University Libraries
University of Nevada, Las Vegas

Computer Science Faculty Publications

Computer Science

12-17-2018

Loosely-Stabilizing Leader Election with Polylogarithmic Convergence Time

Yuichi Sudo
Osaka University, y-sudou@ist.osaka-u.ac.jp

Fukuhito Ooshita
Nara Institute of Science and Technology, f-oosita@is.naist.jp

Hirotsugu Kukugawa
Osaka University, kakugawa@ist.osaka-u.ac.jp

Toshimitsu Masuzawa
Osaka University, masuzawa@ist.osaka-u.ac.jp

Ajoy K. Datta
University of Nevada, Las Vegas, ajoy.datta@unlv.edu

Follow this and additional works at: https://digitalscholarship.unlv.edu/compsci_fac_articles

Part of the Computer Sciences Commons

See next page for additional authors

Authors

Yuichi Sudo, Fukuhito Ooshita, Hirotsugu Kukugawa, Toshimitsu Masuzawa, Ajoy K. Datta, and Lawrence L. Larmore

Loosely-Stabilizing Leader Election with Polylogarithmic Convergence Time

Yuichi Sudo

Graduate School of Information Science and Technology, Osaka University, Japan
y-sudou@ist.osaka-u.ac.jp

Fukuhito Ooshita

Graduate School of Science and Technology, Nara Institute of Science and Technology, Japan
f-oosita@is.naist.jp

Hirotsugu Kakugawa

Graduate School of Information Science and Technology, Osaka University, Japan
kakugawa@ist.osaka-u.ac.jp

Toshimitsu Masuzawa

Graduate School of Information Science and Technology, Osaka University, Japan
masuzawa@ist.osaka-u.ac.jp

Ajoy K. Datta

Department of Computer Science, University of Nevada, Las Vegas, USA
ajoy.datta@unlv.edu

Lawrence L. Larmore

Department of Computer Science, University of Nevada, Las Vegas, USA
lawrence.larmore@unlv.edu

—— **Abstract** ——————————————————————————————

A loosely-stabilizing leader election protocol with polylogarithmic convergence time in the population protocol model is presented in this paper. In the population protocol model, which is a common abstract model of mobile sensor networks, it is known to be impossible to design a self-stabilizing leader election protocol. Thus, in our prior work, we introduced the concept of loose-stabilization, which is weaker than self-stabilization but has similar advantage as self-stabilization in practice. Following this work, several loosely-stabilizing leader election protocols are presented. The loosely-stabilizing leader election guarantees that, starting from an arbitrary configuration, the system reaches a safe configuration with a single leader within a relatively short time, and keeps the unique leader for an sufficiently long time thereafter. The convergence times of all the existing loosely-stabilizing protocols, i.e., the expected time to reach a safe configuration, are polynomial in n where n is the number of nodes (while the holding times to keep the unique leader are exponential in n). In this paper, a loosely-stabilizing protocol with polylogarithmic convergence time is presented. Its holding time is not exponential, but arbitrarily large polynomial in n.

1 Introduction

We consider the *population protocol* (PP) model [3] in this paper. A network called *population* consists of a large number of finite-state automata, called *agents*. Agents often make *interactions* (i.e., pairwise communication) each between a pair of agents by which they update their states. The interactions are opportunistic, that is, they are unknown and unpredictable (or predictable only with probability). Agents are strongly anonymous: they do not have identifiers and they cannot distinguish their neighbors with the same states. As with the majority of studies on population protocols, we assume that the network of agents is a complete graph, and that the scheduler selects an interacting pair of agents at each step uniformly at random.

In this paper, we focus on the problem of self-stabilizing leader election (SS-LE), which is one of the most important and well-studied problems in the PP model. Self-stabilizing leader election requires that starting from any configuration, a population reaches a safe configuration in which exactly one leader exists; and after that, the population keeps that leader forever. These requirements guarantee excellent tolerance against any finite number of transient faults. Since many protocols (whether self-stabilizing or non-stabilizing) in the literature work with the assumption that a unique leader exists [3, 4, 5], SS-LE is a key to improving the fault-tolerance of the PP model itself. However, SS-LE is strictly impossible in the PP model: no protocol can solve SS-LE unless every agent in the population knows the exact size of the population (i.e., the number of agents) [4]. This impossibility comes from a simple partitioning argument[1]. Thus, most of the studies extend (i.e., strengthen) the PP model to circumvent the impossibility. One approach of studies [7, 13, 18] assumes that every agent knows the exact value of n and focuses on the space complexity to solve SS-LE. Another approach [9, 6, 8] solves SS-LE by using *oracles*, which tell every agent whether or not there exists an agent in a leader-state.

To solve SS-LE in a more practical way, our previous work [14] introduces the concept of *loose-stabilization*, which relaxes the closure requirement of self-stabilization but keeps its advantage in practice. Specifically, starting from any initial configuration, the population must reach a safe configuration within a relatively short time; after that, the specification of the problem (the unique leader for leader election) must be sustained for a sufficiently long time, though not necessarily forever. In [14], we gave a loosely-stabilizing leader election (LS-LE) protocol assuming that every agent knows a common upper bound N of n. This protocol is practically equivalent to an SS-LE protocol since it maintains the unique leader for exponential time in n (that is, practically forever) after reaching a safe configuration within $O(N \log N)$ *parallel time*, which we will define later. The assumption that we can use an upper bound N of n is practical because the protocol works correctly even if we make a large overestimation of n, such as $N = 10n$. Recently, Izumi [11] give a method which improves the convergence time of this protocol to linear time, i.e., $O(N)$. In [15, 16, 17], LS-LE protocols are presented for a population where some pairs of agents may not have interactions, i.e., the interaction graph is not complete.

---

[1] Assume that a SS-LE protocol P works without knowledge of exact size n of the population. Then, there must exist $n_1, n_2, (n_1 < n_2)$ such that P works correctly both when $n = n_1$ and when $n = n_2$. Consider an execution of P on the population with size n_2 which starts from a safe configuration where exactly one leader exists. If some $n_2 - n_1$ agents including the unique leader do not interact for a sufficiently long time, the rest of the population consisting of the other n_1 agents must create a new leader to satisfy the convergence requirement of self-stabilization. However, this creation results in two leaders in the population, violating the closure requirement of self-stabilization.

Leader election is easily solved in linear parallel time in the PP model if one does not stick to self- or loose- stabilization. When we design non-stabilizing protocols, we can assume that all agents are in a specific state in the initial configuration. A non-stabilizing leader election protocol was first presented in [3], which elects a unique leader within $O(n)$ parallel time. Recently, Alistarh and Gelashvili [1] give a non-stabilizing leader election protocol with polylogarithmic parallel time. Starting from a specific initial configuration, their protocol elects a unique leader within $O(\log^3 n)$ parallel time, with the assumption that all agents share a common integer $m = \Theta(\log^3 n)$. More recently, Gąsieniec and Staehowiak [10] give an algorithm which converges in $O(\log^2 n)$ parallel time. (The space complexity of the algorithm [10] is also surprisingly small, i.e., $O(\log \log \log n)$ bits.) The key strategies used in these papers to achieve polylogarithmic time are interesting, but we can not utilize them for our purpose because both of them critically depend on the assumption that all agents have the same initial state.

A number of results in the PP model assume the uniformly random scheduler, that is, a pair of agents is chosen uniformly at random to interact at each step [3, 5, 11, 14, 2, 10, 1]. This assumption has been used mainly for evaluating the time complexity of protocols. We also adopt this assumption because the measure of time is crucial in the concept of loose-stabilization. In the PP model, time complexities, such as convergence time and holding time, are often evaluated in *parallel time*, which is defined as the expected number of interactions or steps, divided by n (i.e., the number of agents). This is a natural measure of time because, in practice interactions typically occur in parallel in the population.

1.1 Our Contribution

We now present an LS-LE protocol with polylogarithmic convergence time and polynomial holding time. Here, and for the remainder of this section, when we discuss time complexity, we shall always presume parallel computation. To the best of our knowledge, all previously published LS-LE protocols have at least linear convergence time, and exponential holding time, as shown in Table 1. Our protocol P_{PL} breaks through the barrier of linear convergence time. There is a convergence time/holding time trade-off. Given a parameter $c \geq 1$ and an upper bound N of n, our protocol converges within $O(c \log^3 N)$ time, and has $\Omega(cn^{10c})$ holding time. Although our expected holding time is not exponential in N, it grows as an exponential function of c. Also, the convergence time of P_{PL} does not suffer much from large overestimation of n; it is always $O(c \log^3 n)$ as long as N is polynomial in n. It is worth mentioning that P_{PL} has small space complexity. Each agent needs only $O(\log \log N)$ bits of memory to store all variables of P_{PL}. We can say that this is small space when we consider that any self-stabilizing leader election protocol, which requires knowledge of n, needs the space of at least $\lceil \log n \rceil$ bits [7]. These performances of our protocol cannot be obtained if we require exponential holding time: it is proven by Izumi [11] that any LS-LE protocol whose holding time is exponential requires $\Omega(N)$ convergence time and $\Omega(\log N)$-bit space at each agent.

We obtain a useful tool when analyzing the convergence and holding time of P_{PL}. Let `var` be a variable of some algorithm. Consider that two agents interact and the values of `var` in the two agents change from x and y to x' and y'. We call `var` a *propagating variable* if both $x' \geq \max(x-1, y-1, 0)$ and $y' \geq \max(x-1, y-1, 0)$ are always guaranteed. To the best of our knowledge, all loosely-stabilizing protocol (and possibly some non-stabilizing protocols) use propagating variables. Let z and Δ be any integers. If some agent has value $z + \Delta$ for a propagating variable `var` and Δ is sufficiently large, the propagating property of `var` guarantees that all the agents obtain values larger than z in `var` in a short time;

■ **Table 1** Self/Loosely-stabilizing leader election in the PP model (shown in parallel time).

Protocol	Type	Knowledge	Convergence Time	Holding Time	Agent Space (bits)
[14]	loose-stab.	N	$O(N \log N)$	$\Omega(e^N)$	$O(\log N)$
[11]	loose-stab.	N	$O(N)$	$\Omega(e^N)$	$O(\log N)$
P_{PL}	loose-stab.	N	$O(c \log^3 N)$	$\Omega(cn^{10c})$	$O(\log \log N)$
[7]	self-stab.	(exact) n	polynomial	∞	$O(\log n)$

specifically, in $O(\log n)$ time with high probability. The interesting question is how large this Δ should be. In the analysis of [14], a trivially sufficient value $\Delta = \Theta(n)$ is used. However, this linear value is useless for designing a protocol with polylogarithmic convergence time. In this paper, we prove that $\Delta = \Theta(\log n)$ is sufficient to propagate of values larger than z to the whole population (Lemma 6). This result may seem trivial for experts at the first glance, however, it is not trivial. This is because, whereas every agent participates in each interaction of an execution with probability $2/n$, the probability that a *virtual agent* (defined later) participates in each interaction may not be equal to $2/n$. We prove $\Delta = \Theta(\log n)$ by using three different kinds of random variables and bounding Δ by the sum of them. As we will see later, this result is very helpful in the analysis of the behavior of a protocol with propagating variables.

2 Preliminaries

In this section, we describe our model of computation. We denote the set of integers $\{z \in \mathbb{N} \mid x \leq z \leq y\}$ by $[x, y]$, and denote the n^{th} harmonic number by $H_n = \sum_{k=1}^{n} \frac{1}{k}$. We write the natural logarithm of x as $\ln x$; we indicate the base of other logarithms of x, such as $\log_2 x$.

A *population* is a network consisting of *agents*. We denote the set of all the agents by V and let $n = |V|$. We assume that a population is complete graph, thus every pair of agents (u, v) can interact, where u serves as the *initiator* and v serves as the *responder* of the interaction.

A *protocol* $P(Q, Y, T, \pi_{out})$ consists of a finite set Q of states, a finite set Y of output symbols, a transition function $T : Q \times Q \to Q \times Q$, and an output function $\pi_{out} : Q \to Y$. When two agents interact, T determines their next states according to their current states. The *output* of an agent is determined by π_{out}: the output of an agent in state q is $\pi_{out}(q)$.

A *configuration* is a mapping $C : V \to Q$ that specifies the states of all the agents. We denote the set of all configurations of protocol P by $\mathcal{C}_{\text{all}}(P)$. We say that a configuration C changes to C' by the interaction $e = (u, v)$, denoted by $C \xrightarrow{e} C'$, if $(C'(u), C'(v)) = T(C(u), C(v))$ and $C'(w) = C(w)$ for all $w \in V \setminus \{u, v\}$.

A *schedule* $\gamma = \gamma_0, \gamma_1, \cdots = (u_0, v_0), (u_1, v_1), \ldots$ is a sequence of interactions. A schedule determines which interaction occurs at each time, i.e., interaction γ_t happens at time t under schedule γ. In particular, we consider a *uniformly random scheduler* $\Gamma = \Gamma_0, \Gamma_1, \ldots$ in this paper: each Γ_t is a random variable such that $\Pr(\Gamma_t = (u, v)) = \frac{1}{n(n-1)}$ for any $t \geq 0$ and any distinct $u, v \in V$. Note that we use capital letter Γ for this uniform random scheduler while we refer a deterministic schedule with a lower case such as γ. Given an initial configuration C_0 and a schedule γ, the *execution* of protocol P is defined as $\Xi_P(C_0, \gamma) = C_0, C_1, \ldots$ such that $C_t \xrightarrow{\gamma_t} C_{t+1}$ for all $t \geq 0$. Note that the execution $\Xi_P(C_0, \Gamma) = C_0, C_1, \ldots$ under the uniformly random scheduler is a sequence of configurations where each C_i is a random variable. For a schedule $\gamma = \gamma_0, \gamma_1, \ldots$ and any $t \geq 0$, we say that agent $v \in V$ *participates* in γ_t if v is either the initiator or the responder of γ_t.

The leader election problem requires that every agent should output L or F which means "leader" or "follower" respectively. The specification of leader election, denoted by LE, requires that there exists one agent v such that v is always a leader and all other agents are always followers throughout an execution. We define $\mathrm{EIH}_P(C, LE)$ as the expected number of interactions during which an execution $\Xi_P(C, \Gamma)$ starting from a configuration $C \in \mathcal{C}_{\mathrm{all}}(P)$ keeps LE (i.e., the expected number of interactions until $\Xi_P(C, \Gamma)$ deviates from LE). For any set $\mathcal{S} \subseteq \mathcal{C}_{\mathrm{all}}(P)$ of configurations, we also define $\mathrm{EIC}_P(C, \mathcal{S})$ as the expected number of interactions required for the population to enter a configuration in $\mathcal{S}$ in an execution $\Xi_P(C, \Gamma)$ starting from a configuration $C \in \mathcal{C}_{\mathrm{all}}(P)$. The notation EIH (resp. EIC) stands for the Expected number of Interactions to Hold (resp. Converge).

▶ **Definition 1** (Loose-stabilizing leader election [14]). *Protocol $P(Q, Y, T, \pi_{out})$ is an (α, β)-loosely-stabilizing leader election protocol if there exists a set $\mathcal{S}$ of configurations satisfying the following two inequalities:*

$$\max_{C \in \mathcal{C}_{\mathrm{all}}(P)} \mathrm{EIC}_P(C, \mathcal{S}) \le \alpha \quad \text{and} \quad \min_{C \in \mathcal{S}} \mathrm{EIH}_P(C, LE) \ge \beta.$$

We call a configuration in $\mathcal{S}$ in the above definition *a safe configuration* of P. In terms of parallel time, (α, β)-loosely-stabilizing protocol leader election protocol P reaches a safe configuration within α/n parallel time in expectation and keeps the unique leader for β/n parallel time in expectation thereafter. We call α/n and β/n *the expected holding time* and *the expected convergence time* of P.

Throughout the paper, we will use the following three variants of Chernoff bounds.

▶ **Lemma 2** ([12], Theorems 4.4, 4.5). *Let $X_1, \ldots, X_s$ be independent Poisson trials, and let $X = \sum_{i=1}^s X_i$. Then*

$$\forall \delta, \ 0 \le \delta \le 1 : \ \Pr(X \ge (1 + \delta)\mathbf{E}[X]) \le e^{-\delta^2 \mathbf{E}[X]/3}, \tag{1}$$

$$\forall R \ge 6\mathbf{E}[X] : \ \Pr(X \ge R) \le 2^{-R}, \tag{2}$$

$$\forall \delta, \ 0 < \delta < 1 : \ \Pr(X \le (1 - \delta)\mathbf{E}[X]) \le e^{-\delta^2 \mathbf{E}[X]/2}. \tag{3}$$

3 Protocol P_{PL}

We give a loosely-stabilizing leader election protocol P_{PL}. This protocol uses a given upper bound N on n and has a parameter $c \ge 1$ by which we can adjust the expected convergence time and the expected holding time. As mentioned above, time complexity is measured as *parallel time*, which is defined as the number of interactions divided by n. The expected convergence time of P_{PL} is $O(c \log n \cdot \log^2 N) \subseteq O(c \log^3 N)$ and the expected holding time is $\Omega(cn^{10c+1})$. Thus, we achieve loosely-stabilizing leader election with polylogarithmic convergence time and polynomial holding time by setting $c = \Theta(1)$. For example, assigning $c = 10$ gives $O(\log^3 N)$ convergence time and $\Omega(n^{100})$ holding time.

The pseudo code of P_{PL} is given as Algorithm 1. Each agent has five variables: $\mathtt{leader} \in \{\top, \bot\}$, $\mathtt{shield} \in \{\top, \bot\}$, $\mathtt{virus} \in [0, t_{\mathrm{virus}}]$, $\mathtt{timer}_L \in [0, t_{\mathrm{max}}]$, and $\mathtt{timer}_I \in [0, t_{\mathrm{emit}}]$. The first two variables $\mathtt{leader}$ and $\mathtt{shield}$ are Boolean variables: $v.\mathtt{leader} = \top$ means that v is a leader, $v.\mathtt{shield} = \top$ means that v is *shielded*, which will be explained later. The next three variables $\mathtt{virus}$, $\mathtt{timer}_L$, and $\mathtt{timer}_I$ are count-down timers where their maximum values are $t_{\mathrm{virus}} = 60\lceil \ln N \rceil$, $t_{\mathrm{max}} = 12c \cdot t_{\mathrm{virus}}\lceil \ln N \rceil$, and $t_{\mathrm{emit}} = 12c \cdot t_{\mathrm{virus}}\lceil \ln N \rceil$, respectively (Note that $t_{\mathrm{max}} = t_{\mathrm{emit}}$). The output function π_{out} is defined as follows: an agent with $\mathtt{leader} = \top$ (resp. $\mathtt{leader} = \bot$) outputs L (resp. F).

Algorithm 1 P_{PL}.

Constants:

c, N // given parameters. $N \geq n$ is guaranteed.

$t_{\mathrm{virus}} = 60\lceil \ln N \rceil$

$t_{\max} = 12c \cdot t_{\mathrm{virus}} \lceil \ln N \rceil$

$t_{\mathrm{emit}} = 12c \cdot t_{\mathrm{virus}} \lceil \ln N \rceil$

Variables of each agent:

$\mathtt{leader} \in \{\top, \bot\}$, $\mathtt{shield} \in \{\top, \bot\}$,

$\mathtt{timer}_L \in [0, t_{\max}]$, $\mathtt{virus} \in [0, t_{\mathrm{virus}}]$, $\mathtt{timer}_I \in [0, t_{\mathrm{emit}}]$,

Output function π_{out}:

if $v.\mathtt{leader} = \top$ holds, then the output of agent v is L, otherwise F.

Interaction between initiator a_0 and responder a_1:

1: $a_0.\mathtt{timer}_L \leftarrow a_1.\mathtt{timer}_L \leftarrow \max(a_0.\mathtt{timer}_L - 1,\ a_1.\mathtt{timer}_L - 1, 0)$
2: **for** $i \in \{0,1\}$ such that $a_i.\mathtt{timer}_L = 0$ **do** $a_i.\mathtt{leader} \leftarrow \top$ **endfor**
3: **if** $\exists i \in \{0,1\} : a_i.\mathtt{leader} = \top$ **then** $a_0.\mathtt{timer}_L \leftarrow a_1.\mathtt{timer}_L \leftarrow t_{\max}$ **endif**

4: $a_0.\mathtt{virus} \leftarrow a_1.\mathtt{virus} \leftarrow \max(a_0.\mathtt{virus} - 1,\ a_1.\mathtt{virus} - 1, 0)$
5: **for** $i \in \{0,1\}$ such that $\neg a_i.\mathtt{shield} \wedge (a_i.\mathtt{virus} > 0)$ **do** $a_i.\mathtt{leader} \leftarrow \bot$ **endfor**

6: **for** $i \in \{0,1\}$ **do** $a_i.\mathtt{timer}_I \leftarrow \max(a_i.\mathtt{timer}_I - 1, 0)$ **endfor**
7: **if** $a_0.\mathtt{timer}_I = 0 \wedge a_0.\mathtt{leader} = \top$ **then** $(a_0.\mathtt{virus}, a_0.\mathtt{shield}) \leftarrow (t_{\mathrm{virus}}, \top)$ **endif**
8: **if** $a_1.\mathtt{timer}_I = 0 \wedge a_1.\mathtt{leader} = \top$ **then** $a_1.\mathtt{shield} \leftarrow \bot$ **endif**
9: **for** $i \in \{0,1\}$ such that $a_i.\mathtt{timer}_I = 0$ **do** $a_i.\mathtt{timer}_I \leftarrow t_{\mathrm{emit}}$ **endfor**

Protocol P_{PL} consists of a *timeout mechanism* (Lines 1-3) and a *virus-war mechanism* (Lines 4-9). By using variable $\mathtt{timer}_L$, the timeout mechanism creates a leader when no leader exists in the population. By using variables $\mathtt{timer}_I$, $\mathtt{virus}$, and $\mathtt{shield}$, the virus-war mechanism reduces the number of leaders if there are two or more leaders.

The timeout mechanism of P_{PL} (Lines 1-3) is almost the same as that of the protocol given in [14]. This mechanism uses a leader-timer $\mathtt{timer}_L$, which indicates the possibility of existence of a leader. A leader agent always keeps $\mathtt{timer}_L = t_{\max}$, and resets the leader-timer of the other agent to $t_{\max}$ every time it interacts with a non-leader agent (Line 3). We call this operation *timer reset*. When two non-leaders interact, we take the larger timer value of the two agents, decrease it by one, and substitute the decreased value into the leader-timers of both agents (Line 1). We call this operation *larger value propagation*. When the leader-timer of a non-leader decreases to zero, it suspects that no leader exists in the population, and it becomes a new leader (Line 2). We call this event *timeout*. This mechanism works well for the following reasons: (i) the timeout rarely happens when the population has a leader because the leader-timers of all agents have large values thanks to the timer reset and the larger value propagation, (ii) when no leader exists in the population, a timeout happens and a new leader is created within a short time because there is no possibility of a timer reset. It is proven in [14] that, if $t_{\max} = \Omega(n)$, the timeout rarely happens in the population with at least one leader. In the next section, we show that, with a high probability, the timeout does not happen for a long time when a leader exists, even if the maximum value $t_{\max}$ is polylogarithmic in N, specifically, $t_{\max} = \Theta(\log^2 N)$.

The basic idea of the virus-war mechanism is first presented in [15]. P_{PL} uses this idea, but implements it in a considerably different way in order to reduce the number of leaders to

one within a polylogarithmic parallel time. In the virus-war mechanism (Lines 4-9), every leader tries to kill other leaders by using viruses and become the unique leader. We say that agent v has a virus if $v.\texttt{virus} > 0$, and that v is shielded if $v.\texttt{shield} = \top$. As we will see later, a virus is propagated to the whole population by interactions and kills leaders that are not shielded. Every agent has an individual timer $\texttt{timer}_I$ to create a new virus periodically. This timer is decreased by one every time that agent participates in an interaction (Line 6). When the individual timer of a leader reaches zero at an interaction, its fate differs according to its role in the interaction, initiator or responder. If the agent is an initiator, it succeeds in creating a new virus and becomes shielded, that is, $\texttt{virus} \leftarrow t_{\text{virus}}$ and $\texttt{shield} \leftarrow \top$ (Line 7). If it is a responder, it becomes unshielded i.e., $\texttt{shield} \leftarrow \bot$ (Line 8). Thereafter, the individual timer is reset to the maximum value t_{emit} in both cases (Line 9). A virus spreads by interactions (Line 4). A leader is killed and becomes a non-leader if it catches a virus when it is not shielded (Line 5). The value of $\texttt{virus}$ in an agent corresponds to the TTL (time to live) of the virus that the agent carries. Since it decreases in the larger value propagation fashion, viruses eventually disappear from the population and no virus exists in the following execution until a new virus is created by some leader.

When there are multiple leaders, the virus-war mechanism elects exactly one leader within a short time. Since every agent participates in any interaction Γ_t with probability $2/n$, the individual timer of an agent reaches zero within $O(nt_{\text{emit}})$ interactions with high probability. Therefore, some leader creates a new virus (and becomes shielded) by executing Line 7 within $O(nt_{\text{emit}})$ interactions in expectation and approximately half of leaders are killed by the virus. Therefore, the number of shielded leaders is approximately halved for every $O(nt_{\text{emit}})$ interactions. This rough and intuitive analysis explains why $O(nt_{\text{emit}} \log n) \subseteq O(cn \log^3 N)$ interactions are sufficient to elect a unique leader. On the other hand, we must consider the risk of *suicide*, i.e., the event where a single leader creates a new virus and then becomes unshielded before the virus disappears from the population. This event causes the unique leader to be killed by the virus. A long holding time cannot be achieved if suicides are frequent. Suicide of a single leader v occurs only when a leader v first executes Line 7, starting a new virus while shielding itself, and then later executes Line 8 while the virus is still present, causing itself to be killed shortly thereafter. We can ensure that the frequency of that sequence of events is extremely small by assigning a sufficiently large value to t_{emit} compared to t_{virus}, as we discuss carefully in the next section.

4 Analysis of Convergence and Holding Time

In this section, we prove that P_{PL} is a $(O(cn \log n \cdot \log^2 N), \Omega(cn^{10c+1}))$-loosely-stabilizing leader election protocol. Define the set $\mathcal{S}$ of safe configurations as follows:

$$\mathcal{L}_i = \{C \in \mathcal{C}_{\text{all}}(P_{\text{PL}}) \mid 1 \leq |\{v \in V \mid C(v).\texttt{leader}\}| \leq i\}, \text{ for } i = 1, 2, \ldots, n$$
$$\mathcal{G}_{\text{half}} = \{C \in \mathcal{C}_{\text{all}}(P_{\text{PL}}) \mid \forall v \in V : C(v).\texttt{timer}_L \geq t_{\max}/2\},$$
$$\mathcal{L}_{\text{safe}} = \{C \in \mathcal{C}_{\text{all}}(P_{\text{PL}}) \mid \exists v \in V : C(v).\texttt{leader} \wedge C(v).\texttt{shield} \wedge C(v).\texttt{timer}_I \geq t_{\text{emit}}/2\},$$
$$\mathcal{V}_{\text{clean}} = \{C \in \mathcal{C}_{\text{all}}(P_{\text{PL}}) \mid \forall v \in V : C(v).\texttt{virus} = 0\}.$$
$$\mathcal{S} = \mathcal{L}_1 \cap \mathcal{G}_{\text{half}} \cap (\mathcal{L}_{\text{safe}} \cup \mathcal{V}_{\text{clean}}),$$

We need to prove the following two equalities:

$$\max_{C \in \mathcal{C}_{\text{all}}(P_{\text{PL}})} \text{EIC}_{P_{\text{PL}}}(C, \mathcal{S}) = O(cn \log n \cdot \log^2 N), \tag{4}$$

$$\min_{C \in \mathcal{S}} \text{EIH}_{P_{\text{PL}}}(C, LE) = \Omega\left(cn^{10c+1}\right). \tag{5}$$

Note that $C \in \mathcal{L}_i$ requires that the population in configuration C has at least one leader but does not have more than i leaders. For example, $\mathcal{L}_1$ is the set of configuration where exactly one leader exists in the population and $\mathcal{L}_n$ is the set of configurations where at least one leader exists in the population. Clearly, $\mathcal{L}_1 \subset \mathcal{L}_2 \subset \cdots \subset \mathcal{L}_n$ holds.

In the remainder of this section, we first introduce analytic tools (the notions of epidemic and virtual agents) in Section 4.1. Expected holding time and expected convergence time are analyzed in Sections 4.2 and 4.3, respectively. Since we evaluate the expected convergence time and the expected holding time only asymptotically, it suffices to assume that n is sufficiently large. Specifically, we assume $n \geq 55$; we then have $\lceil \ln N \rceil \geq 5$.

4.1   Tools

The goal of this Section is, intuitively, to prove $\Delta = \Theta(\log n)$ where Δ is the integer introduced at the end of Section 1.1. First, we review the notions of *epidemic* and *virtual agents* presented in [5] and [14] respectively. Most of the definitions in Section 4.1 are borrowed from [14].

By an abuse of notation, we will identify an interaction (u, v) with the set $\{u, v\}$ whenever convenient. Let $\gamma = \gamma_0, \gamma_1, \ldots$ be an infinite sequence of interactions and let r be an agent in V. The *epidemic function* $I_{r,\gamma} : [0, \infty) \to 2^V$ is defined as follows: $I_{r,\gamma}(0) = \{r\}$, and for $t = 1, 2, \ldots, I_{r,\gamma}(t) = I_{r,\gamma}(t-1) \cup \gamma_{t-1}$ if $I_{r,\gamma}(t-1) \cap \gamma_{t-1} \neq \emptyset$; otherwise, $I_{r,\gamma}(t) = I_{r,\gamma}(t-1)$. We say that v is *infected* at time t if $v \in I_{r,\gamma}(t)$ in the epidemic starting from agent r under γ. At time 0, only r is infected; at later steps, an agent becomes infected if it interacts with an infected agent. Once an agent becomes infected, it remains infected thereafter. We define the *infection time* $t_{r,\gamma}(v)$ of agent $v \in V$ to be $\min\{t \geq 0 \mid v \in I_{r,\gamma}(t+1)\}$ if $v \neq r$, and we let $t_{r,\gamma}(r) = -1$. Note that every agent $v \neq r$ is infected by the interaction $\gamma_{t_{r,\gamma}(v)}$.

We now define the *virtual agent* $VA_{r,\gamma}(v)$ of each agent $v \in V$. We assume that all agents eventually become infected, that is, $I_{r,\gamma}(t') = V$ holds for some t'. The virtual agent $VA_{r,\gamma}(v)$ is not defined if no such t' exists.[2] Let v be any agent other than r. We define the *parent* of v as the agent that infects v at time $t_{r,\gamma}(v)$. This parent-child relation defines a spanning tree of G rooted at r. In this tree, we call the unique path from r to v i.e., $v_0(= r), v_1, v_2, \ldots, v_k(= v)$, the *infection path* of v. The virtual agent $VA_{r,\gamma}(v)$ is a virtual entity that migrates from r to v through this infection path. At the beginning, $VA_{r,\gamma}(v)$ stays at r. For every $i \in [0, k-1]$, it migrates from v_i to v_{i+1} at time $t_{r,\gamma}(v_{i+1})$ After reaching $v_k = v$, the virtual agent remains at v. For $t \geq 0$, we say that virtual agent $VA_{r,\gamma}(v)$ *participates* in γ_t if, at time t, $VA_{r,\gamma}(v)$ is at one of the two agents participating in γ_t. For any $t \geq 1$, we define $VI_{r,\gamma}(v, t)$ as the number of interactions in $\gamma_0, \gamma_1, \ldots, \gamma_{t-1}$ that $VA_{r,\gamma}(v)$ participates in. Formally, we define $VI_{r,\gamma}(v, t) = |\{j \in [0, t-1] \mid (v \in \gamma_j \wedge j \geq t_{r,\gamma}(v)) \vee (\exists i \in [0, k-1] : v_i \in \gamma_j \wedge t_{r,\gamma}(v_i) \leq j < t_{r,\gamma}(v_{i+1}))\}|$.

Consider an execution $\Xi_P(C_0, \gamma) = C_0, C_1, \ldots$ of some protocol P that has a propagating variable `var`. Larger value propagation guarantees that the virtual agent $VA_{r,\gamma}(v)$ brings a large value of `var` from r to v when it reaches v through the infection path if r has a sufficiently large value of `var` in C_0 and virtual agent $VA_{r,\gamma}(v)$ participates in sufficiently few interactions in $\gamma_0, \gamma_1, \ldots, \gamma_{t_{r,\gamma}(v)}$. This property is formalized as the following trivial lemma.

▶ **Lemma 3.** *Let $\gamma = \gamma_0, \gamma_1, \ldots$ be a schedule, P be a protocol, and* `var` *be a propagating variable of P. Let $C_0 \in \mathcal{C}_{\mathrm{all}}(P)$ and $\Xi_P(C_0, \gamma) = C_0, C_1, \ldots$. Then $v \in I_{r,\gamma}(t) \Rightarrow C_t(v).\mathtt{var} \geq C_0(r).\mathtt{var} - VI_{r,\gamma}(v, t)$ for any agents $r, v \in V$ and any integer $t \geq 1$.*

---

[2] Such t' exits with probability 1 when the uniformly random scheduler Γ is given.

Angluin et al. [5] prove that the epidemic from any agent $r \in V$ finishes (i.e., all agents are infected) within $\Theta(n \log n)$ interactions with high probability. Furthermore, our previous work [14] gives a concrete lower bound on the probability that the epidemic finishes within a given number of interactions as follows:

▶ **Lemma 4** ([14]). *Let $r \in V$. For any integer $t \geq 1$, $\Pr(I_{r,\Gamma}(2t) \neq V) \leq ne^{-t/n}$.*

From Lemmas 3 and 4, we can achieve the goal of this section by giving a sufficiently tight lower bound of $VI_{r,\gamma}(v,t)$ with high probability. It would be easy if the probabilities of $VA_{r,\gamma}(v) \in \Gamma_t$ for distinct t were independent of each other, and equal to $2/n$, like the probability of $v \in \Gamma_t$. Then, the simple Chernoff bound would give a tight lower bound with high probability. However, $\Pr(VA_{r,\gamma}(v) \in \Gamma_t) = 2/n$ does not hold in general. This is because $VA_{r,\gamma}(v)$ at time t is determined not only by the preceding interactions $\Gamma_0, \Gamma_1, \ldots, \Gamma_{t-1}$ but also by the subsequent interactions. $\Gamma_t, \Gamma_{t+1}, \ldots$. Therefore, we need a careful analysis. The following lemma is one of the main contributions in this paper; it enables us to bound $VI_{r,\Gamma}(v,t)$ probabilistically to a logarithmic order when t is logarithmic in n.

▶ **Lemma 5.** *The following two inequalities hold for any $r, v \in V$, $t \geq 1$, $d \geq 3$, $d' \geq 3$, and $d'' \geq 6$:*

$$\Pr\left(VI_{r,\Gamma}(v,t) \geq 4d\lceil \log_2 n \rceil + 2d'H_n + \frac{4d''t}{n} \;\middle|\; I_{r,\Gamma}(t) = V \right) \leq n^{-d} + n^{-d'} + e^{-\frac{2d''t}{n}},$$

$$\Pr\left(VI_{r,\Gamma}(v,t) \geq 4d\lceil \log_2 n \rceil + 2d'H_n + \frac{8t}{n} \;\middle|\; I_{r,\Gamma}(t) = V \right) \leq n^{-d} + n^{-d'} + e^{-\frac{4t}{3n}}.$$

Proof. Assume that $I_{r,\Gamma}(t) = V$. Let $t_1, t_2, \ldots, t_{n-1}$ be the integers such that $I_{r,\Gamma}(t_i) \neq I_{r,\Gamma}(t_i + 1)$ for each $i \in [1, n-1]$ and $0 \leq t_1 < t_2 < \cdots < t_{n-1} < t$. Let $u_1, u_2, \ldots, u_{n-1}$ be the agents where $t_i = t_{r,\Gamma}(u_i)$. Intuitively, u_i is the $(i+1)^{\text{st}}$ agent to be infected and t_i is its infection time in the epidemic starting from r under Γ. We let $u_0 = r$. Note that both u_i and t_i are random variables. Let Z_1 be the number of interactions that virtual agent $VA_{r,\gamma}(v)$ participates in among the $n-1$ interactions $\Gamma_{t_1}, \Gamma_{t_2}, \ldots, \Gamma_{t_{n-1}}$. Let $Z_2 = VI_{r,\Gamma}(v,t) - Z_1$. As mentioned above, the parent-child relation based on the epidemic starting from r gives the infection path from r to v, denoted by $v_0(=r), v_1, v_2, \ldots, v_k = v$. (Note that k and each v_i are random variables.) For any j ($0 \leq j < t$), let X_j be the indicator variable such that $X_j = 1$ holds if $VA_{r,\Gamma}(v)$ participates in Γ_t, otherwise $X_j = 0$. In other words, we define $X_j = 1 \Leftrightarrow (j \geq t_{r,\Gamma}(v) \wedge v \in \Gamma_j) \vee (\exists i : t_{r,\Gamma}(v_i) \leq j < t_{r,\Gamma}(v_{i+1}) \wedge v_i \in \Gamma_j)$. Let $X = \sum_{j \in \{t_1, t_2, \ldots, t_{n-1}\} \setminus \{t_{r,\Gamma}(v_1), t_{r,\Gamma}(v_2), \ldots, t_{r,\Gamma}(v_k)\}} X_j$. It trivially holds that $Z_1 = \sum_{j \in \{t_1, t_2, \ldots, t_{n-1}\}} X_j = k + X$. Hence, $VI_{r,\Gamma}(v,t) = Z_1 + Z_2 = k + X + Z_2$. In the following, we give probabilistic upper bounds on k, X, and Z_2.

First, we focus on k, the length of the infection path from r to v. The index of v_i ($0 \leq i \leq k$), denoted by x_i, is defined to be the integer such that $v_i = u_{x_i - 1}$, i.e., v_i is the x_i^{th} agent to be infected among the population. For example, $v_0 = r$ is the first infected agent and thus $x_0 = 1$. The parent of each u_j is chosen uniformly at random among $u_0, u_1, \ldots, u_{j-1}$. Therefore, $\Pr(x_i \geq 2x_{i-1}) \geq 1/2$ for $i = 1, 2, \ldots, k$. If $k \geq 4d\lceil \log_2 n \rceil$, then the event $x_i \geq 2x_{i-1}$ must not happen more than $\lceil \log_2 n \rceil$ times for $i = 1, 2, \ldots, 4d\lceil \log_2 n \rceil$.

Hence, letting Y be a binomial random variable such that $Y \sim B(4d\lceil \log_2 n \rceil, 1/2)$, we have

$$\Pr(k \geq 4d\lceil \log_2 n \rceil \mid I_{r,\Gamma}(t) = V) \leq \Pr(Y \leq \lceil \log_2 n \rceil)$$

$$= \Pr\left(Y \leq \frac{\mathbf{E}[Y]}{2d}\right)$$

$$\leq \exp\left(-\frac{1}{2} \cdot \left(\frac{2d-1}{2d}\right)^2 \cdot 2d\lceil \log_2 n \rceil\right)$$

$$\leq \exp\left(-d \ln n \cdot \log_2 e \cdot \left(\frac{2d-1}{2d}\right)^2\right)$$

$$\leq n^{-d},$$

where we use (3) in Lemma 2 for the second inequality and $\forall d \geq 3 : \log_2 e \cdot ((2d-1)/2d)^2 \geq 1$ for the last inequality $(\log_2 e \cdot (5/6)^2 = 1.00187\ldots)$.

Next, we focus on the random number $X = \sum_{j \in \{t_1, t_2, \ldots, t_{n-1}\} \setminus \{t_{r,\Gamma}(v_1), t_{r,\Gamma}(v_2), \ldots, t_{r,\Gamma}(v_k)\}} X_j$. For any $j \in \{t_1, t_2, \ldots, t_{n-1}\} \setminus \{t_{r,\Gamma}(v_1), t_{r,\Gamma}(v_2), \ldots, t_{r,\Gamma}(v_k)\}$, we have $\Pr(X_j = 1) = 1/i$, where $j = t_i$. Therefore, letting $W_1, W_2, \ldots, W_{n-1}$ be a sequence of independent Poisson trials with probability $\Pr(W_i = 1) = 1/i$, we have

$$\Pr(X \geq 2d'H_n \mid I_{r,\Gamma}(t) = V) \leq \Pr\left(\sum_{i=1}^{n-1} W_i \geq 2d'H_{n-1}\right) \leq 2^{-2d'H_{n-1}} \leq e^{-d'H_{n-1}} \leq n^{-d'},$$

where we use $d' \geq 3$, $H_{n-1} = \mathbf{E}[\sum_{i=1}^{n-1} W_i]$, (2) in Lemma 2 for the second inequality, and $H_{n-1} > \ln n$ for the last inequality.

Finally, we focus on Z_2, the number of non-infection interactions that virtual agent $VA_{r,\Gamma}(v)$ participates in among $\Gamma_0, \Gamma_1, \ldots, \Gamma_{t-1}$. Under the condition that no agent is newly infected by an interaction Γ_j, the probability that $VA_{r,\Gamma}(v)$ participates in Γ_j is at most $4/n$ [14]. This is because, letting $m = |I_{r,\Gamma}(j)|$ and $_0C_2 = {}_1C_2 = 0$, the probability is exactly $\frac{m-1}{{}_mC_2 + n-m C_2}$, which is at most $4/n$ regardless of j (See Appendix in [14]). Therefore, letting W' be a binomial random variable such that $W' \sim B(t, 4/n)$, we have

$$\Pr\left(Z_2 \geq \frac{4d''t}{n} \;\middle|\; I_{r,\Gamma}(t) = V\right) \leq \Pr\left(W' \geq \frac{4d''t}{n}\right) \leq 2^{-4d''t/n} \leq e^{-2d''t/n},$$

where we use $d'' \geq 6$ and (2) in Lemma 2 for the second inequality. Similarly, we have

$$\Pr\left(Z_2 \geq \frac{8t}{n} \;\middle|\; I_{r,\Gamma}(t) = V\right) \leq \Pr\left(W' \geq \frac{8t}{n}\right) \leq e^{-4t/3n},$$

where we use (1) in Lemma 2 for the second inequality.

The two inequalities of the lemma follow from the above probabilistic upper bounds on k, X, and Z_2. ◀

From Lemma 3, Lemma 4, and Lemma 5, we obtain Lemma 6 below.

▶ **Lemma 6.** *Let P be any protocol and* `var` *be a propagating variable of P. Let $C_0 \in \mathcal{C}_{\mathrm{all}}(P)$ and $\Xi_P(C_0, \Gamma) = C_0, C_1, \ldots$. Then the following inequalities hold for any $r \in V$, $t \geq 1$,*

$d \geq 3$, $d' \geq 3$, and $d'' \geq 6$:

$$\Pr\left(\forall v \in V : C_{2t}(v).\mathtt{var} > C_0(r).\mathtt{var} - 4d\lceil \log_2 n \rceil - 2d'H_n - \frac{8d''t}{n}\right)$$
$$\geq 1 - n\left(n^{-d} + n^{-d'} + e^{-\frac{4d''t}{n}} + e^{-\frac{t}{n}}\right)$$

$$\Pr\left(\forall v \in V : C_{2t}(v).\mathtt{var} > C_0(r).\mathtt{var} - 4d\lceil \log_2 n \rceil - 2d'H_n - \frac{16t}{n}\right)$$
$$\geq 1 - n\left(n^{-d} + n^{-d'} + e^{-\frac{8t}{3n}} + e^{-\frac{t}{n}}\right).$$

Lemma 6 is formalized for general applications: the lemma can be used for any protocol and any of its propagating variables. In this paper, we use the following two corollaries.

▶ **Corollary 7.** *Let $C_0 \in \mathcal{L}_n$ and $\Xi_{P_{\mathrm{PL}}}(C_0, \Gamma) = C_0, C_1, \ldots$. Then:*

$$\Pr\left(C_{6cn\lceil \ln N \rceil^2} \in \mathcal{G}_{\mathrm{half}}\right) \geq 1 - 3n^{-10c}.$$

Proof. There exists an agent $r \in V$ such that $C_0(r).\mathtt{leader}$ holds since $C_0 \in \mathcal{L}_n$. The value of the leader-timer in r may not be $t_{\max}$ in the initial configuration C_0. However, in the first interaction that r participates in, r resets the timers of both of the interacting agents to $t_{\max}$. Therefore, we can assume that $C_0(r).\mathtt{timer}_L = t_{\max}$. Assigning $t = 3cn\lceil \ln N \rceil^2$, $d = d' = 11c$, and $d'' = 6$ to the first inequality of Lemma 6 yields the result, because we have:

$$4d\lceil \log_2 n \rceil + 2d'H_n + \frac{8d''t}{n} \leq 44c\lceil \log_2 n \rceil + 22c(1 + \ln n) + 144c\lceil \ln N \rceil^2$$
$$\leq 44c(1 + \log_2 e\lceil \ln N \rceil) + 22c(1 + \lceil \ln N \rceil) + 144c\lceil \ln N \rceil^2$$
$$\leq 144c\lceil \ln N \rceil^2 + 86c\lceil \ln N \rceil + 66c$$
$$\leq \frac{t_{\max}}{2},$$

where we use $t_{\max} = 720c\lceil \ln N \rceil^2$ in the last inequality, and we have:

$$n\left(n^{-d} + n^{-d'} + e^{-\frac{4d''t}{n}} + e^{-\frac{t}{n}}\right) \leq n(2n^{-11c} + n^{-72c\lceil \ln N \rceil} + N^{-15c}) \leq 3n^{-10c}. \qquad \blacktriangleleft$$

▶ **Corollary 8.** *Let $C_0 \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{PL}})$ and $\Xi_{P_{\mathrm{PL}}}(C_0, \Gamma) = C_0, C_1, \ldots$. Assume $C_0(r).\mathtt{virus} = t_{\mathrm{virus}}$ for some $r \in V$. Then:*

$$\Pr\left(\forall v \in V : C_{4n\lceil \ln N \rceil}(v).\mathtt{virus} > 0\right) \geq 1 - 2/n.$$

Proof. Assigning $t = 2n\lceil \ln N \rceil$ and $d = d' = 3$ to the second inequality of Lemma 6 we have

$$4d\lceil \log_2 n \rceil + 2d'H_n + \frac{16t}{n} \leq 12\lceil \log_2 n \rceil + 6(1 + \lceil \ln N \rceil) + 32\lceil \ln N \rceil$$
$$\leq 12(1 + \log_2 e\lceil \ln N \rceil) + 6(1 + \lceil \ln N \rceil) + 32\lceil \ln N \rceil$$
$$\leq 56\lceil \ln N \rceil + 18$$
$$< t_{\mathrm{virus}},$$

where we use the assumption $\lceil \ln N \rceil \geq 5$ in the last inequality, and we have

$$n\left(n^{-d} + n^{-d'} + e^{-\frac{8t}{3n}} + e^{-\frac{t}{n}}\right) \leq n(2n^{-3} + N^{-5} + N^{-2}) \leq 2/n,$$

where we use the assumption $n \geq 55$ in the last inequality. $\qquad \blacktriangleleft$

4.2 Expected Holding Time

We prove (5) in this section, that is, $\min_{C \in \mathcal{S}} \text{EIH}_{P_{\text{PL}}}(C, LE) = \Omega\left(cn^{10c+1}\right)$. Let t be a positive integer and $\gamma = \gamma_0, \gamma_1, \ldots$ be a schedule. Let $C_0 \in \mathcal{C}_{\text{all}}(P_{\text{PL}})$ and $\Xi_{P_{\text{PL}}}(C_0, \gamma) = C_0, C_1, \ldots$. We define the indicator variable: $L^+_{C_0, \gamma}(t) = \top$ if a timeout occurs and a new leader is created in the prefix $C_0, C_1, \ldots, C_t$ of length t+1 of the execution that is, at least one of the interactions $\gamma(0), \gamma(1), \ldots, \gamma(t-1)$ causes a timeout in $\Xi_{P_{\text{PL}}}(C_0, \gamma)$; Otherwise we let $L^+_{C_0, \gamma}(t) = \bot$. For convenience, we define $L^+_{C_0, \gamma}(0) = \bot$. Similarly, we define the indicator variable $L^0_{C_0, \gamma}(t)$ as follows: $L^0_{C_0, \gamma}(t) = \top$ if and only if the prefix of length $t + 1$ of the execution includes a configuration where no leader exists. We also define $\tau = 144cn\lceil \ln N \rceil^2 = nt_{\text{emit}}/5$. In the rest of this section, we prove that for any configuration $C_0 \in \mathcal{S}$:

$$\Pr\left(\neg L^+_{C_0, \Gamma}(\tau) \wedge \neg L^0_{C_0, \Gamma}(\tau) \wedge C_\tau \in \mathcal{S}\right) \geq 1 - O(n^{-10c}), \tag{6}$$

where $\Xi_{P_{\text{PL}}}(C_0, \Gamma) = C_0, C_1, \ldots$. From this inequality, we have $\min_{C \in \mathcal{S}} \text{EIH}_{P_{\text{PL}}}(C, LE) \geq (1 - O(n^{-10c})) \cdot (\tau + \min_{C \in \mathcal{S}} \text{EIH}_{P_{\text{PL}}}(C, LE))$, which yields (5), i.e., $\min_{C \in \mathcal{S}} \text{EIH}_{P_{\text{PL}}}(C, LE) = \Omega(cn^{10c+1})$.

Let v be an agent, $t \geq 1$ an integer, and $\gamma = \gamma_0, \gamma_1, \ldots$ a schedule. We denote by $RI_\gamma(v, t)$ the number of interactions in which v participates among the first t interactions of γ (i.e., $\gamma_0, \gamma_1, \ldots, \gamma_{t-1}$). We also define (asynchronous) *round time*. The first round time $\text{RT}_\gamma(1)$ is the minimum t satisfying $\forall v \in V, \exists j \in [0, t-1]: v \in \gamma_j$. For any $i \geq 2$, we define the i^{th} round time $\text{RT}_\gamma(i)$ as the minimum t satisfying $\forall v \in V, \exists j \in [\text{RT}_\gamma(i-1), t-1]: v \in \gamma_j$. For completeness, we define $\text{RT}_\gamma(0) = 0$. Note that, for any $i \geq 1$, every agent $v \in V$ participates in at least one interaction in $\gamma_{\text{RT}_\gamma(i-1)}, \ldots, \gamma_{\text{RT}_\gamma(i)-1}$.

▶ **Lemma 9.** $\Pr(\max_{v \in V} RI_\Gamma(v, \tau) \geq t_{\text{emit}}/2) < N^{-30c+1}$.

Proof. Since each agent v participates in Γ_t with probability $2/n$ for any $t \geq 0$, we have $\mathbf{E}[RI_\Gamma(v, \tau)] = 2\tau/n = \frac{2}{5}t_{\text{emit}}$. Therefore, $\Pr(RI_\Gamma(v, \tau) \geq \frac{1}{2}t_{\text{emit}}) = \Pr(RI_\Gamma(v, \tau) \geq \frac{5}{4}\mathbf{E}[RI_\Gamma(v, \tau)]) \leq \exp(-\mathbf{E}[RI_\Gamma(v, \tau)]/48) = \exp(-6c\lceil \ln N \rceil^2) \leq N^{-30c}$ by the Chernoff bound (1) in Lemma 2 with $\delta = 1/4$, and by the assumption $\lceil \ln N \rceil \geq 5$. Hence, the lemma holds by the union bounds. ◀

To prove (6), we first give lower bounds on the probabilities of $\neg L^+_{C_0, \Gamma}(\tau)$ and $\neg L^0_{C_0, \Gamma}(\tau)$ in Lemmas 10 and 11, respectively. Then, we give lower bounds on the probability of $C_\tau \in \mathcal{S}$ by Lemmas 14 and 15.

▶ **Lemma 10.** *Let* $C_0 \in \mathcal{G}_{\text{half}}$. *Then,* $\Pr(L^+_{C_0, \Gamma}(\tau)) \leq N^{-30c+1}$.

Proof. Since $t_{\text{max}} = t_{\text{emit}}$, and the leader-timer of every agent is no less than $t_{\text{max}}/2$ in $C_0 \in \mathcal{G}_{\text{half}}$, a timeout happens within the first τ interactions with probability at most N^{-30c+1}, by Lemma 9. ◀

▶ **Lemma 11.** *Let* $C_0 \in \mathcal{L}_n \cap (\mathcal{L}_{\text{safe}} \cup \mathcal{V}_{\text{clean}})$. *Then,* $\Pr(L^0_{C_0, \Gamma}(\tau)) \leq N^{-30c+1}$.

Proof. First, consider the case $C_0 \in \mathcal{L}_{\text{safe}}$. The population has a shielded leader v_l whose individual timer is at least $t_{\text{emit}}/2$ at C_0. Therefore, $L^0_{C_0, \Gamma}(\tau)$ holds only if v_l participates in $t_{\text{emit}}/2$ or more interactions and becomes unshielded within the first τ interactions. Next, consider the case $C_0 \in \mathcal{L}_n \cap \mathcal{V}_{\text{clean}}$. One or more leaders exist, but no virus exists in the population in C_0. The leaders are never killed (become non-leaders) until viruses appear in the population. Therefore, $L^0_{C_0, \Gamma}(\tau)$ holds only if some leader, say v_l, creates a new virus.

However, v_l resets its individual timer to t_{emit} and becomes shielded when it creates a new virus. Therefore, $L^0_{C_0,\Gamma}(\tau)$ holds only if v_l participates in t_{emit} or more interactions among the first τ interactions. Thus, in both cases, $L^0_{C_0,\Gamma}(\tau)$ holds only if some agent participates in $t_{\mathrm{emit}}/2$ or more interactions among the first τ interactions. By Lemma 9, this necessary condition holds with probability at most N^{-30c+1}, which yields the lemma. ◄

▶ **Lemma 12.** $\Pr(\mathrm{RT}_\Gamma(i) \geq 2in(1 + \lceil \ln n \rceil)) \leq ne^{-i/4}$ *for any* $i \geq 1$.

Proof. Each round finishes when every agent $v \in V$ has interacted during that round. Consider the case that s ($s \geq 1$) agents have not yet interacted in round j. One of these s agents participates in the next interaction with probability $\frac{{}_sC_2 + s(n-s)}{{}_nC_2} \geq \frac{s}{n}$. Let $X_{1,s}, X_{2,s}, \ldots, X_{i,s}$ be independent random variables each of which corresponds to the number of trials needed to reach the first success where the success probability of each trial is s/n. We have

$$
\Pr(\mathrm{RT}_\Gamma(i) \geq 2in(1 + \lceil \ln n \rceil)) \leq \Pr\left(\sum_{j=1}^i \sum_{s=1}^n X_{j,s} \geq 2in(1 + \lceil \ln n \rceil)\right)
$$

$$
\leq \Pr\left(\sum_{s=1}^n \sum_{j=1}^i X_{j,s} \geq 2inH_n\right)
$$

$$
\leq \Pr\left(\sum_{s=1}^n \sum_{j=1}^i X_{j,s} \geq \sum_{s=1}^n \frac{2in}{s}\right) \leq \sum_{s=1}^n \Pr\left(\sum_{j=1}^i X_{j,s} \geq \frac{2in}{s}\right).
$$

For a binomial random variable $Y_s \sim B(\lceil \frac{2in}{s} \rceil, \frac{s}{n})$, we have $\Pr(\sum_{j=1}^i X_{j,s} \geq \frac{2in}{s}) = \Pr(\sum_{j=1}^i X_{j,s} \geq \lceil \frac{2in}{s} \rceil) \leq \Pr(Y_s \leq i)$. Hence

$$
\Pr\left(\sum_{j=1}^i X_{j,s} \geq \frac{2in}{s}\right) \leq \Pr(Y_s \leq i) \leq \Pr\left(Y_s \leq \frac{1}{2} \cdot \mathbf{E}[Y_s]\right) \leq e^{-\mathbf{E}[Y_s]/8} \leq e^{-i/4},
$$

where we use Chernoff bound given as (3) in Lemma 2 for the third inequality. ◄

▶ **Corollary 13.** $\Pr(\mathrm{RT}_\Gamma(t_{\mathrm{virus}}) \geq \tau) \leq N^{-15c+1}$.

Proof. Since we assume $\lceil \ln N \rceil \geq 5$, Lemma 12 yields $\Pr(\mathrm{RT}_\Gamma(t_{\mathrm{virus}}) \geq \tau) = \Pr(\mathrm{RT}_\Gamma(t_{\mathrm{virus}}) \geq 144cn\lceil \ln N \rceil^2) \leq \Pr(\mathrm{RT}_\Gamma(t_{\mathrm{virus}}) \geq 120cn\lceil \ln N \rceil(1 + \lceil \ln n \rceil)) \leq \Pr(\mathrm{RT}_\Gamma(c \cdot t_{\mathrm{virus}}) \geq 2(c \cdot t_{\mathrm{virus}})n(1 + \lceil \ln n \rceil)) \leq ne^{-c \cdot t_{\mathrm{virus}}/4} \leq N^{-15c+1}$. ◄

▶ **Lemma 14.** *Let* $C_0 \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{PL}})$ *and* $\Xi_{P_{\mathrm{PL}}}(C_0, \Gamma) = C_0, C_1, \ldots$ *. Then,* $\Pr(C_\tau \notin \mathcal{L}_{\mathrm{safe}} \cup \mathcal{V}_{\mathrm{clean}}) \leq 2N^{-15c+1}$.

Proof. Consider the case that every agent has fewer than $t_{\mathrm{emit}}/2$ interactions during the first τ interactions, and the $t_{\mathrm{virus}}^{\mathrm{th}}$ round finishes during the first τ interactions. Thanks to the latter condition, viruses disappear from the population and no virus exists in C_τ, i.e., $C_\tau \in \mathcal{V}_{\mathrm{clean}}$, if no leader creates a new virus within the first τ interactions. If some leader $v_l \in V$ creates a new virus during that period, then $v_l.\mathtt{leader} \wedge v_l.\mathtt{shield} \wedge v_l.\mathtt{timer}_I \geq t_{\mathrm{emit}}/2$ in C_τ, thanks to the former condition, which implies $C_\tau \in \mathcal{L}_{\mathrm{safe}}$. This is because v_l resets its individual timer to t_{emit} and becomes shielded at the time it creates a new virus. The probability that both conditions hold is at least $1 - N^{-30c+1} + N^{-15c+1} > 1 - 2N^{-15c+1}$ by Lemma 9 and Corollary 13. ◄

▶ **Lemma 15.** *Let $C_0 \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{PL}})$ and $\Xi_{P_{\mathrm{PL}}}(C_0, \gamma) = C_0, C_1, \dots$. Then $\Pr(C_\tau \notin \mathcal{G}_{\mathrm{half}} \mid C_{\tau - 6cn\lceil \ln N \rceil^2} \in \mathcal{L}_n) \le 3n^{-10c}$.*

Proof. Immediate from Corollary 7.   ◀

The inequality (6) follows from Lemma 10, Lemma 11, Lemma 14, and Lemma 15. Thus, we obtain the following lemma from the discussion in the beginning of this sub-section.

▶ **Lemma 16.** $\min_{C \in \mathcal{S}} \mathrm{EIH}_{P_{\mathrm{PL}}}(C, LE) = \Omega\left(cn^{10c+1}\right)$.

4.3   Expected Convergence Time

We prove (4) in this section, that is, $\max_{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{PL}})} \mathrm{EIC}_{P_{\mathrm{PL}}}(C, \mathcal{S}) = O(cn \log n \cdot \log^2 N)$. Recall that $\mathcal{S} = \mathcal{L}_1 \cap \mathcal{G}_{\mathrm{half}} \cap (\mathcal{L}_{\mathrm{safe}} \cup \mathcal{V}_{\mathrm{clean}})$. We make use of the fact that $\tau = 144cn\lceil \ln N \rceil^2$.

▶ **Lemma 17.** *Let $C_0 \in \mathcal{L}_{\mathrm{safe}} \cup \mathcal{V}_{\mathrm{clean}}$ and $\Xi_{P_{\mathrm{PL}}}(C_0, \Gamma) = C_0, C_1, \dots$. Then $\Pr(\exists j \in [0, 12\tau\lceil \ln N \rceil] : C_j \in \mathcal{L}_n \cap (\mathcal{L}_{\mathrm{safe}} \cup \mathcal{V}_{\mathrm{clean}})) \ge 1 - O(N^{-1})$.*

Proof. Since $C_0 \in \mathcal{L}_{\mathrm{safe}} \cup \mathcal{V}_{\mathrm{clean}}$, we have $C_0 \in \mathcal{L}_n \cap (\mathcal{L}_{\mathrm{safe}} \cup \mathcal{V}_{\mathrm{clean}})$ if $C_0 \in \mathcal{L}_n$. Hence, it suffices to consider only the case that $C_0 \notin \mathcal{L}_n$. In this case, $C_0 \in \mathcal{V}_{\mathrm{clean}}$ also holds because $C_0 \in \mathcal{L}_{\mathrm{safe}} \cup \mathcal{V}_{\mathrm{clean}}$ and $\mathcal{L}_{\mathrm{safe}} \subset \mathcal{L}_n$. To conclude, we need only consider the case that there exists no virus and no leader in the population at C_0. Staring from such a configuration, the population reaches a configuration in $\mathcal{L}_n \cap (\mathcal{L}_{\mathrm{safe}} \cup \mathcal{V}_{\mathrm{clean}})$ immediately after a new leader is created by the timeout. While no leader exists in the population, $\max_{v \in V} v.\mathtt{timer}_L$ is monotonically non-increasing and decreases at least by one during each asynchronous round. Therefore, the timeout occurs within $t_{\max}$ rounds. Lemma 12 guarantees that $\Pr(\mathrm{RT}_\Gamma(t_{\max}) \ge 2nt_{\max}(1 + \lceil \ln n \rceil) \le ne^{-t_{\max}/4} = O(N^{-1})$. Since $\lceil \ln N \rceil \ge 5$, we have $2nt_{\max}(1 + \lceil \ln n \rceil) = 10\tau(1 + \lceil \ln n \rceil) \le 12\tau\lceil \ln N \rceil$, which yields the lemma.   ◀

▶ **Lemma 18.** *Let $C_0 \in \mathcal{L}_n \cap (\mathcal{L}_{\mathrm{safe}} \cup \mathcal{V}_{\mathrm{clean}})$ and $\Xi_{P_{\mathrm{PL}}}(C_0, \Gamma) = C_0, C_1, \dots$. Then $\Pr(C_\tau \in \mathcal{L}_n \cap \mathcal{G}_{\mathrm{half}} \cap (\mathcal{L}_{\mathrm{safe}} \cup \mathcal{V}_{\mathrm{clean}})) \ge 1 - O(N^{-1})$.*

Proof. Immediate from Lemma 11, Lemma 14, and Lemma 15.   ◀

▶ **Lemma 19.** $\Pr(\min_{v \in V} RI_\Gamma(v, 5\tau) \le t_{\mathrm{emit}}) < N^{-900c+1}$.

Proof. Each agent participates in Γ_t with probability $\frac{2}{n}$, for any $t \ge 0$. Therefore, the Chernoff bound given as (3) in Lemma 2 with $\delta = 1/2$ yields $\Pr(RI_\Gamma(v, 5\tau) \le t_{\mathrm{emit}}) \le e^{t_{\mathrm{emit}}/4} < N^{-900c}$. The lemma holds by the union bounds.   ◀

Intuitively, the following lemma guarantees that the number of leaders decreases at least by half during every 15τ interactions with probability close to $1/4$, and never increases with probability $1 - O(N^{-1})$ after the population enters a configuration in $C_0 \in \mathcal{L}_n \cap \mathcal{G}_{\mathrm{half}} \cap (\mathcal{L}_{\mathrm{safe}} \cup \mathcal{V}_{\mathrm{clean}})$.

▶ **Lemma 20.** *Let $i \in [0, n - 1]$. Let $C_0 \in \mathcal{L}_{1+i} \cap \mathcal{G}_{\mathrm{half}} \cap (\mathcal{L}_{\mathrm{safe}} \cup \mathcal{V}_{\mathrm{clean}})$ and $\Xi_{P_{\mathrm{PL}}}(C_0, \Gamma) = C_0, C_1, \dots$. The following inequalities hold:*

$$\Pr(C_{15\tau} \in \mathcal{L}_{1 + \lfloor i/2 \rfloor}) \ge 1/4 - O(n^{-1}) \tag{7}$$

$$\Pr(C_{15\tau} \in \mathcal{L}_{1+i} \cap \mathcal{G}_{\mathrm{half}} \cap (\mathcal{L}_{\mathrm{safe}} \cup \mathcal{V}_{\mathrm{clean}})) \ge 1 - O(N^{-1}). \tag{8}$$

Proof. In this proof, we use notation "with high probability" to represent "with probability $1 - O(n^{-1})$". By repeated application of Lemma 10, Lemma 11, Lemma 14, and Lemma 15, it holds with high probability that no new leader is created and at least one leader always exists in the population in $C_0, C_1, \ldots, C_{15\tau}$. Moreover, Lemma 19 guarantees that the individual timer of every agent reaches zero in every 5τ interactions with high probability. In particular, for all $v \in V$, $v.\mathtt{timer}_I$ reaches zero at least once in the middle period $C_{5\tau}, C_{5\tau+1}, \ldots, C_{10\tau}$ with high probability. Let V_L be the set of all leaders in C_0 and $i' = |V_L| - 1$. Note that $0 \leq i' \leq i$ since $C_0 \in \mathcal{L}_{1+i}$. Since $\neg L^0_{C_0,\Gamma}(15\tau)$ with high probability, some $v_l \in V_L$ creates a new virus with probability at least $1/2 - O(n^{-1})$. This is because, when the individual timer of a leader reaches zero in an interaction, that leader creates a new virus with probability $1/2$ (i.e., if it is the initiator of the current interaction). The virus propagates to the whole population within $4n \log n$ ($< 5\tau$) interactions thereafter with high probability by Corollary 8. At this time, with probability at least $1/2$, independently of configuration C_0, no less than $\lceil i'/2 \rceil$ agents in $V_L \setminus \{v_l\}$ are unshielded. This is because the individual timer of every $v \in V_L \setminus \{v_l\}$ reaches zero before the time, and then becomes unshielded with probability $1/2$. We make the margin period $C_0, C_1, \ldots, C_{5\tau}$ for this reason. To conclude, no fewer than $\lceil i'/2 \rceil$ leaders in $V_L \setminus \{v_l\}$ are killed and at most $1 + \lfloor i'/2 \rfloor \leq 1 + \lfloor i/2 \rfloor$ leaders survive (i.e., remain leaders) in $C_{15\tau}$, with probability $1/4 - O(n^{-1})$. Therefore, we obtain (7) since $\neg L^+_{C_0,\Gamma}(15\tau) \wedge \neg L^0_{C_0,\Gamma}(15\tau)$ holds with high probability. Inequality (8) is guaranteed simply by applying Lemma 10, Lemma 11, Lemma 14, and Lemma 15 repeatedly (fifteen times). ◄

The following corollary follows from Lemma 20.

▶ **Corollary 21.** *Let $C_0 \in \mathcal{L}_n \cap \mathcal{G}_{\mathrm{half}} \cap (\mathcal{L}_{\mathrm{safe}} \cup \mathcal{V}_{\mathrm{clean}})$ and $\Xi_{P_{\mathrm{PL}}}(C_0, \Gamma) = C_0, C_1, \ldots$. Then, there exists some integer $w = O(\tau \log n)$ such that $\Pr(C_w \in \mathcal{S}) \geq 1 - O\left(\frac{\log n}{N}\right)$.*

▶ **Lemma 22.** $\max_{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{PL}})} \mathrm{EIC}_{P_{\mathrm{PL}}}(C, \mathcal{S}) = O\left(cn \log n \cdot \log^2 N\right)$ *holds.*

Proof. By Lemma 14, Lemma 17, Lemma 18, and Corollary 21, we have the following inequality:

$$\max_{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{PL}})} \mathrm{EIC}_{P_{\mathrm{PL}}}(C, \mathcal{S}) \leq O(\tau \log n) + O\left(\frac{\log n}{N}\right) \cdot \max_{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{PL}})} \mathrm{EIC}_{P_{\mathrm{PL}}}(C, \mathcal{S}).$$

Solving this inequality yields $\max_{C \in \mathcal{C}_{\mathrm{all}} P_{\mathrm{PL}}} \mathrm{EIC}_{P_{\mathrm{PL}}}(C, \mathcal{S}) = O(\tau \log n) = O(cn \log n \cdot \log^2 N)$.
◄

▶ **Theorem 23.** *Protocol P_{PL} is an $(O(cn \log n \cdot \log^2 N), \Omega(cn^{10c+1}))$- loosely-stabilizing leader election protocol.*

Thus, in terms of parallel time, P_{PL} is a loosely-stabilizing leader election algorithm with polylogarithmic convergence time ($O(c \log n \cdot \log^2 N) \subseteq O(c \log^3 N)$) and arbitrarily large polynomial holding time ($\Omega(cn^{10c})$).

5 Conclusion

We have presented a loosely-stabilizing leader election protocol with polylogarithmic convergence time. Given an upper bound N of n and a parameter c, our protocol elects a unique leader in the population within $O(c \log^3 N)$ parallel time starting from any configuration, and keeps the unique leader for $\Omega(cn^{10c})$ parallel time.

─── **References** ───

1 Dan Alistarh and Rati Gelashvili. Polylogarithmic-time leader election in population protocols. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming*, pages 479–491. Springer, 2015.

2 Dan Alistarh, Rati Gelashvili, and Milan Vojnović. Fast and exact majority in population protocols. In *the 34th ACM Symposium on Principles of Distributed Computing*, pages 47–56, 2015.

3 D. Angluin, J Aspnes, Z. Diamadi, M.J. Fischer, and R. Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, 18(4):235–253, 2006.

4 D. Angluin, J. Aspnes, M. J Fischer, and H. Jiang. Self-stabilizing population protocols. *ACM Transactions on Autonomous and Adaptive Systems*, 3(4):13, 2008.

5 Dana Angluin, James Aspnes, and David Eisenstat. Fast computation by population protocols with a leader. *Distributed Computing*, 21(3):183–199, 2008.

6 J. Beauquier, P. Blanchard, and J. Burman. Self-stabilizing leader election in population protocols over arbitrary communication graphs. In *International Conference on Principles of Distributed Systems*, pages 38–52, 2013.

7 S. Cai, T. Izumi, and K. Wada. How to prove impossibility under global fairness: On space complexity of self-stabilizing leader election on a population protocol model. *Theory of Computing Systems*, 50(3):433–445, 2012.

8 D. Canepa and M. G. Potop-Butucaru. Stabilizing leader election in population protocols, 2007. http://hal.inria.fr/inria-00166632.

9 M. J. Fischer and H. Jiang. Self-stabilizing Leader Election in Networks of Finite-State Anonymous Agents. In *International Conference on Principles of Distributed Systems*, pages 395–409, 2006. `doi:10.1007/11945529_28`.

10 Leszek Gąsieniec and Grzegorz Staehowiak. Fast space optimal leader election in population protocols. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2653–2667. SIAM, 2018.

11 T. Izumi. On Space and Time Complexity of Loosely-Stabilizing Leader Election. In *International Colloquium on Structural Information and Communication Complexity*, pages 299–312, 2015.

12 M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis.* Cambridge University Press, 2005.

13 R. Mizoguchi, H. Ono, S. Kijima, and M. Yamashita. On space complexity of self-stabilizing leader election in mediated population protocol. *Distributed Computing*, 25(6):451–460, 2012.

14 Y. Sudo, J. Nakamura, Y. Yamauchi, F. Ooshita, H. Kakugawa, and T. Masuzawa. Loosely-stabilizing leader election in a population protocol model. *Theoretical Computer Science*, 444:100–112, 2012.

15 Y. Sudo, F. Ooshita, H. Kakugawa, and T. Masuzawa. Loosely-Stabilizing Leader Election on Arbitrary Graphs in Population Protocols. In *International Conference on Principles of Distributed Systems*, pages 339–354, 2014.

16 Y. Sudo, F. Ooshita, H. Kakugawa, and T. Masuzawa. Loosely-stabilizing Leader Election on Arbitrary Graphs in Population Protocols without Identifiers nor Random Numbers. In *International Conference on Principles of Distributed Systems*, 2015.

17 Yuichi Sudo, Toshimitsu Masuzawa, Ajoy K Datta, and Lawrence L Larmore. The Same Speed Timer in Population Protocols. In *the 36th IEEE International Conference on Distributed Computing Systems*, pages 252–261, 2016.

18 X. Xu, Y. Yamauchi, S. Kijima, and M. Yamashita. Space Complexity of Self-Stabilizing Leader Election in Population Protocol Based on k-Interaction. In *Symposium on Self-Stabilizing Systems*, pages 86–97, 2013.