Loosely-stabilizing Leader Election on Arbitrary Graphs in Population Protocols

Yuichi Sudo[1,2], Fukuhito Ooshita[2],
Hirotsugu Kakugawa[2], and Toshimitsu Masuzawa[2]

[1] NTT Secure Platform Laboratories
3-9-11 Midori-cho, Musashino, Tokyo, 180-8585, Japan
sudo.yuichi@lab.ntt.co.jp
[2] Graduate School of Information Science and Technology, Osaka University
1-5 Yamadaoka, Suita, Osaka, 565-0871, Japan
{y-sudou, f-oosita, kakugawa, masuzawa}@ist.osaka-u.ac.jp

Abstract. In the population protocol model Angluin et al. proposed in 2004, there exists no self-stabilizing protocol that solves leader election on complete graphs without knowing the exact number of nodes. To circumvent the impossibility, we previously introduced the concept of *loose-stabilization*, which relaxes the closure requirement of self-stabilization. A loosely-stabilizing protocol guarantees that starting from any initial configuration a system reaches a loosely-safe configuration, and after that, the system keeps its specification (e.g. the unique leader) not forever, but for a sufficiently long time. Our previous work presented a loosely-stabilizing protocol that solves the leader election on complete graphs using only the upper bound N of n, not the exact value of n. We take this work one step further in this paper: We propose two loosely-stabilizing protocols that solve leader election for *arbitrary graphs*. One is a deterministic protocol that uses the identifiers of nodes while the other is a probabilistic protocol that works on anonymous networks. Given the upper bounds N and Δ of the number of nodes and the maximum degree of nodes respectively, both protocols keep a unique leader for $\Omega(Ne^N)$ expected steps after entering a loosely-safe configuration. The former enters a loosely-safe configuration within $O(m\Delta N \log n)$ expected steps while the latter does within $O(m\Delta^2 N^3 \log N)$ expected steps where m is the number of edges of the graph.

Keywords: Loose-stabilization, Population protocols, Leader election

1 Introduction

The *population protocol* (PP) model, which was presented by Angluin et al.[1], represents wireless sensor networks of mobile sensing devices that cannot control their movement. Two devices (say *agents*) communicate with each other only when they come sufficiently close to each other (we call this event an *interaction*). One example represented by this model is a flock of birds where each bird is

equipped with a sensing device with a small transmission range; each device can communicate with another device only when the corresponding birds come sufficiently close to each other. This unique but meaningful model has attracted broad attention, and there have been numerous studies involving it.

Self-stabilizing leader election (SS-LE) requires that starting from any configuration, a system (say *population*) reaches a safe-configuration in which a unique leader is elected, and after that, the population has the unique leader forever. Self-stabilizing leader election is important in the PP model because (i) many population protocols in the literature work on the assumption that a unique leader exists [1–3], and (ii) self-stabilization tolerates any finite number of transient faults and this property suits systems consisting of numerous cheap and unreliable nodes. (Such systems are the original motivation of the PP model.) However, there exists strict impossibility of SS-LE in the PP model: no protocol can solve SS-LE for complete graphs, arbitrary graphs, trees, lines, degree-bounded graphs and so on unless the exact size of the graph (the number of agents n) is available [3]. This impossibility holds even if we strengthen the PP model by assigning unique identifies to agents, allowing agents to use random numbers, introducing memory of communication links (mediated population protocols [10]), or allowing more than two agents (k agents) to interact at the same time (the PP_k model [5]).

Accordingly, many studies of SS-LE took either one of the following two approaches. One approach is to accept the assumption that the exact value of n is available and focus on the space complexity of the protocol. Cai et al. [6] proved that n states of each agent is necessary and sufficient to solve SS-LE for a complete graph of n agents. Mizoguchi et al.[12] and Xu et al.[14] improved the space-complexity by adopting the mediated population protocol model and the PP_k model respectively. The other approach is to use *oracles*, a kind of failure detectors. Fischer and Jiang [8] took this approach for the first time. They introduced oracle $\Omega?$ that informs all agents whether at least one leader exists or not and proposed two protocols that solve SS-LE for rings and complete graphs by using $\Omega?$. Beauquier et al.[4] presented an SS-LE protocol for arbitrary graphs that uses two copies of $\Omega?$. Canepa et al.[7] proposed two SS-LE protocols that use $\Omega?$ and consume only 1 bit of each agent: one is a deterministic protocol for trees and the other is a probabilistic protocol for arbitrary graphs although the position of the leader is not static and moves among the agents.

Our previous work [13] took another approach to solve SS-LE. We introduced the concept of loose-stabilization, which relaxes the closure requirement of self-stabilization: we allow protocols to deviate from the specification after following it for a sufficiently long time. Concretely, starting from any initial configuration, the population must reach a loosely-safe configuration within a relatively short time; after that, the specification of the problem (the unique leader) must be kept for a sufficiently long time, though not forever. We then proposed a loosely-stabilizing protocol that solves leader election on complete graphs using only an upper bound N of n, not using the exact value of n. Starting from any configuration, the protocol enters a loosely-safe configuration within $O(nN \log n)$

expected steps. After that, the unique leader is kept for $\Omega(Ne^N)$ expected steps. Since the specification is kept for an exponentially long time, we can say this loosely-stabilizing protocol is practically equivalent to a self-stabilizing leader election protocol. Furthermore, this protocol works on any complete graph whose size is no more than N while protocols using the exact value of n work only on the complete graph of size n.

Some works on population protocols assume the probabilistic distribution regarding the interactions of agents: any interaction occurs uniformly at random [1, 2, 13]. This assumption have been used partly 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.

1.1 Our Contribution

In this paper, we consider loosely-stabilizing leader election for *arbitrary undirected graphs*. We adopt two settings: the population with agent-identifiers as in [9] [3] and the population in which agents can use random numbers for state-transition as in [7]. As mentioned above, no self-stabilizing protocol can solve SS-LE for arbitrary graphs, even in these settings, unless the exact value of n is available. For each setting, we propose two protocols P_{ID} and P_{RD} respectively. Given upper bounds N of n and Δ of the maximum degree of nodes, both protocols keep the unique leader for $\Omega(Ne^N)$ expected steps after entering a loosely-safe configuration. Protocol P_{ID} enters a loosely-safe configuration within $O(mN\Delta \log n)$ expected steps while P_{RD} does within $O(mN^3\Delta^2 \log N)$ expected steps where m is the number of edges of the graph. Both protocols consume only $O(\log N)$ bits of each agent's memory. We can say this space complexity is small because even space optimal self-stabilizing protocols that use exact value of n consume $O(\log n)$ bits of each agent[6, 12]. For simplicity, our protocols are presented for undirected graphs. However, they work on *directed* graphs with slight modification which is discussed in the conclusion.

Angluin et al.[1] proves that for any population protocol P working on complete graphs, there exists a protocol that simulates P on any arbitrary graph. However, this simulation can be achieved assuming that all the agents have the common initial states at the start of the execution. Since we cannot assume the specific initial states (This is the essence of self-stabilization), we cannot translate our previous loosely-stabilizing algorithm[13] for complete graphs to a loosely-stabilizing algorithm that works for arbitrary graphs.

[3] Strictly speaking, our model with identifiers is stronger than the model in [9]. We use identifiers to compare their values while Guerraoui et al.[9] only allow equality-test of identifiers and prohibited any other calculation of identifiers such as value-comparing.

2 Preliminaries

This section defines the model we consider for this paper. The model includes both agent-identifiers and random numbers while protocols P_{ID} and P_{RD} use only one of them. In what follows, we denote set $\{z \in \mathbb{N} \mid x \leq z \leq y\}$ by $[x, y]$.

A *population* is a simple and weakly-connected directed graph $G(V, E, \mathrm{id})$ where V ($|V| \geq 2$) is a set of *agents*, $E \subseteq V \times V$ is a set of directed edges and *id* defines unique identifiers of agents. Each edge represents a possible *interactions* (or communication between two agents): If $(u, v) \in E$, agents u and v can interact with each other where u serves as an *initiator* and v serves as a *responder*. Each agent v has the unique identifier $\mathrm{id}(v) \in I$ ($I = [0, \mathrm{id}_{\max}]$, $\mathrm{id}_{\max} \in O(n^c)$ for constant c). We say that G is undirected if it satisfies $(u, v) \in E \Leftrightarrow (v, u) \in E$. We define $n = |V|$ and $m = |E|$.

A *protocol* $P(Q, Y, I, R, T, O)$ consists of a finite set Q of states, a finite set Y of output symbols, a set of possible identifiers I, a range of random numbers $R \subset \mathbb{N}$, transition function $T : (Q \times I) \times (Q \times I) \times R \to Q \times Q$, and output function $O : (Q \times I) \to Y$. When an interaction between two agents occurs, T determines the next states of the two agents based on the current states of the agents, identifiers of the two agents, and a random number $r \in R$ generated at each interaction. The *output of an agent* is determined by O: the output of agent v with state $q \in Q$ is $O(q, \mathrm{id}(v))$. We assume that the set of possible identifiers I is a given parameter and not subject to protocol design.

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 $\mathcal{C}_{\mathrm{all}}(P)$. We say that configuration C changes to C' by interaction $e = (u, v)$ and integer $r \in R$, denoted by $C \overset{e,r}{\to} C'$, if we have $(C'(u), C'(v)) = T(C(u), \mathrm{id}(u), C(v), \mathrm{id}(v), r)$ and $C'(w) = C(w)$ for all $w \in V \setminus \{u, v\}$. A scheduler determines which interaction occurs at each time. In this paper, we consider a uniformly random scheduler $\Gamma = \Gamma_0, \Gamma_1, \ldots$: each $\Gamma_t \in E$ is a random variable such that $\Pr(\Gamma_t = (u, v)) = 1/m$ for any $t \geq 0$ and any $(u, v) \in E$. We also define the random number sequence as $\Lambda = R_1, R_2, \ldots$: each number $R_t \in R$ is a random variable such that $\Pr(R_t = r) = 1/|R|$ for any $t \geq 0$ and $r \in R$. Given an initial configuration C_0, Γ, and Λ, the *execution* of protocol P is defined as $\Xi_P(C_0, \Gamma, \Lambda) = C_0, C_1, \ldots$ such that $C_t \overset{\Gamma_t, R_t}{\to} C_{t+1}$ for all $t \geq 0$. We denote $\Xi_P(C_0, \Gamma, \Lambda)$ simply by $\Xi_P(C_0)$ when no misunderstanding can arise.

The leader election problem requires that every agent should output L or F which means "leader" or "follower" respectively. We say that a finite or infinite sequence of configurations $\xi = C_0, C_1, \ldots$ preserves a unique leader, denoted by $\xi \in LE$, if there exists $v \in V$ such that $O(C_t(v), \mathrm{id}(v)) = L$ and $O(C_t(u), \mathrm{id}(u)) = F$ for any $t \geq 0$ and $u \in V \setminus \{v\}$. For $\xi = C_0, C_1, \ldots$, the holding time of the leader $\mathrm{HT}(\xi, LE)$ is defined as the maximum $t \in \mathbb{N}$ that satisfies $(C_0, C_1, \ldots, C_{t-1}) \in LE$. We define $\mathrm{HT}(\xi, LE) = 0$ if $C_0 \notin LE$. We denote $\mathbf{E}[\mathrm{HT}(\Xi_P(C), LE)]$ by $\mathrm{EHT}_P(C, LE)$. Intuitively, $\mathrm{EHT}_P(C, LE)$ is the expected number of interactions for which the population keeps the unique leader after protocol P starts from configuration C. For configuration sequence

$\xi = C_0, C_1, \ldots$ and a set of configurations $\mathcal{C}$, we define convergence time $\mathrm{CT}(\xi, \mathcal{C})$ as the minimum $t \in \mathbb{N}$ that satisfies $C_t \in \mathcal{C}$. We define $\mathrm{CT}(\xi, \mathcal{C}) = |\xi|$ if $C_t \notin \mathcal{C}$ for any $t \geq 0$, where $|\xi|$ is the length of ξ. We denote $\mathbf{E}[\mathrm{CT}(\Xi_P(C), \mathcal{C})]$ by $\mathrm{ECT}_P(C, \mathcal{C})$. Intuitively, $\mathrm{ECT}_P(C, \mathcal{C})$ is the expected number of interactions by which the population enters a configuration in $\mathcal{C}$ after P starts from C.

Definition *Protocol $P(Q, Y, I, R, T, O)$ is an (α, β)-loosely-stabilizing leader election protocol if there exists set $\mathcal{S}$ of configurations satisfying two inequalities $\max_{C \in \mathcal{C}_{\mathrm{all}}(P)} \mathrm{ECT}_P(C, \mathcal{S}) \leq \alpha$ and $\min_{C \in \mathcal{S}} \mathrm{EHT}_P(C, LE) \geq \beta$.*

3 Leader Election with Identifiers

This section presents loosely-stabilizing leader election protocol P_{ID}, which works on arbitrary undirected graphs with unique identifiers of agents (Protocol 1). In the protocol description, we regard a state of agents as a collection of *variables* (e.g. `timer`), and denote a transition function as pseudo code that updates variables of initiator x and responder y. We denote the value of variable `var` of agent $v \in V$ by $v.\mathtt{var}$. We also denote the value of `var` in state $q \in Q$ by $q.\mathtt{var}$.

This protocol elects the agent with the minimum identifier, denoted by $v_{\min}$, as the leader. Each agent v tries to find the minimum identifier and stores it on $v.\mathtt{lid}$. At interaction, two agents x and y compare their `lid` and store the smaller value on their `lid` (Lines 3 and 6), by which the smallest identifier $\mathrm{id}(v_{\min})$ eventually spreads to all the agents. Then, after some point, $v_{\min}$ is the unique leader because output function O makes only agents v satisfying $\mathrm{id}(v) = v.\mathtt{lid}$ output L and other agents output F.

However, in the initial configuration, some agents may have false identifiers (or the integers that are not identifiers of any agent in the population) on `lid`. A false identifier may spread to the population instead of $\mathrm{id}(v_{\min})$ if it is smaller than $\mathrm{id}(v_{\min})$. We define $\mathrm{ID} = \{\mathrm{id}(v) \mid v \in V\}$, which is the correct identifiers set (Note that $\mathrm{ID} \subseteq I$). Protocol P_{ID} removes false identifiers $i \notin \mathrm{ID}$ from `lid` of all the agents by the *timeout mechanism*. Specifically, if $x.\mathtt{lid} \neq y.\mathtt{lid}$, we take the timer value of the agent with smaller `lid`, decrease it by one, and substitute the decreased value into $x.\mathtt{lid}$ and $y.\mathtt{lid}$ (Lines 4 and 7). If $x.\mathtt{lid} = y.\mathtt{lid}$, we take the larger value of $x.\mathtt{timer}$ and $y.\mathtt{timer}$, decrease it by one, and substitute the decreased value into $x.\mathtt{lid}$ and $y.\mathtt{lid}$ (Line 9). We call this event *larger value propagation*. If x or y is a leader, both timers are reset to $t_{\max}$ (Line 12). We call this event *timer reset*. When a timer becomes zero, agents x and y suspect that there exists no leader in the population. In this case, they elect the one with a smaller identifier as a leader by substituting $\min(\mathrm{id}(x), \mathrm{id}(y))$ into $x.\mathtt{lid}$ and $y.\mathtt{lid}$ (Line 14). We call this event *timeout*. Agents with false identifiers never experience timer reset; thus, their timers keep on decreasing. Hence, timeout eventually occurs and their `lid`s satisfy $\mathtt{lid} \in \mathrm{ID}$. This mechanism rarely ruins the stability of the unique leader because agents with $\mathtt{lid} \in \mathrm{ID}$ keep high value timers because of timer reset and lager value propagation.

Protocol 1 Leader Election with Identifiers P_{ID}

Variables of each agent:

$\texttt{lid} \in I$, $\texttt{timer} \in [0, t_{\max}]$

Output function O:

if $v.\texttt{lid} = \mathrm{id}(v)$ holds, then the output of agent v is L; Otherwise, F.

Interaction between initiator x and responder y:

1: **if** $x.\texttt{lid} > \mathrm{id}(x)$ **then** $x.\texttt{lid} \leftarrow \mathrm{id}(x)$ **endif**
2: **if** $x.\texttt{lid} < y.\texttt{lid}$ **then**
3: $\quad y.\texttt{lid} \leftarrow x.\texttt{lid}$
4: $\quad x.\texttt{timer} \leftarrow y.\texttt{timer} \leftarrow \max(x.\texttt{timer} - 1,\ 0)$
5: **else if** $x.\texttt{lid} > y.\texttt{lid}$ **then**
6: $\quad x.\texttt{lid} \leftarrow y.\texttt{lid}$
7: $\quad x.\texttt{timer} \leftarrow y.\texttt{timer} \leftarrow \max(y.\texttt{timer} - 1,\ 0)$
8: **else** // $x.\texttt{lid} = y.\texttt{lid}$ at this time
9: $\quad x.\texttt{timer} \leftarrow y.\texttt{timer} \leftarrow \max(x.\texttt{timer} - 1,\ y.\texttt{timer} - 1,\ 0)$
10: **end if**
11: **if** $\mathrm{id}(x) = x.\texttt{lid}$ **or** $\mathrm{id}(y) = y.\texttt{lid}$ **then** // a leader resets timers
12: $\quad x.\texttt{timer} \leftarrow y.\texttt{timer} \leftarrow t_{\max}$
13: **else if** $x.\texttt{timer} = 0$ **then** // a new leader is created at timeout
14: $\quad x.\texttt{lid} \leftarrow y.\texttt{lid} \leftarrow \min(\mathrm{id}(x),\ \mathrm{id}(y))$
15: $\quad x.\texttt{timer} \leftarrow y.\texttt{timer} \leftarrow t_{\max}$
16: **end if**

Complexity Analysis The upper bound $t_{\max}$ of variable `timer` is the only parameter of P_{ID}, which affects the correctness and complexities of the protocol. We assume $t_{\max} \geq 8\delta \max(d,\ 2+\log n)$ where δ is the maximum degree of the agents and d is the diameter of population G. (Note that δ is an even number because G is undirected.) We prove the following equations under this assumption:

$$\max_{C \in \mathcal{C}_{\mathrm{all}}} \mathrm{ECT}_{P_{\mathrm{ID}}}(C, \mathcal{S}_{\mathrm{id}}) = O(m\delta\tau \log n), \tag{1}$$

$$\min_{C \in \mathcal{S}_{\mathrm{id}}} \mathrm{EHT}_{P_{\mathrm{ID}}}(C, LE) = \Omega(\tau e^{\tau}), \tag{2}$$

where $\tau = t_{\max}/(8\delta)$ and $\mathcal{S}_{\mathrm{id}}$ is the set of configurations in which $v.\texttt{lid} = \mathrm{id}(v_{\min})$ and $v.\texttt{timer} > t_{\max}/2$ hold for all $v \in V$ and $v_{\min}.\texttt{timer} = t_{\max}$ holds. When upper bounds N of n and Δ of δ are available and we assign $t_{\max} = 8N\Delta$, protocol P_{ID} is an $(O(m\Delta N \log n), \Omega(Ne^{N}))$-loosely-stabilizing leader election protocol.

First, we analyze the expected holding time. Let $C_0 \in \mathcal{S}_{\mathrm{id}}$ and $\Xi_{P_{\mathrm{ID}}}(C_0) = C_0, C_1, \ldots$. To prove (2), it suffices to show that both $C_0, \ldots, C_{2m\tau} \in LE$ and $C_{2m\tau} \in \mathcal{S}_{\mathrm{id}}$ hold with probability at least $p_{\mathrm{suc}} = 1 - O(ne^{-\tau})$. Then, we have $\min_{C_0 \in \mathcal{S}_{\mathrm{id}}} \mathrm{EHT}_{P_{\mathrm{ID}}}(C_0, LE) \geq 2m\tau/(1 - p_{\mathrm{suc}}) = \Omega(\tau e^{\tau})$.

Note We use some variants of Chernoff bounds for the proofs of Sections 3 and 4. See appendix for those Chernoff bounds.

Lemma 1 *The probability that every $v \in V$ joins only less than $t_{\max}/2$ interactions among $\Gamma_0, \ldots, \Gamma_{2m\tau-1}$ is at least $1 - ne^{-\tau}$.*

Proof For any $v \in V$ and $t \geq 0$, v joins interaction Γ_t with probability at most δ/m. Thus, the number of interactions v joins during the $2m\tau$ interactions is bounded by binomial random variable $X \sim B(2m\tau, \delta/m)$. Applying a variant of Chernoff bound (See Lemma C1 in Appendix), we have

$$\begin{aligned}
\Pr(X \geq t_{\max}/2) &= \Pr(X \geq 2\mathbf{E}[X]) && \because t_{\max} = 8\delta\tau \\
&\leq e^{-\mathbf{E}[X]/3} \\
&= e^{-2\delta\tau/3} && \text{(By Chernoff Bound of Lemma C1)} \\
&\leq e^{-\tau}. && \because \delta \geq 2
\end{aligned}$$

Summing up the probabilities for all $v \in V$ gives the lemma. □

Lemma 2 *Let $C_0 \in \mathcal{L}_{\text{lid}}$ and $\Xi_{P_{\text{ID}}}(C_0) = C_0, C_1, \ldots$. Then, we have the following inequality:*

$$\Pr(\forall v \in V,\ C_{2m\tau}(v).\texttt{timer} > t_{\max}/2) \geq 1 - 2ne^{-\tau}.$$

Proof It suffices to show $\Pr(C_{2m\tau}(v).\texttt{timer} > t_{\max}/2) \geq 1 - 2e^{-\tau}$ for any agent $v \in V$. We denote the shortest path from $v_{\min}$ to v by $(v_0, v_1, \ldots, v_k)$ where $v_0 = v_{\min},\ v_k = v,\ 0 \leq k \leq d$ and $(v_{i-1}, v_i) \in E$ for all $i = 1, \ldots, k$. For any $t \in [0, 2m\tau]$, we define $v_{\text{head}}(t)$ as v_l with maximum $l \in [1, k]$ such that there exist $t_1, t_2, \ldots, t_l$ satisfying $0 \leq t_1 < t_2 < \cdots < t_l < t$ and $\Gamma_{t_i} \in \{(v_{i-1}, v_i), (v_i, v_{i-1})\}$ for $i = 1, 2, \ldots, l$. We define $v_{\text{head}}(t) = v_0$ if such l does not exist. Intuitively, $v_{\text{head}}(t)$ is the head of the agents in path $(v_0, v_1, \ldots, v_k)$ to which a large timer value is propagated from $v_{\min}$. (Remember that $v_{\min}$ resets the timers to $t_{\max}$.) We define $J(t)$ as the number of integers $i \in [0, t]$ such that $v_{\text{head}}(i)$ joins interaction Γ_i. Intuitively, $J(t)$ is the number of interactions that the head agent joins among $\Gamma_0, \ldots, \Gamma_t$. Obviously, we have $C_t(v_{\text{head}}(t)).\texttt{timer} \geq t_{\max} - J(t)$ for any $t \geq 0$.

In what follows, we prove $\Pr(v_{\text{head}}(2m\tau) = v) \geq 1 - e^{-\tau}$ and $\Pr(J(2m\tau) < t_{\max}/2) \geq 1 - e^{-\tau}$, which give $\Pr(C_{2m\tau}(v).\texttt{timer} > t_{\max}/2) \geq 1 - 2e^{-\tau}$. For any $i \in [1, k]$, a pair v_{i-1} and v_i interacts with probability $2/m$ at each interaction. Hence, we can say each interaction makes v_{head} forward with probability $2/m$. Therefore, by letting Z be a binomial random variable such that $Z \sim B(2m\tau, 2/m)$, we have

$$\begin{aligned}
\Pr(v_{\text{head}}(t) = v) &= 1 - \Pr(Z < k) \\
&\geq 1 - \Pr(Z < d) \\
&\geq 1 - \Pr\left(Z < \frac{1}{4} \cdot \mathbf{E}[Z]\right) && \because d \leq \tau = \frac{1}{4} \cdot \mathbf{E}[Z] \\
&\geq 1 - e^{-9\mathbf{E}[Z]/32} && \text{(By Chernoff bound of Lemma C3)} \\
&> 1 - e^{-\tau}.
\end{aligned}$$

The probability that $v_{\text{head}}(t)$ joins interaction Γ_t is at most δ/m regardless of t. Hence, by letting Z' be a binomial random variable such that $Z' \sim B(2m\tau, \delta/m)$,

we have

$$\begin{aligned}\Pr(J(2m\tau) < t_{\max}/2) &> 1 - \Pr(Z' \geq t_{\max}/2) \\ &= 1 - \Pr(Z' \geq 2\mathbf{E}[Z']) \\ &> 1 - e^{-\mathbf{E}[Z']/3} \quad \text{(By Chernoff bound of Lemma C1)} \\ &= 1 - e^{-2\delta\tau/3} \\ &> 1 - e^{-\tau}. \qquad \because \delta \geq 2\end{aligned}$$

Thus, we have shown $\Pr(C_{2m\tau}(v).\texttt{timer} > t_{\max}/2) \geq 1 - 2e^{-\tau}$. □

Lemma 3 $\min_{C \in \mathcal{S}_{\mathrm{id}}} \mathrm{EHT}_{P_{\mathrm{ID}}}(C, LE) = \Omega(\tau e^{\tau})$.

Proof We have $C_0, \ldots, C_{2m\tau} \in LE$ and $C_{2m\tau} \in \mathcal{S}_{\mathrm{id}}$ if $C_0 \in \mathcal{S}_{\mathrm{id}}$ holds, no timeout happens, and any agent interacts at most $t_{\max}/2$ times during $2m\tau$ interactions. Hence, probability p_{suc} discussed in the beginning of this section is at least $1 - 3ne^{-\tau}$ by Lemmas 1 and 2, which leads to the lemma. □

Next, we analyze the expected convergence time. To prove (1), we define two sets of configurations: $\mathcal{C}_{\mathrm{lid}} = \{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{ID}}) \mid \forall v \in V, C(v).\texttt{lid} \in \mathrm{ID}\}$ and $\mathcal{L}_{\mathrm{lid}} = \mathcal{C}_{\mathrm{lid}} \cap \{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{ID}}) \mid C(v_{\min}).\texttt{lid} = \mathrm{id}(v_{\min}) \wedge C(v_{\min}).\texttt{timer} = t_{\max}\}$.

Lemma 4 $\max_{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{ID}})} \mathrm{ECT}_{P_{\mathrm{ID}}}(C, \mathcal{C}_{\mathrm{lid}}) = O(m\delta\tau \log n)$.

Proof Let z be the maximum value of $v.\texttt{timer}$ such that $v.\texttt{lid} \notin \mathrm{ID}$. This z decreases by one every time all interactions of E occur. Thus, it takes at most $\frac{m}{m} + \frac{m}{m-1} + \ldots \frac{m}{1} \leq m(1 + \log m)$ expected steps to decrease z by one. Hence, $\max_{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{ID}})} \mathrm{ECT}_{P_{\mathrm{ID}}}(C, \mathcal{C}_{\mathrm{lid}}) \leq t_{\max} m(1 + \log m) = O(m\delta\tau \log n)$. □

Lemma 5 $\max_{C \in \mathcal{C}_{\mathrm{lid}}} \mathrm{ECT}_{P_{\mathrm{ID}}}(C, \mathcal{L}_{\mathrm{lid}}) = O(m)$.

Proof We have $v_{\min}.\texttt{lid} = \mathrm{id}(v_{\min})$ and $v_{\min}.\texttt{timer} = t_{\max}$ just after $v_{\min}$ interacts in any configuration of $\mathcal{C}_{\mathrm{lid}}$. This takes $O(m)$ expected interactions. □

Lemma 6 $\max_{C \in \mathcal{L}_{\mathrm{lid}}} \mathrm{ECT}_{P_{\mathrm{ID}}}(C, \mathcal{S}_{\mathrm{id}}) = O(m\tau)$.

Proof Sketch Let $C_0 \in \mathcal{L}_{\mathrm{lid}}$ and $\Xi_{P_{\mathrm{ID}}}(C_0) = C_0, C_1, \ldots$. By similar argument to Lemmas 1 and 2, we can prove $\Pr(C_{2m\tau} \in \mathcal{S}_{\mathrm{id}}) > 1 - 2ne^{-\tau}$. Since $C \in \mathcal{L}_{\mathrm{lid}}$ cannot change to $D \notin \mathcal{L}_{\mathrm{lid}}$, we have $\max_{C \in \mathcal{L}_{\mathrm{lid}}} \mathrm{ECT}_{P_{\mathrm{ID}}}(C, \mathcal{S}_{\mathrm{id}}) \leq 2m\tau + 3ne^{-\tau} \cdot \max_{C \in \mathcal{L}_{\mathrm{lid}}} \mathrm{ECT}_{P_{\mathrm{ID}}}(C, \mathcal{S}_{\mathrm{id}})$. Solving this inequality gives the lemma. (See Appendix for the complete proof.) □

The following lemma immediately follows from Lemmas 4, 5, and 6.

Lemma 7 $\max_{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{ID}})} \mathrm{ECT}_{P_{\mathrm{ID}}}(C, \mathcal{S}_{\mathrm{id}}) = O(m\delta\tau \log n)$.

Lemmas 3 and 7 gives the following theorem.

Theorem 1 *Protocol P_{ID} is a $(O(m\delta\tau \log n), \Omega(\tau e^{\tau}))$ loosely-stabilizing leader election protocol for arbitrary graphs when $t_{\max} \geq 8\delta \max(d, 2 + \log n)$.*

Therefore, given upper bound N and Δ of n and δ respectively, we get a $(O(m\Delta N \log n), \Omega(Ne^{N}))$ loosely-stabilizing leader election protocol for arbitrary graphs by assigning $t_{\max} = 8N\Delta$.

4 Leader Election with Random Numbers

This section presents loosely-stabilizing leader election protocol P_{RD}. It works on arbitrary undirected anonymous graphs with a random number generated at each interaction (Protocol 2). Random numbers are used in Line 11: When the protocol enters Line 11, the code is executed with probability $p = 1/|R|$. This is implemented as the code is executed only when a specific number is generated. For example, $p = 0.01$ if we assign $R = [0, 99]$ and treat 0 as a specific number.

Each agent has binary variable $\mathtt{DoA} \in \{\mathrm{DEAD}, \mathrm{ALIVE}\}$ and three timers $\mathtt{timer_L}$, $\mathtt{timer_V}$ and $\mathtt{timer_S}$. The output function defines leaders based on $\mathtt{DoA}$: agent v is a leader if v is alive (or $v.\mathtt{DoA} = \mathrm{ALIVE}$), and a follower if v is dead (or $v.\mathtt{DoA} = \mathrm{DEAD}$). Protocol P_{RD} consists of a timeout mechanism (Lines 1-7) and a virus-war mechanism (Lines 8-14). By using $\mathtt{timer_L}$, the timeout mechanism creates a leader when it is suspected that no leader exists. By using $\mathtt{timer_V}$ and $\mathtt{timer_S}$, the virus-war mechanism reduces the number of leaders.

The timeout mechanism is almost the same as P_{ID}. By the timer reset and the larger value propagation, timeout eventually occurs when no leader exists, and all agents keep high timer values with high probability when one ore more leaders exist. At timeout, a dead agent becomes a leader (Line 5).

In the virus-war mechanism, each leader tries to kill other leaders by viruses and become the unique leader. We say that agent v has a virus if $v.\mathtt{timer_V} > 0$, and v wears a (head) shield if $v.\mathtt{timer_S} > 0$. A leader creates a new virus with probability p when it interacts as an initiator (Line 11). When creating a virus, the agent wears a shield so as not to be killed by the new virus (Line 11). A virus spreads among agents by interactions (Line 8), and an agent is killed when it has a virus without a shield (Lines 13-14). A virus has TTL (time to live), which is memorized on $\mathtt{timer_V}$ and decreased by one at each interaction of its owner (line 8). When $\mathtt{timer_V}$ becomes zero, the virus vanishes and looses the ability to kill agents. A shield also has TTL, which is memorized on $\mathtt{timer_S}$ and decreased by one at each interaction of its owner (Line 9). When $\mathtt{timer_S}$ becomes zero, the shield vanishes and looses the ability to protect its owner from viruses.

The virus-war mechanism correctly works if p is sufficiently small and t_{shld} is sufficiently greater than t_{virus}. Consider the case multiple leaders exist. Since p is small, all viruses and shields eventually vanishes. After that, some agent eventually creates a new virus and shield. The created virus kill all other agents unless some of them also create a new virus and shield before the virus reaches them. Since p is sufficiently small, the probability of the latter is small. Thus, the unique leader is elected within a relatively short time. Even after that, the unique leader keeps on creating new viruses, each of which may kill the leader. However, the leader is not killed for an extremely long time: since $t_{\mathrm{shld}} \gg t_{\mathrm{virus}}$, the leader's shield rarely vanishes before all viruses vanish from the population.

Complexity Analysis We have four parameters in P_{RD}: three upper bounds $t_{\max}$, t_{virus}, and t_{shld} of the timers, and probability p. We assume $t_{\mathrm{virus}} = t_{\max}/2$, $t_{\max} \geq 8\delta \max(d,\ 2 + \log(13n\delta\lceil \log n \rceil))$, $t_{\mathrm{shld}} \geq 2\delta t_{\max}\lceil \log n \rceil$ and $p \leq$

Protocol 2 Leader Election with Random Numbers P_{RD}

Variables of each agent:

$\mathtt{DoA} \in \{\mathrm{DEAD},\ \mathrm{ALIVE}\}$, $\mathtt{timer_L} \in [0, t_{\max}]$, $\mathtt{timer_V} \in [0, t_{\mathrm{virus}}]$, $\mathtt{timer_S} \in [0, t_{\mathrm{shld}}]$

Output function O:

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

Interaction between initiator x and responder y:

$x.\mathtt{timer_L} \leftarrow y.\mathtt{timer_L} \leftarrow \max(x.\mathtt{timer_L} - 1,\ y.\mathtt{timer_L} - 1,\ 0)$
if $x.\mathtt{DoA} = \mathrm{ALIVE}$ **or** $y.\mathtt{DoA} = \mathrm{ALIVE}$ **then**
  $x.\mathtt{timer_L} \leftarrow y.\mathtt{timer_L} \leftarrow t_{\max}$ // a leader resets timer
else if $x.\mathtt{timer_L} = 0$ **then** // a new leader is created at timeout
  $x.\mathtt{DoA} \leftarrow \mathrm{ALIVE}$
  $x.\mathtt{timer_L} \leftarrow y.\mathtt{timer_L} \leftarrow t_{\max}$
end if
$x.\mathtt{timer_V} \leftarrow y.\mathtt{timer_V} \leftarrow \max(x.\mathtt{timer_V} - 1,\ y.\mathtt{timer_V} - 1,\ 0)$
$x.\mathtt{timer_S} \leftarrow \max(0, x.\mathtt{timer_S} - 1)$
if $x.\mathtt{DoA} = \mathrm{ALIVE}$ **then**
  Execute ($x.\mathtt{timer_S} \leftarrow t_{\mathrm{shld}}$, $x.\mathtt{timer_V} \leftarrow t_{\mathrm{virus}}$) with probability p
end if
if $x.\mathtt{timer_V} > 0$ **and** $x.\mathtt{timer_S} = 0$ **then** $x.\mathtt{DoA} \leftarrow \mathrm{DEAD}$ **endif**
if $y.\mathtt{timer_V} > 0$ **and** $y.\mathtt{timer_S} = 0$ **then** $y.\mathtt{DoA} \leftarrow \mathrm{DEAD}$ **endif**

$(4mt_{\mathrm{shld}})^{-1}$. We prove the following equations under this assumption:

$$\max_{C \in \mathcal{C}_{\mathrm{all}}} \mathrm{ECT}_{P_{\mathrm{RD}}}(C, \mathcal{S}_{\mathrm{RD}}) = O(mp^{-1}), \tag{3}$$

$$\min_{C \in \mathcal{S}_{\mathrm{RD}}} \mathrm{EHT}_{P_{\mathrm{RD}}}(C, LE) = \Omega(\tau e^{\tau}), \tag{4}$$

where $\tau = t_{\max}/(8\delta)$ and $\mathcal{S}_{\mathrm{RD}}$ is the set of configurations we define later. When upper bounds N and Δ are available and we assign $t_{\max} = 8N\Delta$, $t_{\mathrm{shld}} = 2\Delta t_{\max}\lceil \log N \rceil$ and $p = (4N^2 t_{\mathrm{shld}})^{-1}$ (*i.e.*, $R = [0, 4N^2 t_{\mathrm{shld}} - 1]$), then P_{RD} is an $(O(m\Delta^2 N^3 \log N), \Omega(Ne^N))$-loosely-stabilizing leader election protocol.

Before proving equations (3) and (4), we define five sets of configurations:

$$\mathcal{G}_{\mathrm{half}} = \{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{RD}}) \mid \exists v \in V,\ C(v).\mathtt{DoA} = \mathrm{ALIVE} \wedge C(v).\mathtt{timer_S} > t_{\mathrm{shld}}/2\},$$
$$\mathcal{V}_{\mathrm{clean}} = \{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{RD}}) \mid \forall v \in V,\ C(v).\mathtt{timer_V} = 0\},$$
$$\mathcal{L}_{\mathrm{half}} = \{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{RD}}) \mid \#_L(C) \geq 1 \ \wedge\ \forall v \in V,\ C(v).\mathtt{timer_L} > t_{\max}/2\},$$
$$\mathcal{L}_{\mathrm{one}} = \{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{RD}}) \mid \#_L(C) = 1\},$$
$$\mathcal{S}_{\mathrm{RD}} = (\mathcal{G}_{\mathrm{half}} \cup \mathcal{V}_{\mathrm{clean}}) \cap \mathcal{L}_{\mathrm{half}} \cap \mathcal{L}_{\mathrm{one}},$$

where $\#_L(C)$ denotes the number of leaders in configuration C. Note that $\mathcal{G}_{\mathrm{half}}$ requires that not all agents but at least one leader has $\mathtt{timer_S}$ more than $t_{\mathrm{shld}}/2$.

First, we analyze the expected holding time. Let $C_0 \in \mathcal{S}_{\mathrm{RD}}$ and $\Xi_{P_{\mathrm{RD}}}(C_0) = C_0, C_1, \ldots$. To prove (4), it suffices to show that both $C_0, \ldots, C_{8m\delta\tau\lceil \log n \rceil} \in LE$ and $C_{8m\delta\tau\lceil \log n \rceil} \in \mathcal{S}_{\mathrm{RD}}$ hold with probability no less than $p_{\mathrm{suc}} = 1 - O(n\delta \log n \cdot e^{-\tau})$. Then, $\min_{C_0 \in \mathcal{S}_{\mathrm{RD}}} \mathrm{EHT}_{P_{\mathrm{RD}}}(C_0, LE) \geq 8m\delta\tau\lceil \log n \rceil \tau/(1 - p_{\mathrm{suc}}) = \Omega(\tau e^{\tau})$.

We define two predicates PROP_i and HALF_i for any $i \geq 0$: $\mathrm{PROP}_i = 1$ if $C_{2m\tau(i+1)}(v).\mathtt{timer_L} > t_i - t_{\max}/2$ for all $v \in V$, otherwise $\mathrm{PROP}_i = 0$, where

$t_i = \max_{v \in V} C_{2m\tau i}(v)$; $\mathrm{HALF}_i = 1$ if every agent joins only less than $t_{\max}/2$ interactions among $\Gamma_{2m\tau i}, \ldots, \Gamma_{2m\tau(i+1)-1}$, otherwise $\mathrm{HALF}_i = 0$. Intuitively, $\mathrm{PROP}_i = 1$ means the maximum value of $\mathtt{timer_L}$ propagates to all the agents well during the $2m\tau$ interactions, and $\mathrm{HALF}_i = 1$ means every agent does not interact so much during the $2m\tau$ interactions.

Lemma 8 *Let $C_0 \in \mathcal{S}_{\mathrm{RD}}$ and $\Xi_{P_{\mathrm{RD}}}(C_0) = C_0, C_1, \ldots$. Then, we have both $C_0, \ldots, C_{8m\delta\tau\lceil \log n \rceil} \in LE$ and $C_{8m\delta\tau\lceil \log n \rceil} \in \mathcal{S}_{\mathrm{RD}}$ if the following conditions hold:*
(A) $\#_L(C_t) \geq 1$ for all $t = 0, \ldots, 8m\delta\tau\lceil \log n \rceil$,
(B) $C_{8m\delta\tau\lceil \log n \rceil} \in \mathcal{G}_{\mathrm{half}} \cup \mathcal{V}_{\mathrm{clean}}$,
(C) $\mathrm{PROP}_i = 1$ for all $i = 0, \ldots, 4\delta\lceil \log n \rceil - 1$, and
(D) $\mathrm{HALF}_i = 1$ for all $i = 0, \ldots, 4\delta\lceil \log n \rceil - 1$.

Proof We have $C_{2m\tau i}(v).\mathtt{timer_L} > t_{\max}/2$ for any $i \in [0, 4\delta\lceil \log n \rceil]$ from (A) and (C). Since no agent interacts more than $t_{\max}/2$ times among each $2m\tau$ interactions (i.e. (D)), timeout does not occur at any interaction $\Gamma_0, \ldots, \Gamma_{8m\delta\tau\lceil \log n \rceil - 1}$, by which we obtain $C_0, \ldots, C_{8m\delta\tau\lceil \log n \rceil} \in LE$. We also obtain $C_{8m\delta\tau\lceil \log n \rceil} \in \mathcal{L}_{\mathrm{half}} \cap \mathcal{L}_{\mathrm{one}} \cap (\mathcal{G}_{\mathrm{half}} \cup \mathcal{V}_{\mathrm{clean}}) = \mathcal{S}_{\mathrm{RD}}$ from above discussion and (B). □

Lemma 9 *The probability that every agent joins only less than $t_{\mathrm{shld}}/2$ interactions as an initiator among $\Gamma_0, \ldots, \Gamma_{8m\delta\tau\lceil \log n \rceil - 1}$ is at least $1 - ne^{-\delta\tau}$.*

Proof For any $v \in V$ and $t \geq 0$, v joins interaction Γ_t as an initiator with probability at most $\delta/(2m)$ since v has at most $\delta/2$ outgoing edges. Thus, the number of interactions v joins as an initiator during the $8m\delta\tau\lceil \log n \rceil$ interactions is bounded by binomial random variable $X \sim B(8m\delta\tau\lceil \log n \rceil, \delta/(2m))$. We have

$$\begin{aligned}
\Pr(X \geq t_{\mathrm{shld}}/2) &\leq \Pr(X \geq 8\delta^2\tau\lceil \log n \rceil) && \because t_{\mathrm{shld}} \geq 16\delta^2\tau\lceil \log n \rceil \\
&= \Pr(X \geq 2\mathbf{E}[X]) \\
&\leq e^{-\mathbf{E}[X]/3} && \text{(By Chernoff Bound of Lemma C2)} \\
&= e^{-4\delta^2\tau\lceil \log n \rceil/3} \\
&= e^{-\delta\tau}.
\end{aligned}$$

Summing up these probabilities gives the lemma. □

Lemma 10 *Let $C_0 \in \mathcal{S}_{\mathrm{RD}}$ and $\Xi_{P_{\mathrm{RD}}}(C_0) = C_0, C_1, \ldots$.*
Then, we have $\Pr(\forall t \in [0, 8m\delta\tau\lceil \log n \rceil - 1],\ \#_L(C_t) \geq 1) \geq 1 - ne^{-\delta\tau}$.

Proof By Lemma 9, it suffices to show that $\#_L(C_t) \geq 1$ holds for all $t \in [0, 8m\delta\tau\lceil \log n \rceil]$ when we assume every agent joins only less than $t_{\mathrm{shld}}/2$ interactions as an initiator among $\Gamma_0, \ldots, \Gamma_{8m\delta\tau\lceil \log n \rceil - 1}$. Since $C_0 \in \mathcal{S}_{\mathrm{RD}}$, we have $C_0 \in \mathcal{G}_{\mathrm{half}} \cup \mathcal{V}_{\mathrm{clean}}$. If $C_0 \in \mathcal{G}_{\mathrm{half}}$, there exists a leader v such that $C_0(v).\mathtt{timer_S} > t_{\mathrm{shld}}/2$. By the assumption, v decrease its $\mathtt{timer_S}$ by at most $t_{\mathrm{shld}}/2$; thus, v is never killed and remains a leader in $C_0, \ldots, C_{8m\delta\tau\lceil \log n \rceil}$. If $C_0 \in \mathcal{V}_{\mathrm{clean}}$, no leader is killed before a new virus is created. Even if some leader u creates a new virus at interaction Γ_t ($0 \leq t < 8m\delta\tau\lceil \log n \rceil$), u wears a new shield at the same time. Hence, u remains a leader in $C_t, \ldots, C_{8m\delta\tau\lceil \log n \rceil}$ by the assumption. □

We define the first round time $\mathrm{RT}_\Gamma(1)$ as the minimum t satisfying $\forall e \in E,\ 0 \leq \exists t' \leq t,\ \Gamma_{t'} = e$. For any $i \geq 2$, we define the i-th round time $\mathrm{RT}_\Gamma(i)$ as the minimum t satisfying $\forall e \in E,\ \mathrm{RT}_\Gamma(i-1) < \exists t' \leq t,\ \Gamma_{t'} = e$. Lemma 12 bounds $\mathrm{RT}_\Gamma(i)$ from above with high probability. To prove the lemma, we firstly prove Lemma 11.

Lemma 11 *Let $v_1, v_2, \ldots, v_l$ be any l $(l < n)$ agents in V. There exists at least $2l$ edges of E that are incident to at least one of the l agents.*

Proof Since $l < n$, there exists agent $r \in V$ that differs from any $v_1, v_2, \ldots, v_l$. Since G is strongly connected, there exists a rooted spanning tree T on G where r is the root agent of T. Then, every v_i $(i \in [1, k])$ has two edges between v_i and the parent agent of v_i in T. (Remind that G is undirected, that is, $(u, v) \in E \Leftrightarrow (v, u) \in E$ for any $u, v \in V$.) These edges are mutually exclusive. Thus, we have $2l$ edges of E that are incident to at least one of the l agents. □

Lemma 12 $\Pr(\mathrm{RT}_\Gamma(i) < 2im\lceil \log n \rceil) \geq 1 - ne^{-i/4}$ *holds for any* $i \geq 1$.

Proof Each round j $(j \geq 1)$ finishes when every agent $v \in V$ interacts in round j. Consider the case that k $(k \geq 1)$ agents have not yet interacted in round j and only $n-k$ agents have interacted in round j. We call the former uninvolved agents and the latter involved agents. If $k < n$, one of the k uninvolved agents joins the next interaction and becomes an involved agent with probability more than $2k/m$ by Lemma 11. If $k = n$, some uninvolved agent joins the next interaction with probability 1. Let $X_{j,k}$ $(j \geq 1,\ k \geq 1)$ be the random variable that corresponds to the number of trials to the first success in which the success probability of each trial is $2k/m$. From the above discussion, we obtain

$$\begin{aligned} \Pr(\mathrm{RT}_\Gamma(i) \geq 2im\lceil \log n \rceil) &\leq \Pr\left(\sum_{j=1}^{i} \left(1 + \sum_{k=1}^{n-1} X_{j,k} \right) \geq 2im\lceil \log n \rceil \right) \\ &\leq \Pr\left(\sum_{k=1}^{n-1} \sum_{j=1}^{i} X_{j,k} \geq 2im\lceil \log n \rceil - i \right). \end{aligned} \tag{5}$$

For binomial random variable $Y_k \sim B(\lceil \frac{im}{k} \rceil, \frac{2k}{m})$, we have $\Pr(\sum_{j=1}^{i} X_{j,k} > \frac{im}{k}) \leq \Pr(\sum_{j=1}^{i} X_{j,k} \geq \lceil \frac{im}{k} \rceil) \leq \Pr(Y_k \leq i)$. Hence, we have

$$\begin{aligned} \Pr\left(\sum_{j=1}^{i} X_{j,k} > \frac{im}{k} \right) &\leq \Pr(Y_k \leq i) \\ &\leq \Pr\left(Y_k \leq \frac{1}{2} \cdot \mathbf{E}[Y_K] \right) \\ &\leq e^{-\mathbf{E}[Y_k]/8} \qquad \text{(By Chernoff Bound of Lemma C2)} \\ &\leq e^{-i/4}. \end{aligned} \tag{6}$$

From Inequalities (5) and (6), we have

$$\begin{aligned}\Pr(\mathrm{RT}_{\Gamma}(i) \geq 2im\lceil \log n \rceil) &\leq \Pr\left(\sum_{k=1}^{n-1}\sum_{j=1}^{i} X_{j,k} \geq 2im\lceil \log n \rceil - i\right) \\ &\leq \Pr\left(\sum_{k=1}^{n-1}\sum_{j=1}^{i} X_{j,k} > \sum_{k=1}^{n-1}\frac{im}{k}\right) \\ &\leq \sum_{k=1}^{n-1}\Pr\left(\sum_{j=1}^{i} X_{j,k} > \frac{im}{k}\right) \\ &\leq ne^{-i/4},\end{aligned}$$

where $\sum_{k=1}^{n-1}\frac{im}{k} \leq im(1+\log n) - i < 2im\lceil \log n \rceil - i$ is used for the second inequality. Thus, $\Pr(\mathrm{RT}_{\Gamma}(i) < 2im\lceil \log n \rceil) \geq 1 - ne^{-i/4}$ holds. □

Lemma 13 *Let* $C_0 \in \mathcal{S}_{\mathrm{RD}}$ *and* $\Xi_{P_{\mathrm{RD}}}(C_0) = C_0, C_1, \ldots$.
Then, we have $\Pr(C_{8m\delta\tau\lceil \log n \rceil} \in \mathcal{G}_{\mathrm{half}} \cup \mathcal{V}_{\mathrm{clean}}) \geq 1 - 2ne^{-\delta\tau}$.

Proof Assume that $\mathrm{RT}_{\Gamma}(t_{\mathrm{virus}}) < 8m\delta\tau\lceil \log n \rceil$ holds and every agent joins only less than $t_{\mathrm{shld}}/2$ interactions as an initiator among $\Gamma_0, \ldots, \Gamma_{8m\delta\tau\lceil \log n \rceil - 1}$. These assumptions lead to $C_{8m\delta\tau\lceil \log n \rceil} \in \mathcal{G}_{\mathrm{half}} \cup \mathcal{V}_{\mathrm{clean}}$ as follows. If a new virus is not created among $\Gamma_0, \ldots, \Gamma_{8m\delta\tau\lceil \log n \rceil - 1}$, then all viruses in the initial configuration vanish during the period since each round decreases the maximum value of $\mathtt{timer_V}$ by at least one. Thus, $C_{8m\delta\tau\lceil \log n \rceil} \in \mathcal{V}_{\mathrm{clean}}$ holds. If some agent v creates a new virus at Γ_t, then v wears a new shield at the same time. Thus, $C_{t+1}(v).\mathtt{timer_S} = t_{\mathrm{shld}}$. Since v interacts as an initiator only less than $t_{\mathrm{shld}}/2$ times among $\Gamma_{t+1}, \ldots, \Gamma_{8m\delta\tau\lceil \log n \rceil - 1}$, we have $C_{8m\delta\tau\lceil \log n \rceil}(v).\mathtt{timer_S} > t_{\mathrm{shld}}/2$, which means $C_{8m\delta\tau\lceil \log n \rceil} \in \mathcal{G}_{\mathrm{half}}$. By $t_{\mathrm{virus}} = 4\delta\tau$ and Lemmas 9 and 12, the probability that the two assumptions hold is at least $1 - 2ne^{-\delta\tau}$. □

Lemma 14 $\Pr(\mathrm{PROP}_i = 1) \geq 1 - 2ne^{-\tau}$ *for any* $i \geq 0$.

Proof The same argument as the proof of Lemma 2 gives the lemma. □

Lemma 15 $\Pr(\mathrm{HALF}_i = 1) \geq 1 - ne^{-\tau}$ *for any* $i \geq 0$.

Proof Each interaction is independent. Thus, Lemma 1 gives the lemma. □

Lemma 16 $\min_{C \in \mathcal{S}_{\mathrm{RD}}} \mathrm{EHT}_{P_{\mathrm{RD}}}(C, LE) = \Omega(\tau e^{\tau})$.

Proof Probability p_{suc}, discussed in the beginning of this section, is at least $1 - 3ne^{-\delta\tau} - 4\delta\lceil \log n \rceil \cdot 3ne^{-\tau} \geq 1 - 13n\delta\lceil \log n \rceil e^{-\tau}$ by Lemmas 8, 10, 13, 14 and 15, which leads to the lemma. □

Next, we analyze the expected convergence time. We define two sets of configurations: $\mathcal{N}_{\mathrm{oVG}} = \{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{RD}}) \mid \forall v \in V,\ C(v).\mathtt{timer_V} = C(v).\mathtt{timer_S} = 0\}$ and $\mathcal{L} = \{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{RD}}) \mid \#_L(C) \geq 1\}$.

Lemma 17 $\max_{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{RD}})} \mathrm{ECT}_{P_{\mathrm{RD}}}(C, \mathcal{S}_{\mathrm{RD}}) = O(mp^{-1})$.

Proof Sketch Probability p, with which a leader creates a virus at each interaction, is sufficiently small ($p < 1/(4mt_{\mathrm{shld}})$). Thus, the probability that all viruses and shields vanish (i.e. the population enters a configuration of $\mathcal{N}_{\mathrm{oVG}}$) within $2mt_{\mathrm{shld}}$ interactions is at least $1-(2mt_{\mathrm{shld}} \cdot p + O(ne^{-\tau})) > 1/2 - O(ne^{-\tau})$. Even if the reached configuration of $\mathcal{N}_{\mathrm{oVG}}$ does not have any leader, the timeout mechanism creates a leader, and the population enters a configuration of $\mathcal{N}_{\mathrm{oVG}} \cap \mathcal{L}$. This takes less than $16m\delta\tau\lceil \log n \rceil$ interactions with probability $1 - O(ne^{-\tau})$. After the population enters into $\mathcal{N}_{\mathrm{oVG}} \cap \mathcal{L}$, additional $\lceil m/p \rceil$ interactions create a new virus with probability $1 - e^{-2}$. Let v be a leader that creates the virus. Since v wears a new shield at the same time, v is not killed and remains a leader during the next $2m\tau$ interactions with probability $1 - O(e^{-\tau})$. On the other hand, the virus spreads to all the agents within these $2m\tau$ interactions with probability $1 - O(ne^{-\tau})$, killing all the agents other than v. A leader other than v may create a new virus during the $2m\tau$ interactions, and survives with a shield. However, this probability is at most $2m\tau \cdot p \leq 1/4$. Hence, v becomes the unique leader within the $2m\tau$ interactions with probability $3/4 - O(ne^{-\tau})$. After the $2m\tau$ interactions, all the agents have $\mathtt{timer}_{\mathtt{L}} > t_{\max}/2$ with probability $1 - O(ne^{-\tau})$ by the larger value propagation, and $v.\mathtt{timer}_{\mathtt{S}} > t_{\mathrm{shld}}/2$ holds with probability $1 - O(ne^{-\tau})$. Hence, the population enters a configuration of $\mathcal{L}_{\mathrm{one}} \cap \mathcal{L}_{\mathrm{half}} \cap \mathcal{G}_{\mathrm{half}} \subset \mathcal{S}_{\mathrm{RD}}$ within the $2m\tau$ interactions with probability $3/4 - O(ne^{-\tau})$. As a result, starting from any configuration, the population enters into $\mathcal{S}_{\mathrm{RD}}$ within $O(mp^{-1})$ interactions with probability $1/4 - e^{-2} - O(ne^{-\tau}) > 0.11 - o(1)$, which gives the lemma. (See Appendix for the complete proof.) □

Lemmas 16 and 17 gives the following theorem.

Theorem 2 *Protocol P_{RD} is a $(O(mp^{-1}), \Omega(\tau e^{\tau}))$ loosely-stabilizing leader election protocol for arbitrary graphs when $t_{\max} \geq 8\delta \max(d, 2 + \log(13n\delta\lceil \log n \rceil))$, $t_{\mathrm{virus}} = t_{\max}/2$, $t_{\mathrm{shld}} \geq 2\delta t_{\max}\lceil \log n \rceil$ and $p \leq (4mt_{\mathrm{shld}})^{-1}$.*

Therefore, given upper bound N and Δ of n and δ respectively, we get a $(O(m\Delta^2 N^3 \log N), \Omega(Ne^N))$ loosely-stabilizing leader election protocol for arbitrary graphs by assigning $t_{\max} = 8N\Delta$, $t_{\mathrm{virus}} = t_{\max}/2$, $t_{\mathrm{shld}} = 2\Delta t_{\max}\lceil \log N \rceil$ and $p = (4N^2 t_{\mathrm{shld}})^{-1}$.

5 Conclusion

We have presented two loosely-stabilizing leader election protocols for arbitrary undirected graphs in the PP model: one works with agent-identifiers and the other works with random numbers. Both protocols keep a unique leader for an exponentially long expected time after entering a loosely-safe configuration. The protocols use only upper bounds N of n and Δ of δ while any self-stabilizing leader election protocol needs the exact knowledge of n. The restriction of the protocols to *undirected* graph is only for simplicity of protocol description and

complexity analysis. The proposed protocols also work on arbitrary *directed* graphs with slight modification: it is only necessary that a responder also executes some actions of an initiator (Line 1 of Protocol 1 and Lines 10-12 of Protocol 2). Our future work is to develop a loosely-stabilizing leader election protocol without agent-identifiers or random numbers for arbitrary graphs. We will also tackle with loosely-stabilizing leader election for some classes of graphs (e.g. rings and trees).

References

1. 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.
2. D. Angluin, J. Aspnes, and D. Eisenstat. fast computation by population protocols with a leader. In *DISC*, pages 61–75, 2006.
3. 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.
4. J. Beauquier, P. Blanchard, and J. Burman. Self-stabilizing leader election in population protocols over arbitrary communication graphs. In *OPODIS*, pages 38–52, 2013.
5. J. Beauquier, J. Burman, L. Rosaz, and B. Rozoy. Non-deterministic population protocols. In *OPODIS*, pages 61–75, 2012.
6. 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.
7. D. Canepa and M. G. Potop-Butucaru. Stabilizing leader election in population protocols. 2007. http://hal.inria.fr/inria-00166632.
8. M. J. Fischer and H. Jiang. Self-stabilizing leader election in networks of finite-state anonymous agents. In *OPODIS*, pages 395–409, 2006.
9. R. Guerraoui and E. Ruppert. Even small birds are unique: Population protocols with identifiers. *Rapport de Recherche CSE-2007-04, Department of Computer Science and Engineering, York University, York, ON, Canada*, 2007.
10. O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Mediated population protocols. *Theoretical Computer Science*, 412(22):2434–2450, 2011.
11. M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005.
12. 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.
13. 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.
14. X. Xu, Y. Yamauchi, S. Kijima, and M. Yamashita. Space complexity of self-stabilizing leader election in population protocol based on k-interaction. In *SSS*, pages 86–97, 2013.

Table 1. Notations and Assumptions for P_{ID}

Notations	
τ :	$t_{\max}/(8\delta)$
$v_{\min}$:	$\mathrm{argmin}_{v \in V}\, \mathrm{id}(v)$
ID :	$\{\mathrm{id}(v) \mid v \in V\}$
$\mathcal{C}_{\mathrm{lid}}$:	$\{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{ID}}) \mid \forall v \in V, C(v).\mathtt{lid} \in \mathrm{ID}\}$
$\mathcal{L}_{\mathrm{lid}}$:	$\mathcal{C}_{\mathrm{lid}} \cap \{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{ID}}) \mid C(v_{\min}).\mathtt{lid} = \mathrm{id}(v_{\min}) \wedge C(v_{\min}).\mathtt{timer} = t_{\max}\}$
$\mathcal{S}_{\mathrm{id}}$:	$\{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{ID}}) \mid \forall v \in V,\ C(v).\mathtt{lid} = \mathrm{id}(v_{\min}) \wedge C(v).\mathtt{timer} > t_{\max}/2 \wedge C(v_{\min}).\mathtt{timer} = t_{\max}\}$
Assumptions	
	$t_{\max} \geq 8\delta \max(d,\ 2 + \log n)$

Appendix

The appendix presents the proofs of the lemmas that have only proof sketches in the main body of this paper. Specifically, we presents the complete proofs of Lemmas 6 and 17 in Sections A and B respectively. In Section C, we show the three variants of Chernoff bounds [11] used in several proofs of this paper. The appendix also presents Tables 1 and 2 that summarize notations and assumptions for protocols P_{ID} and P_{RD} respectively.

A Proofs in Leader Election with Identifiers

Lemma 6 $\max_{C \in \mathcal{L}_{\mathrm{lid}}} \mathrm{ECT}_{P_{\mathrm{ID}}}(C, \mathcal{S}_{\mathrm{id}}) = O(m\tau)$.

Proof Let $C_0 \in \mathcal{L}_{\mathrm{lid}}$ and $\Xi_{P_{\mathrm{ID}}}(C_0) = C_0, C_1, \ldots$. Since $C_t \in \mathcal{L}_{\mathrm{lid}}$ holds for every $t \geq 0$, identifier $\mathrm{id}(v_{\min})$ is the smallest among $\mathtt{lid}$ of all the agents in any configuration $C_0, C_1, \ldots$. Hence, once agent v satisfies $v.\mathtt{lid} = \mathrm{id}(v_{\min})$, then $v.\mathtt{lid} = \mathrm{id}(v_{\min})$ always holds until a timeout occurs at v. Lemma 2 has shown that, with probability at least $1 - 2ne^{-\tau}$, every agent v satisfies $v.\mathtt{lid} = \mathrm{id}(v_{\min})$ within $2m\tau$ interactions, and after that, keeps on satisfying $v.\mathtt{timer} > t_{\max}/2$ at least until $\Gamma_{2m\tau-1}$ finishes. Thus, the probability that $C_{2m\tau}(v).\mathtt{lid} = \mathrm{id}(v_{\min})$ and $C_{2m\tau}(v).\mathtt{timer} > t_{\max}/2$ hold for all $v \in V$ is at least $1 - 2ne^{-\tau}$. Note that $C_{2m\tau}(v_{\min}).\mathtt{timer} = t_{\max}$ holds with probability 1. Hence, we have

$$\max_{C \in \mathcal{L}_{\mathrm{lid}}} \mathrm{ECT}_{P_{\mathrm{ID}}}(C, \mathcal{S}_{\mathrm{id}}) \leq 2m\tau + 2ne^{-\tau} \cdot \max_{C \in \mathcal{L}_{\mathrm{lid}}} \mathrm{ECT}_{P_{\mathrm{ID}}}(C, \mathcal{S}_{\mathrm{id}}).$$

Solving this inequality gives $\max_{C \in \mathcal{L}_{\mathrm{lid}}} \mathrm{ECT}_{P_{\mathrm{ID}}}(C, \mathcal{S}_{\mathrm{id}}) \in O(m\tau)$. □

B Proofs in Leader Election with Random Numbers

We show Lemmas B2, B3 and B5 to prove Lemma 17. Lemma B2 (B3, B5) gives the lower bound of the probability that the population enters from $\mathcal{C}_{\mathrm{all}}(P_{\mathrm{RD}})$

Table 2. Notations and Assumptions for P_{RD}

Notations	
τ :	$t_{\max}/(8\delta)$
$\#_{\mathtt{L}}(C)$:	the number of leaders in configuration C
$\mathcal{G}_{\mathrm{half}}$:	$\{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{RD}}) \mid \exists v \in V,\ C(v).\mathtt{DoA} = \mathrm{ALIVE} \wedge C(v).\mathtt{timer_S} > t_{\mathrm{shld}}/2\}$
$\mathcal{V}_{\mathrm{clean}}$:	$\{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{RD}}) \mid \forall v \in V,\ C(v).\mathtt{timer_V} = 0\}$
$\mathcal{L}_{\mathrm{half}}$:	$\{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{RD}}) \mid \#_{\mathtt{L}}(C) \geq 1 \ \wedge\ \forall v \in V,\ C(v).\mathtt{timer_L} > t_{\max}/2\}$
$\mathcal{L}_{\mathrm{one}}$:	$\{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{RD}}) \mid \#_{\mathtt{L}}(C) = 1\}$
$\mathcal{S}_{\mathrm{RD}}$:	$(\mathcal{G}_{\mathrm{half}} \cup \mathcal{V}_{\mathrm{clean}}) \cap \mathcal{L}_{\mathrm{half}} \cap \mathcal{L}_{\mathrm{one}}$
$\mathcal{N}_{\mathrm{oVG}}$:	$\{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{RD}}) \mid \forall v \in V,\ C(v).\mathtt{timer_V} = C(v).\mathtt{timer_S} = 0\}$
$\mathcal{L}$:	$\{C \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{RD}}) \mid \#_{\mathtt{L}}(C) \geq 1\}$
Assumptions	
	$t_{\max} \geq 8\delta \max(d,\ 2 + \log(13n\delta\lceil \log n \rceil))$
	$t_{\mathrm{virus}} = t_{\max}/2$
	$t_{\mathrm{shld}} \geq 2\delta t_{\max}\lceil \log n \rceil$
	$p \leq (4mt_{\mathrm{shld}})^{-1}$

into $\mathcal{N}_{\mathrm{oVG}}$ (from $\mathcal{N}_{\mathrm{oVG}}$ into $\mathcal{N}_{\mathrm{oVG}} \cap \mathcal{L}$, from $\mathcal{N}_{\mathrm{oVG}} \cap \mathcal{L}$ into $\mathcal{S}_{\mathrm{RD}}$, respectively) within a certain number of interactions. We also show Lemmas B1 and B4 to prove Lemmas B2 and B5 respectively.

Lemma B1 *The probability that every $v \in V$ joins more than t_{shld} interactions as an initiator among $\Gamma_0, \ldots, \Gamma_{2mt_{\mathrm{shld}}}$ is at least $1 - ne^{-t_{\mathrm{shld}}/4}$.*

Proof For any $v \in V$ and $t \geq 0$, v joins interaction Γ_t as an initiator with probability at least $1/m$. Thus, the number of interactions v joins during the $2mt_{\mathrm{shld}}$ interactions is bounded from below by binomial random variable $X \sim B(2mt_{\mathrm{shld}}, 1/m)$. Applying the Chernoff bound of Lemma C2, we have

$$\begin{aligned}\Pr(X \leq t_{\mathrm{shld}}) &= \Pr(X \leq \mathbf{E}[X]/2) \\ &\leq e^{-\mathbf{E}[X]/8} \qquad \text{(By Chernoff Bound of Lemma C2)} \\ &= e^{-t_{\mathrm{shld}}/4}.\end{aligned}$$

Summing up the probabilities for all $v \in V$ gives the lemma. □

Lemma B2 *Let $C_0 \in \mathcal{C}_{\mathrm{all}}(P_{\mathrm{RD}})$ and $\Xi_{P_{\mathrm{RD}}}(C_0) = C_0, C_1, \ldots$. Then, we have $\Pr(C_{2mt_{\mathrm{shld}}} \in \mathcal{N}_{\mathrm{oVG}}) \geq 1 - 2ne^{-\delta\tau} - 2mt_{\mathrm{shld}} \cdot p$.*

Proof First, we show that $C_{2mt_{\mathrm{shld}}} \in \mathcal{N}_{\mathrm{oVG}}$ holds when the following three conditions hold:

(A) every agent $v \in V$ joins more than t_{shld} interactions as an initiator among $\Gamma_0, \ldots, \Gamma_{2mt_{\mathrm{shld}}}$,
(B) $\mathrm{RT}(t_{\mathrm{virus}}) \leq 2mt_{\mathrm{shld}}$, and
(C) no new virus is created during $\Gamma_0, \ldots, \Gamma_{2mt_{\mathrm{shld}}}$.

Until a new virus is created, variable $v.\texttt{timer}_\texttt{S}$ for each $v \in V$ is monotonically non-increasing and it decreases by one every time v interacts as an initiator. Hence, no agent wears a shield in configuration $C_{2mt_{\rm shld}}$ by (A) and (C). Until a new virus is created, the maximum value of all $v.\texttt{timer}_\texttt{V}$ (i.e. $\max_{v \in V} v.\texttt{timer}_\texttt{V}$) is monotonically non-increasing during $\Gamma_0, \ldots, \Gamma_{2mt_{\rm shld}}$ and it decreases at least by one in each round. Hence, no agent has a virus in configuration $C_{2mt_{\rm shld}}$ by (B) and (C). Thus, we have $C_{2mt_{\rm shld}} \in \mathcal{N}_{\rm oVG}$ when (A),(B) and (C) hold.

Next we give lower bounds on probability of (A),(B) and (C). The probability of (A) is at least $1 - ne^{-t_{\rm shld}/4} > 1 - ne^{-\delta\tau}$ from Lemma $B1$. The probability of (B) is at least $1 - ne^{-t_{\rm virus}/4} = 1 - ne^{-\delta\tau}$ from Lemma 12. At each interaction, a new virus is created with probability at most p (with probability exact p when a leader interacts as an initiator and with probability 0 otherwise). Hence, the probability of (C) is at least $1 - 2mt_{\rm shld} \cdot p$. Thus, Conditions (A),(B) and (C) hold with probability at least $1 - 2ne^{-\delta\tau} - 2mt_{\rm shld} \cdot p$. □

Lemma B3 *Let $C_0 \in \mathcal{N}_{\rm oVG}$ and $\Xi_{P_{\rm RD}}(C_0) = C_0, C_1, \ldots$. Then, we have* $\Pr(\exists i \in [0, 16m\delta\tau\lceil \log n \rceil],\ C_i \in \mathcal{N}_{\rm oVG} \cap \mathcal{L}) \geq 1 - 2ne^{-\delta\tau}$.

Proof The lemma trivially holds if C_0 has one or more leaders. Therefore, we consider the case C_0 does not have any leader (i.e. $C_0 \notin \mathcal{L}$). Since followers never create viruses or shields, there exists neither a virus nor a shield until a leader is created. Therefore, the population reaches a configuration of $\mathcal{N}_{\rm oVG} \cap \mathcal{L}$ at the first timeout of execution $\Xi_{P_{\rm RD}}(C_0) = C_0, C_1, \ldots$.

Thus, it suffices to show that a timeout occurs within $16m\delta\tau\lceil \log n \rceil$ interactions with probability at least $1 - 2ne^{-\delta\tau}$. During the period no leader exists, the maximum value of all $v.\texttt{timer}_\texttt{L}$ (i.e. $\max_{v \in V} v.\texttt{timer}_\texttt{L}$) is monotonically non-increasing and decreases at least by one in each round. This means a timeout occurs until $t_{\max}$ rounds finish. By Lemma 12, we have $\Pr(\mathrm{RT}(t_{\max}) < 16m\delta\tau\lceil \log n \rceil) \geq 1 - ne^{-t_{\max}/4} = 1 - ne^{-2\delta\tau}$. □

Lemma B4 *Let $C_0 \in \mathcal{C}_{\rm all}(P_{\rm RD})$ and $\Xi_{P_{\rm RD}}(C_0) = C_0, C_1, \ldots$. Let $t_{\rm init}$ be the maximum value of all $v.\texttt{timer}_\texttt{V}$ in C_0 (i.e. $\max_{v \in V} C(v).\texttt{timer}_\texttt{V}$). Then, we have* $\Pr(\forall v \in V,\ C_{2m\tau}(v).\texttt{timer}_\texttt{V} > t_{\rm init} - t_{\max}/2) > 1 - 2ne^{-\tau}$.

Proof The same argument as the proof of Lemma 2 gives the lemma. □

Lemma B5 *Let $C_0 \in \mathcal{N}_{\rm oVG} \cap \mathcal{L}$ and $\Xi_{P_{\rm RD}}(C_0) = C_0, C_1, \ldots$. Then, we have* $\Pr(\exists i \in [0, \lceil 2mp^{-1} \rceil + 2m\tau],\ C_i \in \mathcal{S}_{\rm RD}) \geq 1 - e^{-2} - 5ne^{-\tau} - 2m\tau \cdot p$.

Proof Let t be the minimum integer (i.e. the first time) such that configuration C_t has a virus. During the period one or more leaders exist, each interaction makes a new virus with probability at least p/m. Hence, the probability of $t < \lceil 2mp^{-1} \rceil$ is at least $1 - (1 - p/m)^{\lceil 2mp^{-1} \rceil} > 1 - e^{-2}$.

Therefore, it suffices to show that $C_{t+2m\tau} \in \mathcal{S}_{\rm RD}$ holds with probability at least $1 - 5ne^{-\tau} - 2m\tau \cdot p$. We denote the leader that creates a virus at interaction Γ_{t-1} by v. Note that, in configuration C_t, only v has a virus and a shield while the other agents do not have viruses or shields. Furthermore, the virus and the

shield of v have the maximum TTL (t_{virus} and t_{shld} respectively in C_t. We have $C_{t+2m\tau} \in \mathcal{S}_{\text{RD}}$ if all the following conditions hold:

(A) every agent has a virus in $C_{t+2m\tau}$,
(B) every agent except for v does not wear a shield in $C_{t+2m\tau}$,
(C) agent v joins only less than $t_{\text{shld}}/2$ interactions as an initiator during $\Gamma_t, \Gamma_{t+1}, \ldots, \Gamma_{t+2m\tau-1}$, and
(D) every agent has $\texttt{timer}_{\texttt{L}}$ larger than $t_{\max}/2$ in $C_{t+2m\tau}$.

By (A) and (B), all agents except for v are dead in $C_{t+2m\tau}$. By (C), v always has a shield larger than $t_{\text{shld}}/2$ during the $2m\tau$ interactions, and hence, is alive (i.e. is a leader) in $C_{t+2m\tau}$. Therefore, $C_{t+2m\tau} \in \mathcal{L}_{\text{one}} \cap \mathcal{G}_{\text{half}}$ holds. Moreover, $C_{t+2m\tau} \in \mathcal{L}_{\text{half}}$ holds by (D). Thus, we have $C_{t+2m\tau} \in \mathcal{L}_{\text{one}} \cap \mathcal{G}_{\text{half}} \cap \mathcal{L}_{\text{half}} \subset \mathcal{S}_{\text{RD}}$ when (A),(B),(C) and (D) hold.

Therefore, it suffices to show that all (A),(B),(C) and (D) hold with probability $1 - 5ne^{-\tau} - 2m\tau \cdot p$. Since $C_t(v).\texttt{timer}_{\texttt{V}} = t_{\text{virus}} = t_{\max}/2$, the probability of (A) is at least $1 - 2ne^{-\tau}$ by Lemma B4. The sufficient condition of (B) is that a new virus is not created during $\Gamma_t, \Gamma_{t+1}, \ldots, \Gamma_{t+2m\tau-1}$. The probability of this condition is at least $1 - 2m\tau \cdot p$. The probability of (C) is at least $1 - ne^{-\delta\tau} > 1 - ne^{-\tau}$ by Lemma 9. Finally, The probability of (D) is at least $1 - 2ne^{-\tau}$ by Lemma 14. Thus, all (A), (B), (C) and (D) hold with probability at least $1 - 5ne^{-\tau} - 2m\tau \cdot p$. □

Lemma 17 $\max_{C \in \mathcal{C}_{\text{all}}(P_{\text{RD}})} \text{ECT}_{P_{\text{RD}}}(C, \mathcal{S}_{\text{RD}}) = O(mp^{-1})$.

Proof By Lemmas B2, B3 and B5, starting from any configuration of $\mathcal{C}_{\text{all}}(P_{\text{RD}})$, the population enters a configuration of $\mathcal{S}_{\text{RD}}$ within $2mt_{\text{shld}} + 16m\delta\tau\lceil \log n \rceil + \lceil 2mp^{-1} \rceil + 2m\tau$ interactions with probability at least $1 - 2ne^{-\delta\tau} - 2mt_{\text{shld}} \cdot p - 2ne^{-\delta\tau} - e^{-2} - 5ne^{-\tau} - 2m\tau \cdot p$. The former expression is at most $\lceil (2m+1) \cdot p^{-1} \rceil$ and the latter expression is at least $1 - 3mt_{\text{shld}} \cdot p - 6ne^{-\tau} - e^{-2} > 1 - 3/4 - 6e^{-2}/26 - e^{-2} > 0.08$. Hence, we have

$$\max_{C \in \mathcal{C}_{\text{all}}(P_{\text{RD}})} \text{ECT}_{P_{\text{RD}}}(C, \mathcal{S}_{\text{RD}}) \leq \lceil (2m+1)p^{-1} \rceil + 0.92 \cdot \max_{C \in \mathcal{C}_{\text{all}}(P_{\text{RD}})} \text{ECT}_{P_{\text{RD}}}(C, \mathcal{S}_{\text{RD}}).$$

Solving this inequality gives $\max_{C \in \mathcal{C}_{\text{all}}(P_{\text{RD}})} \text{ECT}_{P_{\text{RD}}}(C, \mathcal{S}_{\text{RD}}) = O(mp^{-1})$. □

C Chernoff Bounds

Lemma C1 (from Eq. (4.2) in [11]) *The following inequality holds for any binomial random variable X:*

$$\Pr(X \geq 2\mathbf{E}[X]) \leq e^{-\mathbf{E}[X]/3}.$$

Lemma C2 (from Eq. (4.5) in [11]) *The following inequality holds for any binomial random variable X:*

$$\Pr(X \leq \mathbf{E}[X]/2) \leq e^{-\mathbf{E}[X]/8}.$$

Lemma C3 (from Eq. (4.5) in [11]) *The following inequality holds for any binomial random variable X:*

$$\Pr(X \leq \mathbf{E}[X]/4) \leq e^{-9\mathbf{E}[X]/32}.$$