

# Polylogarithmic-Time Leader Election in Population Protocols

Dan Alistarh  
Microsoft Research

Rati Gelashvili  
MIT

## Abstract

Population protocols are networks of finite-state agents, interacting randomly, and updating their state using simple rules. Despite their extreme simplicity, these systems have been shown to cooperatively perform complex computational tasks, such as simulating register machines to compute standard arithmetic functions. The election of a unique *leader agent* is a key requirement in such computational constructions. Yet, the fastest currently known population protocol for electing a leader only has *polynomial* convergence time.

In this paper, we give the first population protocol for leader election with *poly-logarithmic* convergence time. The protocol structure is quite simple: each node has an associated value, and is either a *leader* (still in contention) or a *minion* (following some leader). A leader keeps incrementing its value and “defeats” other leaders in one-to-one interactions, and will drop from contention and become a minion if it meets a leader with higher value. Importantly, a leader also drops out if it meets a *minion* with higher absolute value. While these rules are quite simple, the proof that this algorithm achieves poly-logarithmic convergence time is non-trivial. In particular, the argument combines careful use of concentration inequalities with anti-concentration bounds, showing that the leaders’ values become spread apart as the execution progresses, which in turn implies that straggling leaders get quickly eliminated. We complement our analysis with empirical results, showing that our protocol converges fast, even for large network sizes.

## 1 Introduction

Recently, there has been significant interest in modeling and analyzing interactions arising in biological or bio-chemical systems through an algorithmic lens. Several interesting computational models have been proposed for such networks, for example, the cellular automata model [Neu66], the stone-age distributed computing model [EW13], or the population model [AAD<sup>+</sup>06].

In particular, population protocols [AAD<sup>+</sup>06], which are the focus of this paper, consist of a set of  $n$  finite-state nodes interacting in pairs, where each interaction may update the states of both participants. The goal is to have all nodes converge on an output value, which represents the result of the computation, usually a predicate on the initial state of the nodes. The set of interactions occurring at each step is assumed to be decided by an adversarial scheduler, which is usually subject to some fairness conditions. The standard scheduler when computing convergence bounds is the *probabilistic (uniform random) scheduler* [AAE08, PVV09, DV12, MNRS14], which picks the next pair to interact uniformly at random in each step. We adopt this probabilistic scheduler model in this paper. (Some references refer to this model as the *probabilistic* population model.) The fundamental measure of convergence is *parallel time*, defined as the number of scheduler steps until convergence, divided by  $n$ .<sup>1</sup>

The class of predicates computable by population protocols is now well-understood [AAD<sup>+</sup>06, AAE06, AAER07] to consist precisely of *semilinear predicates*, i.e. predicates definable in first-order

---

<sup>1</sup>An alternative definition is when reactions occur in parallel according to a Poisson process [PVV09, DV12].

Presburger arithmetic. The first such construction was given in [AAD<sup>+</sup>06], and later improved in terms of convergence time in [AAE06]. A parallel line of research studied the computability of deterministic functions in chemical reaction networks, which are also instances of population protocols [CDS14]. All three constructions fundamentally rely on the election of a single initial *leader* node, which co-ordinates phases of computation.

Reference [AAD<sup>+</sup>06] gives a simple protocol for electing a leader from a uniform population, based on the natural idea of having leaders eliminate each other directly through symmetry breaking. Unfortunately, this strategy takes at least linear parallel time in the number of nodes  $n$ : for instance, once this algorithm reaches *two* surviving leaders, it will require  $\Omega(n^2)$  additional interactions for these two leaders to meet. Reference [AAE06] proposes a significantly more complex protocol, conjectured to be sub-linear, and whose convergence is only studied experimentally. These references posit the existence of a sublinear-time population protocol for leader election as a “pressing” open problem in the area. In fact, the existence of a poly-logarithmic leader election protocol would imply that *any semilinear predicate* is computable in poly-logarithmic time by a *uniform population* [AAE06].

**Contribution:** In this paper, we solve this problem by proposing a new population protocol for leader election, which converges in  $O(\log^5 n)$  expected parallel time. Our protocol, called *LM* for *Leader-Minion*, roughly works as follows. Throughout the execution, each node is either a *leader*, meaning that it can still win, or a *minion*, following some leader. Each node state is associated to some *absolute value*, which is a positive integer, and with a *sign*, positive if the node is still in contention, and negative if the node has become a minion.

If two leaders meet, the one with the larger absolute value survives, and increments its value, while the other drops out, becoming a minion, and adopting the other node’s value, but with a negative sign. (If both leaders have the same value, they both increment it and continue.) If a leader meets a minion with *smaller* absolute value than its own, it increments its value, while the minion simply adopts the leader’s value, but keeps the negative sign. Conversely, if a leader meets a minion with *larger* absolute value than its own, then the leader will drop out of contention, adopting the minion’s value, with negative sign. Finally, if two minions meet, they update their values to the maximum absolute value between them, but with a *negative* sign.

These rules ensure that, eventually, a single leader survives. While the protocol is relatively simple, the proof of poly-logarithmic time convergence is non-trivial. In particular, the efficiency of the algorithm hinges on the minion mechanism, which ensures that a leader with high absolute value can eliminate other contenders in the system, without having to directly interact with them.

Roughly, the argument is based on two technical insights. First, consider two leaders at a given time  $T$ , whose (positive) values are at least  $\Theta(\log^2 n)$  apart. Then, we show that, within  $O(\log^2 n)$  parallel time from  $T$ , the node holding the smaller value has become a minion, with high probability. Intuitively, this holds since 1) this node will probably meet either the other leader or one of its minions within this time interval, and 2) it cannot increase its count fast enough to avoid defeat. For the second part of the argument, we show via anti-concentration that, after parallel time  $\Theta(\log^4 n)$  in the execution, the values corresponding to an arbitrary pair of nodes will be separated by at least  $\Omega(\log^2 n)$ .

We ensure that the values of nodes cannot grow beyond a certain threshold, and set the threshold in such a way that the total number of states is  $\Theta(\log^5 n)$ . We show that with high probability the leader will be elected before the values of the nodes reach the threshold. In the other case, remaining leaders with threshold values engage in a *backup* dynamics where minions are irrelevant and leaders defeat each other when they meet based on random binary indicators which are set using the randomness of the scheduler. This process is slower but correct and only happens with very low probability, allowing to conclude that the *LM* algorithm converges to a single leader within  $O(\log^5 n)$  parallel time, both with high probability and in expectation.

In population protocols, in every interaction, one node is said to be the *initiator*, the other is the *responder*, and the state update rules can use this distinction. This would allow a leader (the initiator in the interaction) to defeat another leader with the same value (responder in the interaction), in our protocol, and could also simplify the backup dynamics of our algorithm. However, our *LM* algorithm has the nice property that the state update rules can be made completely symmetric with regards to the initiator and responder roles.<sup>2</sup>

Summing up, we give the first poly-logarithmic time protocol for electing a leader from a uniform population. We note that there exists a trivial  $\Omega(\log n)$  parallel time lower bound, as  $\Omega(n \log n)$  interactions are necessary to ensure that each node interacts at least once, and clearly no algorithm may converge to a single leader before this time. Our algorithm needs  $O(\text{polylog } n)$  memory states per node, i.e.  $O(\log \log n)$  memory bits per node. We complement our analysis with empirical data, suggesting that the convergence time of our protocol is close to logarithmic, and that in fact the asymptotic constants are small.

**Related Work:** We restrict our attention to work in the population model. The framework of population protocols was formally introduced in reference [AAD<sup>+</sup>06], to model interactions arising in biological, chemical, or sensor networks. It sparked research into its computational power [AAD<sup>+</sup>06, AAE06, AAER07], and into the time complexity of fundamental tasks such as majority [AAE08, PVV09, DV12, MNRS14, AGV15], and leader election [AAD<sup>+</sup>06, AAE06].<sup>3</sup> References interested in *computability* consider an adversarial scheduler which is restricted to be *fair*, e.g., where each agent interacts with every other agent infinitely many times. For complexity bounds, the standard scheduler is *uniform*, scheduling each pair uniformly at random at each step, e.g., [AAE08, PVV09, DV12, MNRS14]. This model is also known as the *probabilistic* population model.

To the best of our knowledge, no population protocol for electing a leader with sub-linear convergence time was known before our work. References [AAD<sup>+</sup>06, AAE06, CDS14] present leader-based frameworks for population computations, assuming the existence of such a node. The existence of such a sub-linear protocol is stated as an open problem in [AAD<sup>+</sup>06, AAE06]. Reference [DH13] proposes a *leader-less* framework for population computation.

A parallel line of research studied *self-stabilizing* population protocols, e.g., [AAFJ06, FJ06, SNY<sup>+</sup>10], that is, protocols which can converge to a correct solution from an arbitrary initial state. It is known that stable leader election is impossible in such systems [AAFJ06]; references [FJ06, SNY<sup>+</sup>10] circumvent this impossibility by relaxing the problem semantics. Our algorithm is not affected by this result since it is not self-stabilizing.

## 2 Preliminaries

**Population Protocols:** We assume a population consisting of  $n$  agents, or nodes, each executing as a deterministic state machine with states from a finite set  $Q$ , with a finite set of input symbols  $X \subseteq Q$ , a finite set of output symbols  $Y$ , a transition function  $\delta : Q \times Q \rightarrow Q \times Q$ , and an output function  $\gamma : Q \rightarrow Y$ . Initially, each agent starts with an input from the set  $X$ , and proceeds to update its state following interactions with other agents, according to the transition function  $\delta$ . For simplicity of exposition, we assume that agents have identifiers from the set  $V = \{1, 2, \dots, n\}$ , although these identifiers are not known to agents, and not used by the protocol.

<sup>2</sup>For this reason, *LM* algorithm works for  $n > 2$  nodes, because to elect a leader among two nodes it is necessary to rely on the initiator-responder role distinction.

<sup>3</sup>The best known upper bound for deterministic majority is of  $O(\log n \log s + \log n/(\epsilon s))$  parallel time [AGV15], where  $n$  is the number of nodes,  $s$  is the number of states per node, and  $\epsilon$  is the initial node difference between the two input states. The two problems are complementary, and no complexity-preserving transformations exist, to our knowledge.

The agents' interactions proceed according to a directed *interaction graph*  $G$  without self-loops, whose edges indicate possible agent interactions. Usually, the graph  $G$  is considered to be the complete graph on  $n$  vertices, a convention we also adopt in this paper.

The execution proceeds in *steps*, or *rounds*, where in each step a new edge  $(u, w)$  is chosen uniformly at random from the set of edges of  $G$ . Each of the two chosen agents updates its state according to the function  $\delta$ .

**Parallel Time:** The above setup considers sequential interactions; however, in general, interactions between pairs of distinct agents are independent, and are usually considered as occurring in parallel. In particular, it is customary to define one unit of *parallel time* as  $n$  consecutive steps of the protocol.

**The Leader Election Problem:** In the *leader election* problem, all agents start in the same initial state  $A$ , i.e. the only state in the input set  $X = \{A\}$ . The output set is  $Y = \{Win, Lose\}$ .

A population protocol solves leader election within  $\ell$  steps with probability  $1 - \phi$ , if it holds with probability  $1 - \phi$  that for any configuration  $c : V \rightarrow Q$  reachable by the protocol after  $\geq \ell$  steps, there exists a unique agent  $i$  such that, (1) for the agent  $i$ ,  $\gamma(c(i)) = Win$ , and, (2) for any agent  $j \neq i$ ,  $\gamma(c(j)) = Lose$ .

### 3 The Leader Election Algorithm

In this section, we describe the *LM* leader election algorithm. The algorithm has an integer parameter  $m > 0$ , which we set to  $\Theta(\log^5 n)$ . Each state corresponds to an integer value from the set  $\{-m, -m+1, \dots, -2, -1, 1, 2, m-1, m, m+1\}$ . Respectively, there are  $2m+1$  different states. We will refer to states and values interchangeably. All nodes start in the same state corresponding to value 1.

The algorithm, specified in [Figure 1](#), consists of a set of simple deterministic update rules for the node state. In the pseudocode, the node states before an interaction are denoted by  $x$  and  $y$ , while their new states are given by  $x'$  and  $y'$ . All nodes start with value 1 and continue to interact according to these simple rules. We prove that all nodes except one will converge to negative values, and that convergence is fast with high probability. This solves the leader election problem since we can define  $\gamma$  as mapping only positive states to *Win* (a leader).

Since positive states translate to being a leader according to  $\gamma$ , we call a node a *contender* if it has a positive value, and a *minion* otherwise. We present the algorithm in detail below. The state updates (i.e. the transition function  $\delta$ ) of the *LM* algorithm are completely symmetric, that is, the new state  $x'$  depends on  $x$  and  $y$  (lines 2-4) exactly as  $y'$  depends on  $y$  and  $x$  (lines 5-7).

If a node is a contender and has absolute value not less than the absolute value of the interaction partner, then the node remains a contender and updates its value using the *contend-priority* function (lines 3 and 6). The new value will be one larger than the previous value except when the previous value was  $m+1$ , in which case the new value will be  $m$ .

If a node had a smaller absolute value than its interaction partner, or was a minion already, then the node will be a minion after the interaction. It will set its value using the *minion-priority* function, to either  $-\max(|x|, |y|)$ , or  $-m$  if the maximum was  $m+1$  (lines 4 and 7).

Values  $m+1$  and  $m$  are treated exactly the same way by minions (essentially corresponding to  $-m$ ). These values serve as a binary tie-breaker among the contenders that ever reach the value  $m$ , as will become clear from the analysis.

### 4 Analysis

In this section, we provide a complete analysis of our leader election algorithm.

---

**Parameters:**

$m$ , an integer  $> 0$ , set to  $\Theta(\log^5 n)$

**State Space:**

$LeaderStates = \{1, 2, \dots, m-1, m, m+1\}$ ,

$MinionStates = \{-1, -2, \dots, -m+1, -m\}$ ,

**Input:** States of two nodes,  $x$  and  $y$

**Output:** Updated states  $x'$  and  $y'$

**Auxiliary Procedures:**

$is-contender(x) = \begin{cases} \text{true} & \text{if } x \in LeaderStates; \\ \text{false} & \text{otherwise.} \end{cases}$

$contend-priority(x, y) = \begin{cases} m & \text{if } \max(|x|, |y|) = m+1; \\ \max(|x|, |y|) + 1 & \text{otherwise.} \end{cases}$

$minion-priority(x, y) = \begin{cases} -m & \text{if } \max(|x|, |y|) = m+1; \\ -\max(|x|, |y|) & \text{otherwise.} \end{cases}$

```
1 procedure update( $x, y$ )
2   if  $is-contender(x)$  and  $|x| \geq |y|$  then
3      $x' \leftarrow contend-priority(x, y)$ 
4   else  $x' \leftarrow minion-priority(x, y)$ 
5   if  $is-contender(y)$  and  $|y| \geq |x|$  then
6      $y' \leftarrow contend-priority(x, y)$ 
7   else  $y' \leftarrow minion-priority(x, y)$ 
```

---

Figure 1: The state update rules for the *LM* algorithm.

**Notation:** Throughout this proof, we denote the set of  $n$  nodes executing the protocol by  $V$ . We measure execution time in discrete steps (rounds), where each step corresponds to an interaction. The *configuration* at a given time  $t$  is a function  $c : V \rightarrow Q$ , where  $c(v)$  is the state of the node  $v$  at time  $t$ . (We omit the explicit time  $t$  when clear from the context.) We call a node *contender* when the value associated with its state is positive, and a *minion* when the value is negative. As previously discussed, we also assume that  $n > 2$ .

We first prove that the algorithm never eliminates all contenders and that a configuration with a single contender means that a leader is elected.

**Lemma 4.1.** *There is always at least one contender in the system. After an execution reaches a configuration with only a single node  $v$  being a contender, then from this point,  $v$  will have  $c(v) > 0$  in every reachable future configuration  $c$ , and there may never be another contender.*

*Proof.* By the structure of the algorithm, a node starts as a contender and may become a minion during an execution, but a minion may never become a contender again. Moreover, an absolute value associated with the state of a minion node can only increase to an absolute value of an interaction partner. Finally, an absolute value may never decrease except from  $m+1$  to  $m$ .

Let us assume for contradiction that an execution reaches a configuration where all nodes are minions. Consider this time point  $T$  and let the maximum absolute value of the nodes be  $u$ . Because the minions cannot increase the maximum absolute value in the system, there must have been a contender with value  $u$  during the execution before time  $T$ . In order for this contender to have become a minion by time  $T$ , it must have interacted with another node with an absolute value strictly larger than  $u$ . However, an absolute value of a node never decreases except from  $m+1$  to  $m$ , and despite existence of a larger absolute value than  $u$  before time  $T$ ,  $u$  was the largest absolute value at time  $T$ , thus  $u$  must be equal to  $m$ . But after such an interaction the second node that was in the state  $m+1$  remains a contender with value  $m$ . Now it must have interacted with yet another node of value  $m+1$  before time  $T$  in order to itself become a minion, but then this another

node is left as a contender with value  $m$  and the same reasoning applies to it. Our proof follows by infinite descent.

Consequently, whenever there is a single contender in the system, it must have the largest absolute value. Otherwise, it could interact with a node with a larger absolute value and become a minion itself, contradicting the proof above that all nodes may never be minions. Because of this invariant, the only contender may never become a minion and the minions can never become contenders again.  $\square$

Now we turn our attention to the convergence speed (assuming  $n > 2$ ) of the *LM* algorithm. Our goal is bound the number of rounds necessary to eliminate all except a single contender. In order for a contender to get eliminated, it must come across a larger value of another contender, the value possibly conducted through a chain of multiple minions via multiple interactions.

We first show by a rumor spreading argument that if the difference between the values of two contenders is large enough, then the contender with the smaller value will become a minion within the next  $O(n \log^2 n)$  rounds, with high probability. Then we use anti-concentration bounds to establish that for any two fixed contenders, given that no absolute value in the system reaches  $m$ , after every  $O(n \log^4 n)$  rounds the difference between their values is large enough with constant probability.

**Lemma 4.2.** *Fix a constant  $\gamma \geq 8$ , and consider two contender nodes with values  $u_1$  and  $u_2$ , where  $u_1 - u_2 \geq 4\gamma \log^2 n$  at time  $T$ . Then, after  $\gamma n \log^2 n$  rounds from  $T$ , the node that initially held the value  $u_2$  will be a minion with probability at least  $1 - \frac{2}{n^{\gamma/8}}$ .*

*Proof.* Call a node that has an absolute value of at least  $u_1$  an *up-to-date* node, and *out-of-date* otherwise. Initially, at least one node is up-to-date. We will show that if in some configuration  $x < n$  nodes are up-to-date, after a *phase* of  $\frac{\gamma n(n-1) \log n}{4x(n-x)}$  communication rounds, at least  $x + 1$  nodes will be up-to-date with probability at least  $1 - \frac{1}{n^{\gamma/4}}$ . Up-to-date nodes may never become out-of-date. On the other hand, an out-of-date node becomes up-to-date itself after an interaction with any up-to-date node. If we have  $x$  up-to-date nodes, in each round, the probability that an out-of-date node interacts with an up-to-date node increasing the number of up-to-date nodes to  $x + 1$ , is  $\frac{2x(n-x)}{n(n-1)}$ . To upper bound the probability that such an interaction never happens during a phase of  $\frac{\gamma n(n-1) \log n}{4x(n-x)}$  rounds, we can consider a random variable  $Y \sim \text{Bin}\left(\frac{\gamma n(n-1) \log n}{4x(n-x)}, \frac{2x(n-x)}{n(n-1)}\right)$  and use Chernoff bound to establish:

$$\Pr[Y \leq 0] = \Pr\left[Y \leq \frac{\gamma}{2} \log n (1 - 1)\right] \leq \exp\left(-\frac{\gamma \log n}{2 \cdot 2}\right) \leq 2^{-\gamma \log n/4} < \frac{1}{n^{\gamma/4}},$$

Now applying Union Bound over at most  $n$  phases we get that with probability at least  $1 - \frac{n}{n^{\gamma/4}} \geq 1 - \frac{1}{n^{\gamma/8}}$ , after at most

$$\sum_{x=1}^{n-1} \frac{\gamma n(n-1) \log n}{4x(n-x)} \leq \frac{\gamma(n-1) \log n}{4} \sum_{x=1}^{n-1} \left(\frac{1}{x} + \frac{1}{n-x}\right) \leq \gamma n \log^2 n$$

rounds, all nodes will have value at least  $u_1$ . Notice that when a node that was initially holding value  $u_2$  becomes up-to-date and gets an absolute value  $u_1$  from an interaction, it must become a minion by the structure of the algorithm if its value before this interaction was still strictly smaller than  $u_1$ . Thus, we only need to show that the probability of selecting the node that initially had value  $u_2$  at least  $4\gamma \log^2 n$  times (so that its value can reach  $u_1$ ) during these  $\gamma n \log^2 n$  rounds is at most  $\frac{1}{n^{\gamma/8}}$ . The claim then follows by Union Bound.

In each round, the probability to select this node that initially held  $u_2$  is  $2/n$ . Let us describe the number of times it is selected in  $\gamma n \log^2 n$  rounds by considering a random variable  $Z \sim$

$\text{Bin}(\gamma n \log^2 n, 2/n)$ . By Chernoff Bound, the probability being selected at least  $4\gamma \log^2 n$  times in these rounds is at most:

$$\Pr [Z \geq 4\gamma \log^2 n] = \Pr [Z \geq (2\gamma \log^2 n) (1 + 1)] \leq \exp \left( -\frac{2\gamma}{3} \log^2 n \right) \leq \frac{1}{n^{\gamma/8}}$$

finishing the proof.  $\square$

The next Lemma shows that, after  $\Theta(n \log^4 n)$  rounds, the difference between the values of any two given contenders is high, with reasonable probability.

**Lemma 4.3.** *Fix an arbitrary time  $T$ , and a constant  $\gamma \geq 1$ . Consider any two contender nodes at time  $T$ , and time  $T_1$  which is  $32\gamma^2 n \log^4 n$  rounds after  $T$ .*

*If no absolute value of any node reaches  $m$  at any time until  $T_1$ , then, with probability at least  $\frac{1}{24} - \frac{1}{n^{8\gamma}}$ , at time  $T_1$ , either at least one of the two nodes have become minions, or the absolute value of the difference of the two nodes' values is at least  $4\gamma \log^2 n$ .*

*Proof.* We will assume that no absolute value reaches  $m$  at any point until time  $T_1$  and that the two nodes are still contenders at  $T_1$ . We should now prove that the difference of values is large enough.

Consider the  $32\gamma^2 n \log^4 n$  rounds following time  $T$ . If a round involves an interaction with exactly one of the two fixed nodes we call it a *spreading* round. For each round, probability that it will be spreading is  $\frac{4(n-2)}{n(n-1)}$  which for  $n > 2$  is at least  $2/n$ . So, we can describe the number of spreading rounds among the  $32\gamma^2 n \log^4 n$  rounds by considering a random variable  $X \sim \text{Bin}(32\gamma^2 n \log^4 n, 2/n)$ . Then, by Chernoff Bound, the probability of having at most  $32\gamma^2 \log^4 n$  spreading rounds is at most

$$\begin{aligned} \Pr [X \leq 32\gamma^2 \log^4 n] &= \Pr \left[ X \leq 64\gamma^2 \log^4 n \left( 1 - \frac{1}{2} \right) \right] \\ &\leq \exp \left( -\frac{64\gamma^2 \log^4 n}{2^2 \cdot 2} \right) \leq 2^{-8\gamma^2 \log^4 n} < \frac{1}{n^{8\gamma}}, \end{aligned}$$

Let us from now on focus on the high probability event that there are at least  $32\gamma^2 \log^4 n$  spreading rounds between times  $T$  and  $T_1$ , and prove that the desired difference will be large enough with probability  $\frac{1}{24}$ . This implies the claim by Union Bound with the above event (note that for  $n > 2$ ,  $\frac{1}{n^{8\gamma}} < \frac{1}{24}$  holds).

We assumed that both nodes remain contenders during the whole time, hence in each spreading round, a value of exactly one of them, with probability  $1/2$  each, increases by one. Without the loss of generality assume that at time  $T$ , the value of the first node was larger than or equal to the value of the second node. Let us now focus on the sum  $S$  of  $k$  independent uniformly distributed  $\pm 1$  Bernoulli trials  $x_i$  where  $1 \leq i \leq k$ , where each trial corresponds to a spreading round and outcome  $+1$  means that the value of the first node increased, while  $-1$  means that the value of the second node increased. In this terminology, we are done if we show that  $\Pr[S \geq 4\gamma \log^2 n] \geq \frac{1}{24}$  for  $k \geq 32\gamma^2 \log^4 n$  trials.

However, we have that:

$$\Pr[S \geq 4\gamma \log^2 n] \geq \frac{\Pr[|S| \geq 4\gamma \log^2 n]}{2} = \frac{\Pr[|S^2| \geq 16\gamma^2 \log^4 n]}{2} \quad (1)$$

$$\geq \frac{\Pr[|S^2| \geq k/2]}{2} = \frac{\Pr[|S^2| \geq \mathbb{E}[S^2]/2]}{2} \quad (2)$$

$$\geq \frac{1}{2^2 \cdot 2} \frac{\mathbb{E}[S^2]^2}{\mathbb{E}[S^4]} \geq \frac{1}{24} \quad (3)$$

where **1** follows from the symmetry of the sum with regards to the sign, that is, from  $\Pr[S > 4\gamma \log^2 n] = \Pr[S < -4\gamma \log^2 n]$ . For **2** we have used that  $k \geq 32\gamma^2 \log^4 n$  and  $\mathbb{E}[S^2] = k$  (more

about this below). Finally, to get [3](#) we use Paley-Zygmund inequality and the fact that  $\mathbb{E}[S^4] = 3k(k-1) + k \leq 3k^2$ . Evaluating  $\mathbb{E}[S^2]$  and  $\mathbb{E}[S^4]$  is simple by using definition of  $S$  and the linearity of expectation. The expectation of each term then is either 0 or 1 and it suffices to count the number of terms with expectation 1, which are exactly the terms where each multiplier is raised to an even power.  $\square$

Now we are ready to prove the convergence speed with high probability

**Theorem 4.4.** *There exists a constant  $\alpha$ , such that for any constant  $\beta \geq 3$  following holds: If we set  $m = \alpha\beta \log^5 n = \Theta(\log^5 n)$ , the algorithm elects a leader (i.e. reaches a configuration with a single contender) in at most  $O(n \log^5 n)$  rounds (i.e. parallel time  $O(\log^5 n)$ ) with probability at least  $1 - 1/n^\beta$ .*

*Proof.* Let us fix  $\gamma \geq 8$  large enough such that

$$1/24 - 1/n^{8\gamma} - 2/n^{\gamma/8} \geq 1/50. \quad (4)$$

Let  $\beta$  be any constant  $\geq 3$  and take  $\alpha = 800 \cdot (33\gamma^2)$ . We set  $m = \alpha\beta \log^5 n$  and consider the first  $\frac{\alpha\beta n \log^5 n}{4}$  communication rounds of the algorithm execution. For any fixed node, the probability that it interacts in each round is  $2/n$ . Let us describe the number of times a given node interacts within the first  $\frac{\alpha\beta n \log^5 n}{4}$  rounds by considering a random variable  $B \sim \text{Bin}(\frac{\alpha\beta n \log^5 n}{4}, 2/n)$ . By Chernoff Bound, the probability being selected at least  $m$  times in these rounds is at most:

$$\begin{aligned} \Pr[B \geq m] &= \Pr\left[B \geq \frac{\alpha\beta \log^5 n}{2} \cdot (1+1)\right] \\ &\leq \exp\left(-\frac{\alpha\beta}{6} \log^5 n\right) \leq 2^{-\frac{\alpha\beta}{6} \log^5 n} \leq \frac{1}{n^{800\beta}} \end{aligned}$$

Taking Union Bound over all  $n$  nodes, with probability at least  $1 - \frac{n}{n^{800\beta}}$ , all nodes interact strictly less than  $m$  times during the first  $\frac{\alpha\beta n \log^5 n}{4}$  rounds.

Let us from now on focus on the high probability event in above, which means that all absolute values are strictly less than  $m$  during the first  $\frac{\alpha\beta n \log^5 n}{4} = 200\beta \cdot (33\gamma^2)n \log^5 n$  rounds. For a fixed pair of nodes, this allows us to apply [Lemma 4.3](#) followed by [Lemma 4.2](#) (with parameter  $\gamma$ )  $\frac{200\beta \cdot (33\gamma^2)n \log^5 n}{32\gamma^2 n \log^4 n + \gamma n \log^2 n} \geq 200\beta \log n$  times. Each time, first by [Lemma 4.3](#), after  $32\gamma^2 n \log^4 n$  rounds with probability at least  $\frac{1}{24} - \frac{1}{n^{8\gamma}}$  the nodes end up with values at least  $4\gamma \log^2 n$  apart, in which case, after the next  $\gamma n \log^2 n$  rounds, by [Lemma 4.2](#), one of the nodes becomes a minion with probability at least  $1 - \frac{2}{n^{\gamma/8}}$ . By union bound and [\(4\)](#), each of the  $200\beta \log n$  times if both nodes are contenders, we get probability at least  $1/50$  that one of the nodes becomes a minion. Consider a random variable  $W \sim \text{Bin}(200\beta \log n, 1/50)$ . By Chernoff bound the probability that both nodes in a given pair are still contenders after the first  $\frac{\alpha\beta n \log^5 n}{4}$  rounds is at most:

$$\Pr[W \leq 0] = \Pr[W \leq 4\beta \log n (1-1)] \leq \exp\left(-\frac{4\beta \log n}{2}\right) \leq 2^{-2\beta \log n} < \frac{1}{n^{2\beta}},$$

By Union Bound over all  $n(n-1)/2 < n^2$  pairs, the probability that for each pair of nodes, one of them becomes a minion after the first  $\frac{\alpha\beta n \log^5 n}{4}$  is at least  $1 - \frac{n^2}{n^{2\beta}}$ . This implies that only one node can be a contender.

Finally, combining with the conditioned event that none of the nodes interact  $m$  or more times gives that after the first  $\frac{\alpha\beta n \log^5 n}{4} = O(n \log^5 n)$  rounds there must be a single contender with probability at least  $1 - \frac{n^2}{n^{2\beta}} - \frac{n}{n^{800\beta}} \geq 1 - \frac{1}{n^\beta}$  for  $\beta \geq 3$ . A single contender means that leader is elected by [Lemma 4.1](#).  $\square$

Finally, we can prove the expected convergence bound



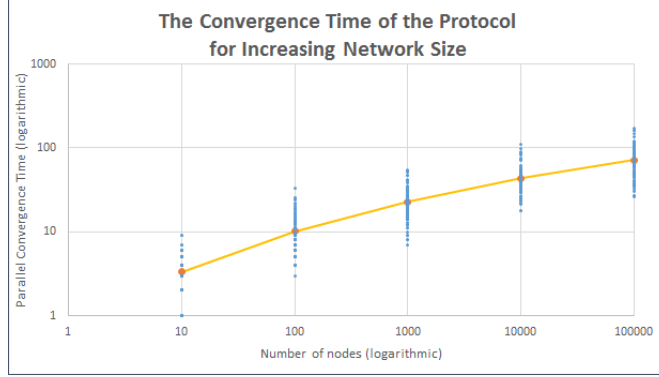


Figure 2: The performance of the *LM* protocol. Both axes are logarithmic. The dots represent the results of individual experiments (100 for each network size), while the solid line represents the mean value for each network size.

**Theorem 4.5.** *There is a setting of parameter  $m$  of the algorithm such that  $m = \Theta(\log^5 n)$ , such that the algorithm elects the leader in expected  $O(n \log^5 n)$  rounds of communication (i.e. parallel time  $O(\log^5 n)$ ).*

*Proof.* Let us prove that from any configuration, the algorithm elects a leader in expected  $O(n \log^5 n)$  rounds. By [Lemma 4.1](#), there is always a contender in the system and if there is only a single contender, then a leader is already elected. Now in a configuration with at least two contenders consider any two of them. If their values differ, then with probability at least  $1/n^2$  these two contenders will interact in the next round and the one with the lower value will become a minion (after which it may never be a contender again). If the values are the same, then with probability at least  $1/n$ , one of these nodes will interact with one of the other nodes in the next round, leading to a configuration where the values of our two nodes differ<sup>4</sup>, from where in the next round, independently, with probability at least  $1/n^2$  these nodes meet and one of them again becomes a minion. Hence, unless a leader is already elected, in any case, in every two rounds, with probability at least  $1/n^3$  the number of contenders decreases by 1.

Thus the expected number of rounds until the number of contenders decreases by 1 is at most  $2n^3$ . In any configuration there can be at most  $n$  contenders, thus the expected number of rounds until reaching a configuration with only a single contender is at most  $2(n-1)n^3 \leq 2n^4$  from any configuration.

Now using [Theorem 4.4](#) with  $\beta = 4$  we get that with probability at least  $1 - 1/n^4$  the algorithm converges after  $O(n \log^5 n)$  rounds. Otherwise, with probability at most  $1/n^4$  it ends up in some configuration from where it takes at most  $2n^4$  expected rounds to elect a leader. The total expected number of rounds is therefore also  $O(n \log^5 n) + O(1) = O(n \log^5 n)$ , i.e. parallel time  $O(\log^5 n)$ .  $\square$

## 5 Experiments and Discussion

**Empirical Data:** We have also measured the convergence time of our protocol for different network sizes. (Figure 2 presents the results in the form of a log-log plot.) The protocol converges to a single leader quite fast, e.g., in less than 100 units of parallel time for a network of size  $10^5$ . This

<sup>4</sup>This is always true, even when the new value is not larger, for instance when the values were equal to  $m + 1$ , the new value of one of the nodes will be  $m \neq m + 1$ .

suggests that the constants hidden in the asymptotic analysis are small. The shape of the curve confirms the poly-logarithmic behavior of the protocol.

**Discussion:** We have given the first population protocol to solve leader election in poly-logarithmic time. Together with the results of [AAE06], the existence of our protocol implies that population protocol can compute any semi-linear predicate on their input in time  $O(n \log^5 n)$ , with high probability. Our algorithm requires  $O(\log \log n)$  bits of memory per node.

Our result opens several avenues for future research. The first concerns *lower bounds*. We conjecture that the lower bound for leader election in population protocols is  $\omega(\log n)$ . Further, empirical data suggests that the analysis of our algorithm can be further tightened, cutting off logarithmic factors. It would also be interesting to see if there is a trade-off between the amount of memory available per node and the running time of the protocol.

## References

- [AAD<sup>+</sup>06] Dana Angluin, James Aspnes, Zoë Diamadi, Michael J Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed computing*, 18(4):235–253, 2006.
- [AAE06] Dana Angluin, James Aspnes, and David Eisenstat. Stably computable predicates are semilinear. In *Proceedings of PODC 2006*, pages 292–299, 2006.
- [AAE08] Dana Angluin, James Aspnes, and David Eisenstat. A simple population protocol for fast robust approximate majority. *Distributed Computing*, 21(2):87–102, July 2008.
- [AAER07] Dana Angluin, James Aspnes, David Eisenstat, and Eric Ruppert. The computational power of population protocols. *Distributed Computing*, 20(4):279–304, November 2007.
- [AAFJ06] Dana Angluin, James Aspnes, Michael J Fischer, and Hong Jiang. Self-stabilizing population protocols. In *Principles of Distributed Systems*, pages 103–117. Springer, 2006.
- [AGV15] Dan Alistarh, Rati Gelashvili, and Milan Vojnovic. Fast and exact majority in population protocols. Technical Report Available Online as MSR-TR-2015-13, 2015. Submitted to PODC 2015.
- [CDS14] Ho-Lin Chen, David Doty, and David Soloveichik. Deterministic function computation with chemical reaction networks. *Natural computing*, 13(4):517–534, 2014.
- [DH13] David Doty and Monir Hajiaghayi. Leaderless deterministic chemical reaction networks. In *DNA Computing and Molecular Programming*, pages 46–60. Springer, 2013.
- [DV12] Moez Draief and Milan Vojnovic. Convergence speed of binary interval consensus. *SIAM Journal on Control and Optimization*, 50(3):1087–1109, 2012.
- [EW13] Yuval Emek and Roger Wattenhofer. Stone age distributed computing. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing*, PODC '13, pages 137–146, New York, NY, USA, 2013. ACM.
- [FJ06] Michael Fischer and Hong Jiang. Self-stabilizing leader election in networks of finite-state anonymous agents. In *Principles of Distributed Systems*, pages 395–409. Springer, 2006.

- [MNRS14] George B. Mertzios, Sotiris E. Nikolettseas, Christoforos Raptopoulos, and Paul G. Spirakis. Determining majority in networks with local interactions and very small local memory. In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, ICALP '14, pages 871–882, 2014.
- [Neu66] John Von Neumann. *Theory of Self-Reproducing Automata*. University of Illinois Press, Champaign, IL, USA, 1966.
- [PVV09] Etienne Perron, Dinkar Vasudevan, and Milan Vojnovic. Using three states for binary consensus on complete graphs. In *INFOCOM 2009, IEEE*, pages 2527–2535. IEEE, 2009.
- [SNY<sup>+</sup>10] Yuichi Sudo, Junya Nakamura, Yukiko Yamauchi, Fukuhito Ooshita, Hirotsugu Kakugawa, and Toshimitsu Masuzawa. Loosely-stabilizing leader election in population protocol model. In *Structural Information and Communication Complexity*, pages 295–308. Springer, 2010.