorial problems suitable for population protocols.

In this paper we study leader election problem where in the final configuration a unique agent must converge to a *leader state* and every other agent has to stabilise in a *follower state*. While the problem is quite well understood and represented in the literature only recently it received greater attention in the context of population protocols, partly due to several developments in a related model of chemical reactions [CCDS14, Dot14]. In particular, in the follow-up work of Doty and Soloveichik [DS15] we learn that leader election cannot be solved in sublinear time when agents are equipped with a fixed (constant) number of states. On the other hand Alistarh and Gelashvili [AG15] proposed an alternative leader election protocol operating in time $\mathcal{O}(\log^3 n)$ and utilising $\mathcal{O}(\log^3 n)$ states. In more recent work [AAE⁺17] Alistarh *et al.* consider a trade-off between the number of states utilised by agents and the time complexity of the solution. They provide a separation argument distinguishing between slowly stabilising population protocols which utilise $o(\log \log n)$ states and rapidly stabilising protocols requiring $\Omega(\log \log n)$ states. This result coincides nicely with another fundamental observation due to Chatzigiannakis *et al.* [CMN⁺11] which shows that population protocols utilising $o(\log \log n)$ states can only cope with semi-linear predicates while presence of $\mathcal{O}(\log n)$ states enables computation of symmetric predicates. Another recent development includes a protocol which elects the leader in time $\mathcal{O}(\log^2 n)$ whp and in expectation utilising $\mathcal{O}(\log^2 n)$ states per agent [BCER17]. The number of states was later reduced to $\mathcal{O}(\log n)$ by Alistarh, Aspnes, Gelashvili in [AAG18] and by Berenbrink *et al.* in [BKKO18] through the application of two types of synthetic coins. Please refer to Table 1 for the summary of past results.

The recent progress in leader election is also aligned with an improved understanding of *phase clocks* capable of counting parallel time approximately. The relevant work includes *leaderless* phase clocks proposed by Alistarh *et al.* [AAG18] and *junta-driven* phase clocks utilised in $\mathcal{O}(\log^2 n)$ -time space-optimal leader election algorithm [GS18]. The concept of phase clocks is also closely related to *oscillators* used to model behaviour of periodic dynamic systems. In [CGK⁺15] Czyzowicz *et al.* provide a thorough study of 3-state oscillators in Lotka-Volterra systems, and in further work [DK18] Dudek and Kosowski consider information dissemination with authoritative sources.

Our results: In this paper we propose the first leader election protocol stabilising in time $o(\log^2 n)$

assuming the asymptotically optimal number of states at each agent. More precisely, we propose a new $\mathcal{O}(\log n \log \log n)$ -time protocol in which each agent operates on $\mathcal{O}(\log \log n)$ states. The solution is always correct but the improved performance refers to the expected time, i.e., the high probability is guaranteed only in time $\mathcal{O}(\log^2 n)$ as in [GS18]. We would like to emphasise that our new algorithm was the first space-efficient population protocol which broke the $\mathcal{O}(\log^2 n)$ time barrier in leader election, see [GSU18]. Similarly to other protocols using non-constant number of states is *non-uniform*, i.e., it requires some rough knowledge of n , e.g., to set the size of the phase clock. Very recently Sudo et al. in [SOI⁺18] proposed a $\mathcal{O}(\log n)$ expected time leader election protocol which assumes $\mathcal{O}(\log n)$ states at each node. **Protocol overview:** Our main aim is to further speed up the space optimal protocol in [GS18]. The first part of the protocol in [GS18] elects a junta of $n^{1-\varepsilon}$ leaders in parallel time $\mathcal{O}(\log n)$ whp utilising the optimal number of $\Theta(\log \log n)$ states. The second part reduces the junta to one leader by repeatedly flipping coins and multi-broadcasting the results by the leader candidates $\Theta(\log n)$ times. This second part is implemented using phase clocks powered by the junta in parallel time $\mathcal{O}(\log^2 n)$ whp, and utilising only a constant memory in each agent.

Thus the main challenge was to overcome the time complexity bottleneck of the second part, which is now replaced by a faster protocol utilising $\Theta(\log \log n)$ states. Please note also that here the implementation detail differs from [GS18]. The first part is executed on a subpopulation different to the one from which the actual leader will be elected. This is to maintain the desired bound $\Theta(\log \log n)$ on states utilised by agents. In addition, the new protocol which replaced the second part is split into two epochs.

The first epoch is responsible for fast elimination which reduces the number of leader candidates to $\mathcal{O}(\log n)$. This is done with the help of levels assigned to each agent in the first part generating $\log \log n$ asymmetric coins. The result of flipping coin at level i is 1 if the relevant agent interacts with another on a level larger than i , and 0 otherwise. Each of these coin flips applied subsequently a constant number of times to the leader candidate subpopulation reduces its size to $\mathcal{O}(\log n)$ in time $\mathcal{O}(\log n \log \log n)$ whp.

The second epoch is the final elimination with the goal to reduce the number of leader candidates from $\mathcal{O}(\log n)$ to one. Unfortunately the method proposed in [GS18] does not allow to select one leader, even from a pack of two, in parallel time substantially better than $\log^2 n$ utilising small memory. In contrast, the elimination by repeated coin flips picks one of the two leader candidates in expected constant number of iterations. This gives expected parallel time $\mathcal{O}(\log n)$ for two remaining candidates. This observation leads to the design of a Las Vegas type protocol which always computes a single leader and works in expected time $\mathcal{O}(\log n \log \log n)$. On the conclusion, when a single leader is elected its can still move between different leader states, as the followers can hover between different follower states. The main difficulty in this epoch is to avoid full elimination of leaders candidates, as a direct application of repeated coin flips may end up in total elimination. And this would not be allowed in a Las Vegas type algorithm. **Methods used:** The new leader election algorithm utilises partition of all agents into three subpopulations including *coins* (C) responsible for generation of asymmetric coins with $\log \log n$ bias levels, *leaders* (L) among which the unique leader is eventually drawn, and *inhibitors* (I) designated to maintain variable-rate phase clocks. A division into sub-populations to reduce space usage at and to accomodate for different roles of agents was previously used in [AAG18, GPI6]. The actions of agents in our new protocol are synchronised by phase clocks. The first application of phase clocks (goverened by a unique leader) in population protocols refers to [AAE08b]. More recently, a novel concept of leader-less phase clocks was introduced in [AAG18], and two-level phase clocks driven by junta (group of agents) were utilised successfully in [GS18]. The latter synchronisation mechanism is also adopted in this paper. We also use here synthetic coin-flips as the main symmetry breaking mechanism, where the extraction source of random bits refers to the scheduler. In [AAE⁺17] the authors explored protocols based on uniform coins while in [BKKO18] one can find the first use of coins with non-constant bias. This paper proposes the first approach in which a larger spectrum of $(\log \log n)$ coins, varying in degree of asymmetry and characterised by $\text{poly}(n)$ bias, is utilised.

Please note that while the fast leader election algorithm presented here builds upon some ideas from [GS18] the significant time improvement is feasible due to several new developments which come in different “styles and flavours”. Firstly, we need biased coins which allow us to elect in time $\mathcal{O}(\log n \log \log n)$ a small (logarithmic size) family of active leader candidates. Also, during the execution of the protocol we handle low probability “out-of-sync” errors differently by creating a family of $\mathcal{O}(\log \log n)$ independent guarantees which certify that at least one active leader candidate remains. We also need to guarantee unique leader election during the final reduction stage. And while the adopted coin-flipping mechanism enables the relevant reduction in expected time $\mathcal{O}(\log n \log \log n)$, we also have to certify the correctness

Paper	States	Time	Runtime
[AG15]	$\mathcal{O}(\log^3 n)$	$\mathcal{O}(\log^3 n)$ $\mathcal{O}(\log^4 n)$	expected w.h.p.
[AAE ⁺ 17]	$\mathcal{O}(\log^2 n)$	$\mathcal{O}(\log^{5.3} n \cdot \log \log n)$ $\mathcal{O}(\log^{6.3} n)$	expected w.h.p.
[BCER17]	$\mathcal{O}(\log^2 n)$	$\mathcal{O}(\log^2 n)$	w.h.p.
[AAG18]	$\mathcal{O}(\log n)$	$\mathcal{O}(\log^2 n)$	expected
[BKKO18]	$\mathcal{O}(\log n)$	$\mathcal{O}(\log^2 n)$	w.h.p.
[GS18]	$\mathcal{O}(\log \log n)$	$\mathcal{O}(\log^2 n)$	w.h.p.
This work	$\mathcal{O}(\log \log n)$	$\mathcal{O}(\log n \cdot \log \log n)$	expected
[SOI ⁺ 18]	$\mathcal{O}(\log n)$	$\mathcal{O}(\log n)$	expected

Table 1: Leader election via population protocols.

of the reduction process using extra $\mathcal{O}(\log \log n)$ states. We achieve this goal by creating a family of $\log \log n$ consecutive signals which are used as certification points during the execution of the protocol. The expected delay between different signals is growing exponentially. We believe this construction is new in the context of population protocols.

Related work: Leader election is one of the fundamental problems in Distributed Computing besides broadcasting, mutual-exclusion, consensus, see, e.g., an excellent text book by Attiya and Welch [AW04]. The problem was originally studied in networks with nodes having distinct labels [Lan77], where an early work focuses on the ring topology in synchronous [FL87, HS80] as well as in asynchronous models [Bur80, Pet82]. Also, in networks populated by mobile agents the leader election was studied first in networks with labelled nodes [HKM⁺08]. However, very often leader election is used as a powerful symmetry breaking mechanism enabling feasibility and coordination of more complex protocols in systems based on uniform (indistinguishable) entities. There is a large volume of work [Ang80, AS91, ASW88, BSV⁺96, BV99, YK89, YK96] on leader election in anonymous networks. In [YK89, YK96] we find a good characterisation of message-passing networks in which leader election is feasible when the nodes are anonymous. In [YK89], the authors study the problem of leader election in general networks under the assumption that node labels are not unique. In [FKK⁺04], the authors study feasibility and message complexity of leader election in rings with possibly non-unique labels, while in [DP04] the authors provide solutions to a generalised leader election problem in rings with arbitrary labels. The work in [FP11] focuses on the time complexity of leader election in anonymous networks where this complexity is expressed in terms of multiple network parameters. In [DP14], the authors study feasibility of leader election for anonymous agents that navigate in a network asynchronously. Another important study on trade-offs between the time complexity and knowledge available in anonymous trees can be found in recent work of Glacet *et al.* [GMP16]. Finally, a good example of recent extensive studies on the exact space complexity in related models refers to plurality consensus. In particular, in [BFGK16] Berenbrink *et al.* proposed a plurality consensus protocol for C original opinions converging in $\mathcal{O}(\log C \log \log n)$ synchronous rounds using only $\log C + \mathcal{O}(\log \log C)$ bits of local memory. They also show a slightly slower solution converging in $\mathcal{O}(\log n \log \log n)$ rounds utilising only $\log C + 4$ bits of the local memory. They also pointed out that any protocol with local memory $\log C + \mathcal{O}(1)$ has the worst-case running time $\Omega(k)$. In [GP16] Ghaffari and Parter propose an alternative algorithm converging in time $\mathcal{O}(\log C \log n)$ in which all messages and the local memory are bounded to $\log C + \mathcal{O}(1)$ bits. Some work on utilisation of random walk in plurality consensus can be found in [BCN⁺15, GHMS15]. More information can be found in two recent surveys on algorithmic advances in population protocols [AG18, ER18].

2 Preliminaries

We study here protocols defined on populations of identical agents in which a dedicated *random scheduler* connects (sequentially or in parallel) agents in pairs *uniformly at random*. We assume that all n agents start in the same initial state. We adopt the classical model of population protocols [AAD⁺04, AAE08a] in which each interaction refers to an ordered pair of agents (responder, initiator). Each interaction

triggers an update of states in both agents according to some predefined deterministic *transition function*, where the update of relevant states is denoted by $A + B \rightarrow C + D$. We focus on two complexity measures including *space complexity* defined as the *number of states* utilised by each agent, and *time complexity* reflecting the total number of interactions required to stabilise the population protocol. We also consider *parallel time* defined as the total number of interactions divided by the size of the population. This time measure can be also seen as the local time observed by an agent, i.e., the number of pairwise interactions in which the agent is involved in. We aim at protocols formed of $\mathcal{O}(n \cdot \text{poly log } n)$ interactions equivalent to the parallel running time $\mathcal{O}(\text{poly log } n)$.

In order to maintain clarity of presentation each state has a name drawn from either a fixed size set of suitable names or a small range of integer values. However, when it is clear from the context we tend to omit the name of this field. Moreover, since each node belongs to exactly one of 3 sub-populations, for simplicity we shorten the notation omitting the part `role =` and writing for example $\langle \dots \rangle$ instead of $\langle \text{role} = \text{C}, \dots \rangle$. This notation allows us to refer only to the relevant fields, i.e., those affected during one particular type of interaction. One should keep in mind also that interactions may trigger several non-conflicting rules. For example, rules of transition of clocks happen in parallel to the rules of transition of coins.

Consider an event X , and let $\eta > 0$ be some predefined constant. We say that an event occurs with *negligible* probability, if there is an integer n_0 , s.t., the probability of this event for $n > n_0$ is at most $n^{-\eta}$. An event occurs *with high probability* (whp) if its probability is at least $1 - n^{-\eta}$ for $n > n_0$. If the event refers to a behaviour of an algorithm, we say it occurs with high probability if the constants used in the algorithm can be fine-tuned so that the probability of this event is at least $1 - n^{-\eta}$. Analogously an event X occurs *with very high probability* (vw hp) if for any $a > 0$ there exists an integer n_a such that event X occurs with probability at least $1 - n^{-a}$ when $n > n_a$. In particular, if an event occurs with probability $1 - n^{-\omega(1)}$, it occurs with very high probability.

3 Phase clock

The actions of our leader election protocol are coordinated by a phase clock utilising junta of clock leaders. A similar approach can be found in [GS18]. The junta leaders are drawn from sub-population *coins* denoted by C . With the help of the phase clock every agent in C keeps track of `phase` $\in \{0, 1, \dots, \Gamma - 1\}$, for a suitable large constant Γ , and maintains its `timemode` $\in \{\text{injunta}, \text{follower}\}$. Let $+_{\Gamma}$ denote addition modulo Γ and

$$\max_{\Gamma}(x, y) = \begin{cases} \max(x, y) & \text{if } |x - y| \leq \Gamma/2, \\ \min(x, y) & \text{if } |x - y| > \Gamma/2. \end{cases}$$

The transition rules of interaction with respect to the phase clock include:

$$\begin{aligned} \langle \text{follower}, \text{phase} = t_1 \rangle + \langle \text{phase} = t_2 \rangle &\rightarrow \\ &\langle \text{follower}, \text{phase} = T_1 \rangle + \langle \text{phase} = t_2 \rangle \\ \langle \text{injunta}, \text{phase} = t_1 \rangle + \langle \text{phase} = t_2 \rangle &\rightarrow \\ &\langle \text{injunta}, \text{phase} = T_2 \rangle + \langle \text{phase} = t_2 \rangle, \end{aligned}$$

where $T_1 = \max_{\Gamma}(t_1, t_2)$, $T_2 = \max_{\Gamma}(t_1, t_2 +_{\Gamma} 1)$. Agents are initialised to $\langle \text{follower}, \text{phase} = 0 \rangle$. During execution of coin preprocessing protocol, see Section 5, some agents in C become *junta members*. We say that the phase clock *passes through 0* whenever its current phase x is reduced in absolute terms. We denote this transition by $\xrightarrow{0}$.

Definition 3.1 (c.f., [GS18]). *Passes through 0 of agents a and b are equivalent if they both occur in a period when the respective agent's clock phases x_a and x_b satisfy $3\Gamma/4 <_{\Gamma} x_a, x_b <_{\Gamma} \Gamma/4$.*

Theorem 3.2 (c.f., Theorem 3.1 and Fact 3.1 in [GS18]). *For any constant $\varepsilon, \eta, d > 0$, there exists a constant Γ , s.t., if the number of junta members is at most $n^{1-\varepsilon}$ at any time whp $1 - n^{-\eta}$, the following conditions hold whp until each agent completes n^{η} passes through 0:*

- All passes through 0 form equivalence classes for all agents and the number of interactions between the closest passes through 0 in different equivalence classes is at least $d \cdot n \log n$.
- The number of interactions between two subsequent passes through 0 in any agent is $\mathcal{O}(n \log n)$.

A period between two agent's passes through zero is called a *round*. By Theorem 3.2, the rounds for different agents form equivalence classes whp that are referred to as *rounds of the protocol*. The *updated* agent is the one which acts as responder during the relevant interaction. Interactions with both start- and end-phase in $\{0, 1, \dots, \Gamma/2 - 1\}$ are denoted by $\xrightarrow{\text{early}}$ and those with start- and end-phase in $\{\Gamma/2, \dots, \Gamma - 1\}$ are denoted by $\xrightarrow{\text{late}}$. Finally, applying Theorem 3.2 and with Γ being twice as big as required by Theorem 3.2, we can guarantee that passes through 0 and through $\Gamma/2$ form strictly separate equivalence classes.

4 High level description

An execution of our algorithm consists of three consecutive epochs whp. These include the *initialisation* epoch, the *fast elimination* epoch and the *final elimination* epoch. For the case when any epoch fails, which happens with negligible probability, we use as backup the slow leader election protocol working in time $\mathcal{O}(n \log n)$ [AAE08a]. During the initialisation epoch the whole population is divided into sub-populations, where the descriptor $\text{role} \in \{\text{C}, \text{I}, \text{L}\}$ differentiates agents between the three sub-populations of *coins*, *inhibitors* and *leaders* respectively. At the start of the protocol all agents are subject to symmetry breaking rules. Each agent gets assigned to one of the three roles (or gets deactivated), and this role is never changed. The two symmetry breaking rules adopted during the initial partition process are as follows:

$$0 + 0 \rightarrow X + \text{L}, \quad X + X \rightarrow \text{C} + \text{I}, \quad (1)$$

where 0 describes an agent before initialisations, and X refers to an intermediate stage before entering sub-population C or I.

During the initialisation epoch a *junta* of size at most $n^{0.77}$ is elected from C whp, which allows to start the phase clock using this junta as clock leaders. This phase clock synchronises all actions of our algorithm until it concludes. In our approach it is important to terminate the initialisation epoch and in turn to stabilise the roles of agents in time $\mathcal{O}(\log n)$. With this in mind we adopt two extra rules, s.t., whenever a node in state 0 or X reaches the end of the first round, it deactivates itself:

$$0 + \star \xrightarrow{0} \text{D} + \star, \quad X + \star \xrightarrow{0} \text{D} + \star, \quad (2)$$

where D denotes *deactivated* agents that, except for passing clock state, do not play any meaningful role in the leader election protocol.

Lemma 4.1. *With high probability, only $\mathcal{O}(n/\log n)$ agents are not initialised in the course of the protocol, i.e., $n - \mathcal{O}(n/\log n)$ agents join C, I or L during the first $\mathcal{O}(n \log n)$ interactions.*

Due to space limitation the proof of Lemma 4.1 is only available in Appendix A, where one can also find the other missing proofs. Using Lemma 4.1 one can immediately conclude that during the first round all $\mathcal{O}(n/\log n)$ not initialised agents, i.e., those not given roles C, L or I, become deactivated with high probability by rule (2). Below we explain functionality of the three adopted sub-populations.

Coin Agents in this group differentiate themselves into non-empty levels $0, 1, 2, \dots, \Phi$, where $\Phi = \log \log n - 3$. The number of agents on level Φ is at most $n^{0.77}$ and these agents form the junta running the phase clock. The levels are also used to simulate $\Phi + 1$ types of asymmetric coins, s.t., if the probability of drawing heads by ℓ -th coin is q the probability of drawing heads by $(\ell + 1)$ -st coin is roughly q^2 . In terms of implementation, tossing ℓ -th asymmetric coin is realised by an agent interacting with another agent as a responder. And the outcome is *heads* if the initiator is a coin on level ℓ or higher.

Leader Agents in this group are leader candidates, i.e., each agent in this group has a chance to become the unique leader. In due course the number of candidates is reduced to one. The main challenge is in fast but also safe candidate elimination, i.e., we need to guarantee that our protocol does not eliminate all candidates.

Inhibitor Agents in this group are split into $\Psi = \Theta(\log \log n)$ distinct subgroups, with the expected cardinality $\Theta(n/2^i)$ of each subgroup i . I.e., the sizes of subgroups span from $\Theta(n)$ to $\Theta(n/\log^c n)$, for some constant $c > 0$. We target with this system of subgroups specific points in time, where the i th subgroup is responsible for phase clock round $2^{\Theta(i)}$ in expectation. This system of groups is used to guide through the final elimination process when we safely reduce the number of leaders from $\mathcal{O}(\log n)$ to a single one.

The first (initialisation) epoch generates at least one leader candidate and with high probability the number of candidates is almost $n/2$. The protocol will eventually elect a single leader among all leader candidates in L during the second and the third epoch. The second epoch related to fast elimination reduces the number of *active* (not withdrawn yet) leader candidates to $\mathcal{O}(\log n)$ agents in time $\mathcal{O}(\log n \log \log n)$ with high probability. The fast elimination uses the sub-population C to simulate assorted biased coins.

The third epoch eliminates all but one competitor which becomes the unique leader, and is successful in achieving this goal whp. In the third epoch a standard leader elimination by fair coin tosses is performed. Eliminated leaders transition from active into a passive state which is still considered the active leader state. During this process subsequent subsets $1, 2, 3, \dots, \Psi$ of subpopulation l perform computations of a slowing down clock, whose i -th tick is in phase clock round $2^{\Theta(i)}$. This tick requires usage of active leaders so i -th tick is not possible if there are none of them. Because of this such a tick is used to transition passive leaders into followers, since it assures them about the existence of at least one active leader. This process guarantees that at least one leader candidate survives while it still works fast in expectation. The third epoch elects a single leader in $\mathcal{O}(\log n \log \log n)$ expected time and in $\mathcal{O}(\log^2 n)$ time with high probability.

Independently, we simultaneously perform actions of the slow constant-space leader election protocol, in which if two leader candidates interact, exactly one of them gets eliminated. This slow protocol causes a slow depletion of leader candidates, which does not have a noticeable effect on the second and the third epochs of the protocol. This depletion assures election of one leader in cases when either the phase clock gets desynchronised or all leader candidates become passive (marked for elimination) during the last two epochs.

The leader candidates elimination process during the second and the third epoch works as follows. The protocol operates in consecutive rounds, each taking time $\mathcal{O}(\log n)$. For each agent a round is defined as the time between two subsequent passes of the phase clock through value zero. In the first half of each round still active leader candidates flip a coin to decide whether they intend to survive (*heads*) this round or not (*tails*). If any *heads* are drawn during this round, the relevant information is distributed (via one-way epidemic [AAE08a]) to all agents during the second half of the round. This results in elimination of all active candidates which drew *tails*. However, if no *heads* are drawn the round is considered *void*.

In the fast elimination process we utilise asymmetric coins implemented through interactions with agents in diverse population of C . The first asymmetric coin Φ is used 4 times to reduce the population of active leader candidates to size at most $n/n^{0.77}$. Further we use each of asymmetric coins $\Phi - 1, \Phi - 2, \dots, 1$ exactly twice. Using a biased coin with *heads* coming with probability q guarantees whp reduction of active leader candidates by a factor close to q . On the conclusion of this process (all coins are used) the number of active leader candidates is down to $\mathcal{O}(\log n)$ whp.

In contrast, in the third epoch symmetric almost fair coins are used in the elimination process indefinitely. This results in elimination of all but a single leader candidate in $\mathcal{O}(\log \log n)$ rounds in expectation. In order to guarantee that the protocol is *always* correct, i.e., the last alive leader candidate is never eliminated, we use the support of agents in sub-population l .

5 Coins

Let $\Phi = \lfloor \log \log n \rfloor - 3$. The states of coins, i.e., agents belonging to sub-population C store the following information: $\text{level} \in \{0, 1, 2, \dots, \Phi\}$, reflecting the level of asymmetry, and $\text{mode} \in \{\text{adv}, \text{stop}\}$ indicating whether a coin is still willing to increment its level. We also need an extra constant space to store the current state of the phase clock. When formed after application of split rule (1) each coin is initialized to $C\langle \text{level} = 0, \text{adv} \rangle$.

Coin preprocessing In what follows we introduce the rules governing level incrementation. Note that these closely resemble the rules from *forming junta* protocol proposed in [GS18].

$$\begin{aligned}
& C\langle \text{level} = x, \text{adv} \rangle + Y \rightarrow \\
& C\langle \text{level} = x, \text{stop} \rangle + Y, \text{ for } Y \neq C, \\
& C\langle \text{level} = x, \text{adv} \rangle + C\langle \text{level} = y \rangle \rightarrow \\
& C\langle \text{level} = x, \text{stop} \rangle + C\langle \text{level} = y \rangle, \text{ for } x > y, \\
& C\langle \text{level} = x, \text{adv} \rangle + C\langle \text{level} = y \rangle \rightarrow
\end{aligned}$$

$$C\langle \text{level} = x + 1, \text{adv} \rangle + C\langle \text{level} = y \rangle, \text{ for } x \leq y, x < \Phi.$$

Once the level of a coin in C reaches Φ it stops growing. Moreover, we give name *injunta* to all coins which managed to reach $\text{level} = \Phi$. In order to characterise properties of coins we formulate a series of lemmas. Let n_C be the total number of coins. By Lemma 4.1 and rules (1) and (2), $n_C = \frac{n}{4} - \mathcal{O}(n/\log n)$ with very high probability. Let C_ℓ be the number of coins which reach level ℓ or higher. The value of C_ℓ depends on the execution thread of the protocol. We first observe that $n_C = C_0$, and further estimates on C_ℓ , for $\ell > 0$ are determined by Lemmas 5.1 (upper bound) and 5.2 (lower bound).

Lemma 5.1 (Lemma 4.2. in [GS18]). *Assume $n^{-1/3} \leq q < 1$ and $C_\ell = q \cdot n$, then $C_{\ell+1} \leq \frac{11}{10}q^2 \cdot n$ with very high probability.*

The lower bound argument (similar to the proof of Lemma 5.1) is given below.

Lemma 5.2. *Assume $n^{-1/3} \leq q < 1$ and $C_\ell = q \cdot n$, then $C_{\ell+1} \geq \frac{9}{20}q^2 \cdot n$ whp.*

Proof. Each coin contributing to value C_ℓ arrives at level ℓ during some interaction t . These coins arrive sequentially. Consider $(i+1)$ -st coin v that got to level ℓ . At the time the coin arrives there are already i coins on levels $\ell' \geq \ell$. Consider the first interaction τ succeeding t in which coin v acts as the responder. During this interaction the initiator is a coin on level $\ell' \geq \ell$ with probability $p_\tau \geq i/n$. Thus v moves to level $\ell+1$ with probability at least i/n as otherwise the responder would end up in state $(\ell, 0)$ and would not contribute to $C_{\ell+1}$. Consider now the sequence of C_ℓ such interactions τ , in which each of C_ℓ coins act as responder after getting to level ℓ . We can attribute to this sequence a binary 0-1 sequence σ of length C_ℓ , s.t., if during interaction τ a coin ends up in state $(\ell, 0)$, the respective entry in σ becomes 0, and otherwise this entry becomes 1 (this happens with probability at least p_τ). The expected number of these 1s is at least $\sum_i i/n = (C_\ell - 1)C_\ell/2n = (q^2 \cdot n - q)/2$. And by Chernoff bound $C_{\ell+1} < \frac{9}{20}q^2 \cdot n$ with very high probability. \square

Lemma 5.3. *For n large enough and $\Phi = \lfloor \log \log n \rfloor - 3$ we have $n^{0.45} \leq C_\Phi \leq n^{0.77}$ whp.*

Proof. We start with $9n/40 \leq n_C = C_0 \leq n/4$ with very high probability. By Lemma 5.1 and Lemma 5.2 iterated ℓ times we conclude that with very high probability

$$(9/20)^{2^{\ell+1}-1} \cdot \frac{n}{2^{2^{\ell+1}}} \leq C_\ell \leq (11/10)^{2^\ell-1} \cdot \frac{n}{2^{2^{\ell+2}}}.$$

Note that if we adopt $\Phi = \lfloor \log \log n \rfloor - 3$, we get

$$\begin{aligned} C_\Phi &\geq (9/20)^{2^{\Phi+1}} \cdot \frac{n}{2^{2^{\Phi+1}}} \geq n \cdot (9/40)^{2^{\log \log n - 2}} \geq \\ &\frac{n}{2^{2.2 \cdot 2^{\log \log n / 4}}} \geq n/n^{0.55} = n^{0.45}. \end{aligned}$$

On the other hand

$$\begin{aligned} C_\Phi &\leq (11/10)^{2^\Phi} \cdot \frac{n}{2^{2^{\Phi+2}}} \leq n \cdot (11/160)^{2^{\log \log n - 4}} \leq \\ &\frac{n}{2^{3.8 \cdot 2^{\log \log n / 16}}} \leq n/n^{0.23} = n^{0.77}. \end{aligned} \quad \square$$

Lemma 5.4 (Analogue of Lemma 4.5. in [GS18]). *The bounds from Lemma 5.1, Lemma 5.2 and Lemma 5.3 hold after $\mathcal{O}(n \log n)$ interactions.*

Proof. (Sketch) In coin preprocessing protocol we need to stabilise first sub-population of coins in time $\mathcal{O}(\log n)$. The time complexity analysis of the remaining part of the protocol is analogous to the one used in forming junta protocol in [GS18]. \square

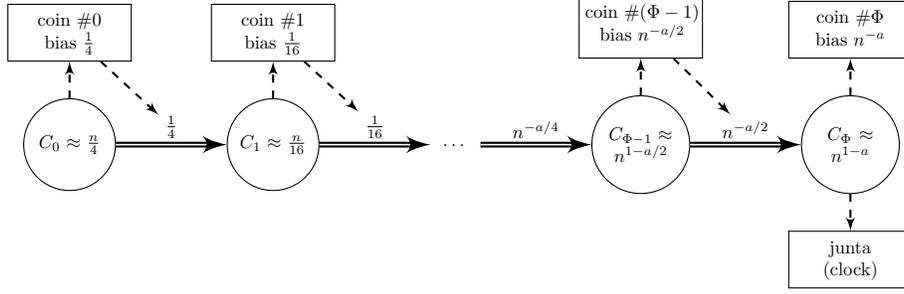


Figure 1: An idealized scheme of coin sub-populations and their relation to biased coins. In the picture $0.23 \leq a \leq 0.55$. Solid lines denote evolution of the population and dashed lines refer to the relevant functionality.

6 Fast elimination

The goal in fast elimination epoch is to reduce the number of active leader candidates to $\mathcal{O}(\log n)$ whp. We also guarantee that at least one agent remains in the group of active leaders **A** whp. All other leader candidates join group **P** of passive agents.

The state of each leader candidate in this epoch consists of: $\text{cnt} \in \{0, 1, \dots, 2\Phi + 3\}$, $\text{leadermode} \in \{\mathbf{A}, \mathbf{P}, \mathbf{W}\}$ (in fast elimination **W** standing for withdrawn is not used), $\text{flip} \in \{\text{none}, \text{heads}, \text{tails}\}$, $\text{void} \in \{\text{true}, \text{false}\}$ (telling whether the round is void), and a constant number of phase clock values. Each leader candidate is initialised at the beginning of the first round of the second epoch to $\mathbf{L}\langle \text{cnt} = 2\Phi + 3, \mathbf{A}, \text{none}, \text{void} = \text{true} \rangle$.

After the first round of the phase clock, when the roles of all agents are fixed and levels of all coins are computed whp, agents enter the fast elimination epoch. This is ensured by starting the counter at one larger than the intended number of coin uses. At the beginning of the fast elimination all leader candidates are active (**A**). In fast elimination we use the sub-population **C** of coins as the source of Φ different types of asymmetric coins. The coin result is generated, when a leader candidate interacts with another agent acting as the responder. The outcome of using ℓ -th biased coin is *heads* when the interaction refers to a coin on level at least ℓ , and *tails* otherwise. When $C_{\ell} = q \cdot n$, the probability of drawing *heads* at this level is q . Thus when there are substantially more than $1/q$ active leader candidates almost certainly at least one of them has to draw *heads*. In turn the number of active leader candidates will be reduced by factor of $1/q$ in expectation. On the other hand, if the number of active leader candidates does not exceed $1/q$, no agent may draw *heads*. In order to have good understanding of the situation the agents with *heads* drawn inform others (using one-way epidemic) about this fact. Thus if an agent draws *tails* and receives a message about other agent(s) having *heads*, it can safely become passive (**P**). This elimination cycle can be carried in one round in time $\mathcal{O}(\log n)$.

During fast elimination active leader candidates utilise coins, s.t., each coin $1, 2, \dots, \Phi - 2, \Phi - 1$ is used exactly twice and coin Φ is applied four times. In other words, the elimination process can be represented by a sequence $(\gamma)_1^{2\Phi+2} = [1, 1, 2, 2, \dots, \Phi - 1, \Phi - 1, \Phi, \Phi, \Phi, \Phi]$ which tells us which coin level is used with what cnt value. In total, the elimination process operates in $\mathcal{O}(\log \log n)$ rounds translating to parallel time $\mathcal{O}(\log n \log \log n)$. We are also able to guarantee reduction of the number of remaining active leader candidates to $\mathcal{O}(\log n)$ whp.

The following transitions are used in the second epoch. When the phase clock passes through zero we have

$$\begin{aligned} \mathbf{L}\langle \text{cnt} = x \rangle + \star &\xrightarrow{0} \\ \mathbf{L}\langle \text{cnt} = x - 1, \text{none}, \text{void} = \text{true} \rangle + \star, &\text{ for } x \geq 1. \end{aligned} \quad (3)$$

When $x = 1$, at the end of the round we move to the third epoch. Otherwise, in the first half of the round application of the coin from the current level $\gamma(x)$ is guaranteed whp, for all active leader

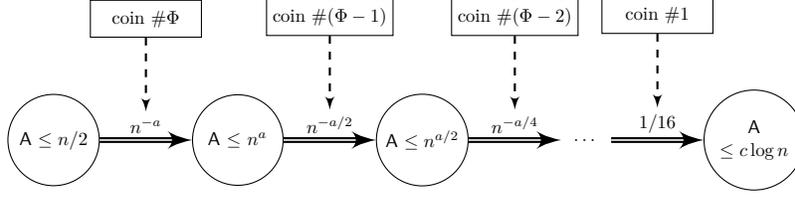


Figure 2: An idealised scheme of the fast elimination process.

candidates. For $x \neq 2\Phi + 3$:

$$\begin{aligned} &L\langle A, \text{cnt} = x, \text{none} \rangle + C\langle \text{level} = y \rangle \xrightarrow{\text{early}} \\ &L\langle A, \text{cnt} = x, \text{heads}, \text{void} = \text{false} \rangle + C\langle \text{level} = y \rangle, \end{aligned} \quad (4)$$

$$\begin{aligned} &L\langle A, \text{cnt} = x, \text{none} \rangle + C\langle \text{level} = y \rangle \xrightarrow{\text{early}} \\ &L\langle A, \text{cnt} = x, \text{tails} \rangle + C\langle \text{level} = y \rangle, \\ &L\langle A, \text{cnt} = x, \text{none} \rangle + Y \xrightarrow{\text{early}} \\ &L\langle A, \text{cnt} = x, \text{tails} \rangle + Y, \end{aligned} \quad (5)$$

when $\gamma(x) \leq y$ and $\gamma(x) > y, Y \neq C$ respectively.

In the second half of the round broadcast (via one-way epidemic) informing about drawn heads works as follows

$$\begin{aligned} &L\langle A, \text{tails}, \text{void} = \text{true} \rangle + L\langle \text{void} = \text{false} \rangle \xrightarrow{\text{late}} \\ &L\langle P, \text{tails}, \text{void} = \text{false} \rangle + L\langle \text{void} = \text{false} \rangle, \end{aligned} \quad (6)$$

$$\begin{aligned} &L\langle \text{void} = \text{true} \rangle + L\langle \text{void} = \text{false} \rangle \xrightarrow{\text{late}} \\ &L\langle \text{void} = \text{false} \rangle + L\langle \text{void} = \text{false} \rangle. \end{aligned} \quad (7)$$

The following lemmas guarantee the correctness of the second epoch whp.

Lemma 6.1. *There exists a constant $c > 0$, s.t., for any $q < 1$ when $N \geq c \log n/q$ agents toss an asymmetric coin resulting in heads with probability q , the following holds:*

1. *none of the agents draws heads with a negligible probability, and*
2. *more than $2q \cdot N$ agents draw heads with a negligible probability.*

Proof. The probability that all agents draw *tails* is at most $(1-q)^{c \log n/q} \leq e^{-c \log n} = n^{-c}$. The expected number of agents which draw *heads* is $q \cdot N \geq c \log n$. By Chernoff bound the probability that more than $2q \cdot N$ agents draw *heads* is smaller than $e^{-Nq} \leq n^{-c}$. \square

Lemma 6.2. *Applying coin (from level) Φ four times and then coins $\Phi - 1, \Phi - 2, \dots, \ell + 1, \ell$ twice reduces the number of active leader candidates to at most $c \log n/q$, where q is the probability of tossing heads by coin $\ell \geq 1$.*

Proof. Induction is on ℓ . In the base case when $\ell = \Phi$ we apply coin Φ four times. By Lemma 5.3 we have $q \geq n^{-0.23}$. So applying this coin four times gives reduction of the number of active leader candidates to at most $\max\{16n \cdot n^{-4 \cdot 0.23}, c \log n/q\} = c \log n/q$ whp.

Now assume the thesis holds for level $\ell + 1$ and we prove it for ℓ . By inductive hypothesis and Lemma 5.2, after application of coin $\ell + 1$ twice there are at most $c \log n/q' \leq 20c \log n/9q^2$ active leader candidates (q' is the counterpart of q at level $\ell + 1$). Applying this coin twice gives further reduction of active leader candidates to at most $\max\{80c \log n/9, c \log n/q\} = c \log n/q$ whp. \square

7 Final elimination

The protocol executes $\Theta(\log \log n)$ rounds of fast elimination (applying coins from level Φ down to 1) concluding with $\mathcal{O}(\log n)$ active leaders left whp. All other leader candidates become passive. The remaining task consists in electing a single leader out of the remaining active $\mathcal{O}(\log n)$ candidates. In the final configuration there must be exactly one leader and $n - 1$ followers and this situation should be maintained forever (although the configuration can evolve indefinitely where the unique leader alternates between allowed leader states and followers move between dedicated follower states). The main idea behind the solution is to iterate the following process. Each of the remaining candidates picks some value at random, agents compute the maximum of these values in time $\mathcal{O}(\log n)$, and only the owners of the largest value remain leader candidates. There can be at most $\mathcal{O}(\log \log n)$ values that one could pick from as we do not allow agents to operate on more states. In fact, we limit this choice to set $\{0, 1\}$ as selecting random bits is easier, and drawing from a larger range does not necessarily provide us with a substantial gain in terms of time complexity. Instead, we concentrate on the expected running time of unique leader election by observing that the number of bits to be drawn from $\{0, 1\}$ (to conclude this process) is expected to be $\mathcal{O}(\log \log n)$. This translates to the expected running time $\mathcal{O}(\log \log n \log n)$. Note that to obtain a unique leader whp each agent has to pick $\mathcal{O}(\log n)$ bits, which translates to time $\mathcal{O}(\log^2 n)$. This is why, instead, we focus on construction of a Las Vegas type algorithm which always elects exactly one leader in faster expected time $\mathcal{O}(\log \log n \log n)$.

The most challenging problem to overcome during each iteration of leader elimination (when the next value from set $\{0, 1\}$ is drawn and information about drawn 1's in spread) is to prevent all remaining leader candidates from becoming followers, which could happen due to phase clock desynchronisation with a negligible probability. In fact a similar setback can also happen in the fast elimination epoch. If elimination was equivalent to becoming a follower, we could accidentally cull all leaders, which is not permissible in a Las Vegas algorithm. In order to prevent this from happening we use the leader modes A (active) for not eliminated leaders, P (passive) for eliminated leaders and we also introduce the W (withdrawn) mode whose holders are followers that were initially in leader subpopulation. The active and passive candidates may still become the unique leader, and we use a joint term *alive* candidates for these two groups.

In order to pick the value for an active leader candidate we utilise coin (from level) 0. All alive candidates keep track of a counter **drag** which is increasing during the final elimination epoch. This counter is ticking at rate which is slowing down in time. The time elapsing between its i -th and $(i + 1)$ -st incrementation is on average $\Theta(4^i \log n)$. Only active candidates can increment this counter. If a passive candidate detects an increment of the counter value (wrt to its own), it transitions into withdrawn state. This process is now safe because withdrawing agents have enough evidence that there are still active candidates with a higher **drag** value. In other words, we have a strong guarantee that all alive candidates do not transition into withdrawn state. This is very important should the phase clock get unexpectedly desynchronised. And indeed, the first increment of **drag** counter makes all candidates that became passive during the fast elimination withdrawn. The counter **drag** operates during the first $\Theta(n \log^2 n)$ interactions, when one leader is selected whp, so it has only $\Theta(\log \log n)$ states. In slowing down this counter we rely on inhibitor agents (I) preprocessed at the same time as coins (C). The counter **drag** assures that if only one active candidate remains in iteration T , then all other candidates become followers before iteration $\mathcal{O}(T)$ whp. Use of counter **drag** is a new technique that makes it possible to achieve expected stabilisation time $\mathcal{O}(\log \log n \log n)$ by withdrawing all passive candidates soon after a single active leader remains whp. The methods utilised in paper [GS18] do not guarantee this effect and thus protocols from the latter have the expected stabilisation time $\mathcal{O}(\log^2 n)$.

Preprocessing begins with the first pass through 0 of the phase clock. Inhibitor agents keep track of $\text{drag} \in \{0, 1, \dots, \Psi\}$, $\text{mode} \in \{\text{adv}, \text{stop}\}$ (flag whether agent is *advancing* or *stopped*) and elevation $\text{elev} \in \{\text{low}, \text{high}\}$. The agents are initialized to $I(\text{drag} = 0, \text{adv}, \text{low})$, and **drag** counts how many subsequent successful coin flips they managed to obtain.

$$\begin{aligned}
 I(\text{drag} = x, \text{adv}) + Y &\xrightarrow{\text{late}} I(\text{drag} = x + 1, \text{adv}) + Y, \\
 &\text{for } Y \neq C, \\
 I(\text{drag} = x, \text{adv}) + C &\xrightarrow{\text{late}} I(\text{drag} = x, \text{stop}) + C.
 \end{aligned}$$

We denote by n_I the total number of inhibitor agents in I. By Lemma 4.1, $n_I = n/4 - \mathcal{O}(n/\log n)$.

Let D_ℓ be the number of agents that reach **drag** ℓ .

Lemma 7.1. *After the first round of the clock $D_\ell = n4^{-\ell}(1 \pm o(1))$ whp.*

Slowed-down inhibitor communication: Inhibitor agents get activated through interaction with the leader agents which reached the appropriate **drag** value, and this communication is done via one-way epidemic (between inhibitors of the same **drag**), i.e.,

$$\begin{aligned} I\langle \text{drag} = x, \text{stop}, \text{low} \rangle + L\langle A, \text{drag} = x \rangle &\rightarrow \\ I\langle \text{drag} = x, \text{stop}, \text{high} \rangle + L\langle A, \text{drag} = x \rangle, & \quad (8) \\ I\langle \text{drag} = x, \star \rangle + I\langle \text{drag} = x, \text{high} \rangle &\rightarrow \\ I\langle \text{drag} = x, \text{high} \rangle + I\langle \text{drag} = x, \text{high} \rangle. & \end{aligned}$$

Safe withdrawal: All active leader candidates with **drag** > 0 are subject to coin-flipping rules (4) and (5). More precisely, in the first half of the round each of them draws the coin from level 0 whp. As rules (6) and (7) apply to agents with coin-flips resulting in success inform (via one-way epidemic) other agents accordingly.

We give below an updated reset rule (analogue of (3)) observing that this rule does not change the **drag** value, as well as updated rules for leaders:

$$\begin{aligned} L\langle \star, \text{void} = \star \rangle + \star &\xrightarrow{0} L\langle \text{none}, \text{void} = \text{true} \rangle + \star, \\ L\langle \star, \text{drag} = x \rangle + L\langle \text{drag} = y \rangle &\rightarrow \\ L\langle W, \text{drag} = y \rangle + L\langle \text{drag} = y \rangle, & \quad \text{for } x < y, \quad (9) \end{aligned}$$

$$\begin{aligned} L\langle A, \text{heads}, \text{drag} = x \rangle + I\langle \text{drag} = x, \text{high} \rangle &\rightarrow \\ L\langle A, \text{heads}, \text{drag} = x + 1 \rangle + I\langle \text{drag} = x, \text{high} \rangle. & \quad (10) \end{aligned}$$

Let $A \leq c \log n$ be the number of active leaders with **drag** $= \ell$. Let T_ℓ be a random variable denoting the number of interactions between the first occurrence of an active leader candidate with **drag** $= \ell$ and the first occurrence of an active leader candidate with **drag** $= \ell + 1$.

Lemma 7.2. *There exist constants $c_1, c_2 > 0$ such that for $\ell \leq \Psi$ we have $\Pr[T_\ell \leq c_1 4^\ell n \log n] \leq n^{-0.5}$ and whp $T_\ell \leq c_2 4^\ell n \log n$.*

Proof. Consider the first interaction t in which a leader candidate assumes **drag** $= \ell$. We are to prove inequalities on the number of interactions T_ℓ till the first interaction t' in which a leader candidate assumes **drag** $= \ell + 1$.

We start with the first inequality, which is in fact a lower bound on T_ℓ . When the first leader with **drag** $= \ell$ occurs it starts propagation (via one-way epidemic) of state **high** amongst inhibitors with **drag** $= \ell$. In the context of the lower bound argument, we can consider a situation in which $A \leq c \log n$ informed agents spread rumour to D_ℓ uninformed agents in the population. By Lemma 7.1 $D_\ell = \Theta(n/4^\ell)$. We observe the following. If the informed part of population is of size x , then a single interaction increments this size with probability approximately $\frac{D_\ell}{n} \cdot \frac{x}{n} = 4^{-\ell} \cdot \frac{x}{n} (1 \pm o(1))$. Thus, when $A < x < 1/2 \cdot D_\ell$ it takes $\Theta(n \cdot 4^\ell)$ interactions to go from x informed agents to $2x$, with high probability (which follows from Chernoff bound). So it takes $\Theta(4^\ell n \log n)$ interactions, i.e., more than $c_1 4^\ell n \log n$ for some $c_1 > 0$, to reach the sub-population of inhibitors that are **high** with **drag** $= \ell$ of cardinality $n^{0.4}$. During this time the probability of having an interaction incrementing value **drag** to $\ell + 1$ is $\Theta(4^\ell n^{-0.6} \log^2 n)$. For n large enough this probability is smaller than $n^{-0.5}$.

With respect to the upper bound, consider the same communication process. Observe that if a single inhibitor with **drag** $= \ell$ gets **high**, all inhibitors with **drag** $= \ell$ get **high** in $\mathcal{O}(4^\ell n \log n)$ subsequent interactions, with high probability. In further $\mathcal{O}(4^\ell n \log n)$ interactions some active leader with **drag** $= \ell$ interacts with one of these inhibitors whp. Thus there is a constant $c_2 > 0$ such that $T_\ell \leq c_2 4^\ell n \log n$ whp. \square

We now bound the time the protocol needs to elect a single leader whp. Recall that at the beginning of the last epoch, there are at most $c \log n$ active leaders, for a constant $c > 0$.

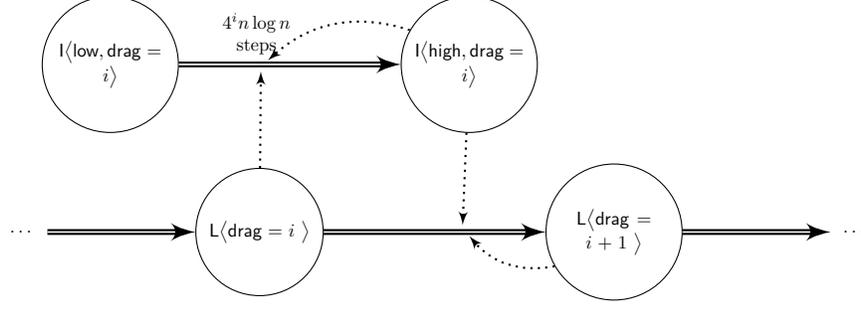


Figure 3: The implementation of slowing down drag counter, where dotted arrows indicate enabled transitions.

Lemma 7.3. *Assume that the preprocessing, the phase clock and coin propagations work properly. After $\mathcal{O}(\log \log n)$ rounds in expectation and $\mathcal{O}(\log n)$ rounds with high probability the number of active leaders is reduced from $c \log n$ to 1.*

Proof. Let sequence F_0, F_1, \dots, F_i contain the number of active leaders after i rounds of elimination. Let B be the number of rounds needed to obtain a single active leaders, that is $B = \min\{i : F_i = 1\}$. Let F'_i be as follows: for $i \leq B$, $F'_i = F_i$ and otherwise $F'_i = (5/6)^{i-B}$. Let $p = (1 - \mathcal{O}(1/\log n))/4$ be the probability of drawing heads and $q = 1 - p$.

Let A_i be the event that all leader candidates drew tails which happens with probability q^{F_i} . When A_i occurs $F_{i+1} = F_i$. Otherwise, with probability $1 - q^{F_i}$, the value of F_{i+1} corresponds to the number of successes during consecutive F_i coin-flips. Thus, $\mathbb{E}[F_{i+1} | F_i] = F_i(p + q^{F_i})$. Assume that $F_i \geq 2$. Then $\mathbb{E}[F_{i+1} | F_i] \leq F_i \cdot (13/16 + o(1)) \leq F_i \cdot 5/6$ for large enough n . Thus, by the definition of F'_i , we have $\mathbb{E}[F'_{i+1} | F'_i] \leq F'_i \cdot 5/6$. In turn $\Pr[B > i] = \Pr[F_i > 1] = \Pr[F'_i > 1] < \mathbb{E}[F'_i] = F_0 \cdot (5/6)^i$. Since $F_0 \leq c \log n$, we get

$$\Pr[B > \log_{6/5}(c \log n \cdot n^\eta)] < n^{-\eta}, \text{ and}$$

$$\mathbb{E}[B] = \sum_{i=0}^{\infty} \Pr[B > i] \leq \sum_{i=0}^{\infty} \min(1, F_0 \cdot (5/6)^i) =$$

$$\mathcal{O}(\log F_0) + \mathcal{O}(1). \quad \square$$

We now compute the time needed to elect a single leader, i.e., to preserve a single agent in state A and to change states of all agents in state P to W.

Lemma 7.4. *Assume that the preprocessing, the phase clock and coin propagations work properly and that exactly one active leader enters $\mathbf{drag} = \Psi$. After time $\mathcal{O}(\log n \log \log n)$ in expectation and time $\mathcal{O}(\log^2 n)$ with high probability there is exactly one leader remaining in state A and all other candidates are moved to W.*

Proof. Let T be the number of rounds it takes to go from $c \log n$ to 1 of agents in state A. By Lemma 7.3 T is $\mathcal{O}(n \log n \log \log n)$ in expectation and $\mathcal{O}(n \log^2 n)$ with high probability. It is enough to consider any round $T' > T$, s.t., between T and T' all leader agents increase their \mathbf{drag} , since such interaction moves all P to W. Let x be the highest \mathbf{drag} value achieved by a P candidate. This candidate moves to state W as soon as it encounters a higher value of \mathbf{drag} in another candidate. In the proof we use value T_x defined just before Lemma 7.2.

Note that by Lemma 7.3 \mathbf{drag} value x is smaller than Ψ . By Lemma 7.2 we also have $\sum_{y=1}^{\Psi} T_y = \mathcal{O}(n \log^2 n)$ whp. Finally, the value of Ψ is propagated amongst leaders in $\mathcal{O}(n \log n)$ interactions whp, which completes the proof of the time bound obtained whp.

Let $T_A = \Theta(n \log n \log \log n)$ be the number of interactions of first two epochs. In order to obtain the improved time bound in expectation we observe that by Lemma 7.2 there exists an integer constant k such that for any y : $T_y + T_{y+1} + \dots + T_{y+k} \geq c_1 4^y n \log n$ whp, where c_1 is the constant defined in Lemma 7.2. Because of this $T \geq T_A + T_1 + \dots + T_x = \Omega(n \log n (\log \log n + 4^x))$ whp. Note that $T_{x+1} \leq c_2 4^x n \log n$ whp, and the time of propagation of value $x + 1$ amongst leaders is $\mathcal{O}(n \log n)$. Thus $T' = \mathcal{O}(n \log n (\log \log n + 4^x))$. Finally, whp the extra time needed to withdraw all passive agents increases

the total number of interactions at most by a constant factor. With remaining negligible probability the expected value of T' is at most the average number of interactions to get all leader candidates to $\text{drag} = \Psi$ which is $\mathcal{O}(n \log^2 n)$. \square

8 Slow backup protocol

We have shown earlier that with high probability our protocol behaves according to the scheme described in previous sections. In this case the protocol elects a single leader in expected time $\mathcal{O}(\log n \log \log n)$ and with high probability in time $\mathcal{O}(\log^2 n)$. However, we still need to provide a full guarantee that new $\mathcal{O}(\log n \log \log n)$ -time protocol elects a unique leader, even if the phase clock gets desynchronised at some point.

And indeed, the successful conclusion is guaranteed by running simultaneously (in the background) the slow elimination protocol from [AAD⁺04] in which an encounter of two leader candidates results in transition of exactly one of them to a withdrawn (follower) state. This must be done without disrupting neither of the three epochs of the fast leader election protocol presented above. We achieve this by introducing a *seniority order* on alive leader candidates to break ties during direct encounters when always more senior leaders survive. More precisely, we say that $L\langle A \rangle$ and $L\langle P \rangle$ agents are mapped to the leader in the output, and $L\langle W \rangle$, C, I, X, D and 0 states are mapped to non-leaders. We also adopt an extra interaction rule for when two agents $A, B \in \{L\langle A \rangle, L\langle P \rangle\}$ meet, B changes its state to $L\langle W \rangle$ when A is more senior to B , i.e.,

$$A + B \rightarrow A + L\langle W \rangle. \quad (11)$$

In addition, the seniority order gives preference to agents with higher drag , and if tied $L\langle A \rangle$ beats $L\langle P \rangle$. Finally, the agent with a smaller level wins, and heads wins with none and tails.

We first observe that with high probability rule (11) may only speed up the elimination process analysed in Sections 6 and 7, since it reduces the number of $L\langle A \rangle$ agents, and whp this rule never eliminates during one round *all* agents with heads.

Lemma 8.1. *Throughout the execution of leader election protocol there is always at least one agent in state $L\langle A \rangle$ or $L\langle P \rangle$.*

Proof. The leader candidates in state $L\langle A \rangle$ are formed by application of rule (1). Only rules (9) and (11) can change states of agents from $L\langle A \rangle$ and $L\langle P \rangle$ to $L\langle W \rangle$. However, neither of these rules can eliminate the last agent of this type which possesses the highest value of drag . \square

We now arrive in the main result of this paper.

Theorem 8.2 (Main result). *The leader election protocol presented in this paper always elects a unique leader. The protocol concludes in $\mathcal{O}(\log n \log \log n)$ expected time, and time $\mathcal{O}(\log^2 n)$ with high probability.*

Proof. By Lemma 5.4 in $\mathcal{O}(\log n)$ rounds whp we elect a junta of the appropriate size as indicated by Lemma 5.3. This junta starts the phase clock. By Lemma 6.2, fast elimination epoch leaves $\mathcal{O}(\log n)$ active leaders in $\mathcal{O}(\log n \cdot \log \log n)$ parallel time, with high probability. By Lemma 7.4, slow elimination epoch leaves a single leader in expected parallel time $\mathcal{O}(\log n \cdot \log \log n)$ and in parallel time $\mathcal{O}(\log^2 n)$ with high probability. In addition, by Lemma 8.1 we never eliminate all leader candidates, and rule (11) guarantees that whp in $\mathcal{O}(n)$ rounds in expectation and in $\mathcal{O}(n \log n)$ whp a single leader is chosen, which does not affect the overall running time. \square

References

- [AAD⁺04] D. Angluin, J. Aspnes, Z. Diamadi, M.J. Fischer, and R. Peralta. Computation in networks of passively mobile finite-state sensors. In *PODC*, pages 290–299, 2004.
- [AAE08a] D. Angluin, J. Aspnes, and D. Eisenstat. A simple population protocol for fast robust approximate majority. *Distributed Computing*, 21(2):87–102, 2008.
- [AAE08b] Dana Angluin, James Aspnes, and David Eisenstat. Fast computation by population protocols with a leader. *Distributed Computing*, 21(3):183–199, 2008.
- [AAE⁺17] D. Alistarh, J. Aspnes, D. Eisenstat, R. Gelashvili, and R.L. Rivest. Time-space trade-offs in population protocols. In *SODA*, pages 2560–2579, 2017.
- [AAG18] D. Alistarh, J. Aspnes, and R. Gelashvili. Space-optimal majority in population protocols. In *SODA*, pages 2221–2239, 2018.
- [AG15] D. Alistarh and R. Gelashvili. Polylogarithmic-time leader election in population protocols. In *ICALP*, pages 479–491, 2015.
- [AG18] Dan Alistarh and Rati Gelashvili. Recent algorithmic advances in population protocols. *SIGACT News*, 49(3):63–73, 2018.
- [AGV15] D. Alistarh, R. Gelashvili, and M. Vojnović. Fast and exact majority in population protocols. In *PODC*, pages 47–56, 2015.
- [Ang80] D. Angluin. Local and global properties in networks of processors (extended abstract). In *STOC*, pages 82–93, 1980.
- [AS91] H. Attiya and M. Snir. Better computing on the anonymous ring. *J. Algorithms*, 12(2):204–238, 1991.
- [ASW88] H. Attiya, M. Snir, and M.K. Warmuth. Computing on an anonymous ring. *J. ACM*, 35(4):845–875, 1988.
- [AW04] H. Attiya and J. Welch. *Distributed Computing: Fundamentals, Simulations and Advanced Topics*. John Wiley & Sons, 2004.
- [BCER17] A. Bilke, C. Cooper, R. Elsässer, and T. Radzik. Brief announcement: Population protocols for leader election and exact majority with $O(\log^2 n)$ states and $O(\log^2 n)$ convergence time. In *PODC*, pages 451–453, 2017.
- [BCN⁺15] L. Becchetti, A.E.F. Clementi, E. Natale, F. Pasquale, and R. Silvestri. Plurality consensus in the gossip model. In *SODA*, pages 371–390, 2015.
- [BEF⁺18] Petra Berenbrink, Robert Elsässer, Tom Friedetzky, Dominik Kaaser, Peter Kling, and Tomasz Radzik. A population protocol for exact majority with $o(\log^5/3 n)$ stabilization time and $\theta(\log n)$ states. In *32nd International Symposium on Distributed Computing, DISC 2018, New Orleans, LA, USA, October 15-19, 2018*, pages 10:1–10:18, 2018.
- [BFGK16] P. Berenbrink, T. Friedetzky, G. Giakkoupis, and P. Kling. Efficient plurality consensus, or: the benefits of cleaning up from time to time. In *ICALP*, pages 136:1–136:14, 2016.
- [BKKO18] P. Berenbrink, D. Kaaser, P. Kling, and L. Otterbach. Simple and efficient leader election. In *SOSA*, pages 9:1–9:11, 2018.
- [BSV⁺96] P. Boldi, Sh. Shammah, S. Vigna, B. Codenotti, P. Gemmel, and J. Simon. Symmetry breaking in anonymous networks: Characterizations. In *ISTCS*, pages 16–26, 1996.
- [Bur80] J.E. Burns. A formal model for message passing systems. Technical Report TR-91, Computer Science Department, Indiana University, September 1980.
- [BV99] P. Boldi and S. Vigna. Computing anonymously with arbitrary knowledge. In *PODC*, pages 181–188, 1999.
- [CCDS14] H.-L. Chen, R. Cummings, D. Doty, and D. Soloveichik. Speed faults in computation by chemical reaction networks. In *DISC*, pages 16–30, 2014.
- [CGK⁺15] J. Czyzowicz, L. Gąsieniec, A. Kosowski, E. Kranakis, P.G. Spirakis, and P. Uznański. On convergence and threshold properties of discrete lotka-volterra population protocols. In *ICALP*, pages 393–405, 2015.

- [CMN⁺11] I. Chatzigiannakis, O. Michail, S. Nikolaou, A. Pavlogiannis, and P.G. Spirakis. Passively mobile communicating machines that use restricted space. *Theor. Comput. Sci.*, 412(46):6469–6483, 2011.
- [DK18] B. Dudek and A. Kosowski. Universal protocols for information dissemination using emergent signals. In *STOC*, to appear, 2018.
- [Dot14] D. Doty. Timing in chemical reaction networks. In *SODA*, pages 772–784, 2014.
- [DP04] S. Dobrev and A. Pelc. Leader election in rings with nonunique labels. *Fundam. Inform.*, 59(4):333–347, 2004.
- [DP14] D. Dereniowski and A. Pelc. Leader election for anonymous asynchronous agents in arbitrary networks. *Distributed Computing*, 27(1):21–38, 2014.
- [DS15] D. Doty and D. Soloveichik. Stable leader election in population protocols requires linear time. In *DISC*, pages 602–616, 2015.
- [DV12] M. Draief and M. Vojnović. Convergence speed of binary interval consensus. *SIAM J. Control and Optimization*, 50(3):1087–1109, 2012.
- [ER18] Robert Elsässer and Tomasz Radzik. Recent results in population protocols for exact majority and leader election. *Bulletin of the EATCS*, 126, 2018.
- [Fis83] M.J. Fischer. The consensus problem in unreliable distributed systems (a brief survey). In *FCT*, pages 127–140, 1983.
- [FKK⁺04] P. Flocchini, E. Kranakis, D. Krizanc, F.L. Luccio, and N. Santoro. Sorting and election in anonymous asynchronous rings. *J. Parallel Distrib. Comput.*, 64(2):254–265, 2004.
- [FL87] G.N. Frederickson and N.A. Lynch. Electing a leader in a synchronous ring. *J. ACM*, 34(1):98–115, 1987.
- [FP11] E.G. Fusco and A. Pelc. How much memory is needed for leader election. *Distributed Computing*, 24(2):65–78, 2011.
- [GHM⁺16] L. Gąsieniec, D.D. Hamilton, R. Martin, P.G. Spirakis, and G. Stachowiak. Deterministic population protocols for exact majority and plurality. In *OPODIS*, pages 14:1–14:14, 2016.
- [GHMS15] L. Gąsieniec, D.D. Hamilton, R. Martin, and P.G. Spirakis. The match-maker: Constant-space distributed majority via random walks. In *SSS*, pages 67–80, 2015.
- [GMP16] Ch. Glacet, A. Miller, and A. Pelc. Time vs. information tradeoffs for leader election in anonymous trees. In *SODA*, pages 600–609, 2016.
- [GP16] M. Ghaffari and M. Parter. A polylogarithmic gossip algorithm for plurality consensus. In *PODC*, pages 117–126, 2016.
- [GS18] L. Gąsieniec and G. Stachowiak. Fast space optimal leader election in population protocols. In *SODA*, pages 2653–2667, 2018.
- [GSU18] Leszek Gąsieniec, Grzegorz Stachowiak, and Przemysław Uznanski. Almost logarithmic-time space optimal leader election in population protocols. *CoRR*, abs/1802.06867, 2018.
- [HKM⁺08] M. Amine Haddar, A. Hadj Kacem, Y. Métivier, M. Mosbah, and M. Jmaiel. Electing a leader in the local computation model using mobile agents. In *AICCSA*, pages 473–480, 2008.
- [HS80] D.S. Hirschberg and J.B. Sinclair. Decentralized extrema-finding in circular configurations of processors. *Commun. ACM*, 23(11):627–628, 1980.
- [Lan77] G. Le Lann. Distributed systems - towards a formal approach. In *IFIP Congress*, pages 155–160, 1977.
- [MCS11] O. Michail, I. Chatzigiannakis, and P.G. Spirakis. *New Models for Population Protocols*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, 2011.
- [MNRS14] G.B. Mertzios, S.E. Nikolettseas, Ch. Raptopoulos, and P.G. Spirakis. Determining majority in networks with local interactions and very small local memory. In *ICALP*, pages 871–882, 2014.
- [Pet82] G.L. Peterson. An $o(n \log n)$ unidirectional algorithm for the circular extrema problem. *ACM Trans. Program. Lang. Syst.*, 4(4):758–762, October 1982.

- [SOI⁺18] Yuichi Sudo, Fukuhito Ooshita, Taisuke Izumi, Hirotugu Kakugawa, and Toshimitsu Masuzawa. Logarithmic expected-time leader election in population protocol model. *CoRR*, abs/1812.11309, 2018.
- [YK89] M. Yamashita and T. Kameda. Electing a leader when processor identity numbers are not distinct (extended abstract). In *Distributed Algorithms, 3rd International Workshop, Nice, France, September 26-28, 1989, Proceedings*, pages 303–314, 1989.
- [YK96] M. Yamashita and T. Kameda. Computing on anonymous networks: Part 1 - characterizing the solvable cases. *IEEE Trans. Parallel Distrib. Syst.*, 7(1):69–89, 1996.

A Appendix (missing proofs)

Lemma 4.1. *With high probability, only $\mathcal{O}(n/\log n)$ agents are not initialised in the course of the protocol, i.e., $n - \mathcal{O}(n/\log n)$ agents join C, I or L during the first $\mathcal{O}(n \log n)$ interactions.*

Proof. By Theorem 3.2, the first round of the phase clock is completed with high probability during the first $d \cdot n \log n$ interactions, for some constant d . We first show that after $4 \cdot n \log n$ interactions at most $n/\log n$ not yet initialised in state 0 agents remain. Let \mathcal{X} be the random variable denoting the number of agents in state 0.

Assume $\mathcal{X} = \alpha n$, for some $\alpha > 0$. We prove that the number of interactions it takes to reduce \mathcal{X} by a factor of 2 is at most $4n/\alpha$, with very high probability. Let σ be a 0-1 sequence of length $4n/\alpha$ referring to the relevant $4n/\alpha$ interactions. In this sequence an entry is set to 1 if during the corresponding interaction the number of not yet initialised agents is reduced, and 0 otherwise. For as long as $\mathcal{X} > \alpha n/2$, the probability of having 1 at each position in σ is at least $\alpha^2/4$ and in turn the expected number of 1s in σ is at least αn . Thus by Chernoff bound the number of 1s in σ is at least $\alpha n/2$ with very high probability. This implies that at least $\alpha n/2$ agents in state 0 get initialised. And iterating this process $\log \log n$ times we get reduction of agents in state 0 to $n/\log n$ in at most $\mathcal{O}(n \log n)$ consecutive interactions. A similar reasoning can be used for agents in the intermediate state X in the next (subsequent to reduction of agents in state 0) $\mathcal{O}(n \log n)$ interactions. Thus one can conclude that after $\mathcal{O}(n \log n)$ initial interactions the number of not yet initialised agents is at most $2n/\log n$. \square

Lemma 7.1. *After the first round of the clock $D_\ell = n4^{-\ell}(1 \pm o(1))$ whp.*

Proof. Let $D'_\ell = D_\ell + \dots + D_\Psi$ be the number of inhibitor agents reaching slowness ℓ or higher and $p = \frac{n_c}{n}$ be the ratio of coins in the population. By Lemma 4.1 after $\mathcal{O}(n \log n)$ interactions of the first round $p = \frac{1}{4} - \mathcal{O}(1/\log n)$ and remains stable with high probability. An inhibitor agent reaches level ℓ by a series of ℓ successful synthetic coin flips, which happens with probability $p^\ell = 4^{-\ell}(1 - \ell \cdot \mathcal{O}(1/\log n))$. By Chernoff bound we have $D'_\ell = n_I \cdot p^\ell \pm \mathcal{O}(\sqrt{p^\ell \cdot n_I \log n_I})$, with high probability. We have $\ell = \mathcal{O}(\log \log n)$, $n_I = \Theta(n)$ and $D_\ell = D'_\ell - D'_{\ell+1}$, for $\ell < \Psi$ and $D_\Psi = D'_\Psi$. Thus there exists $D'_\ell = 4^{-\ell} \cdot n \cdot (1 \pm o(1))$ and the claimed bound holds.

In addition, we observe that with high probability during the initial $\Theta(n \log n)$ interactions each inhibitor agent experiences $\Omega(\log n)$ interactions with coin agents, determining its **drag** during the second round of the clock. \square