



HAL
open science

Optimal Time and Space Leader Election in Population Protocols

Petra Berenbrink, George Giakkoupis, Peter Kling

► **To cite this version:**

Petra Berenbrink, George Giakkoupis, Peter Kling. Optimal Time and Space Leader Election in Population Protocols. STOC 2020 - 52nd Annual ACM Symposium on Theory of Computing, Jun 2020, Chicago, United States. pp.1-29, 10.1145/3357713.3384312 . hal-02545348

HAL Id: hal-02545348

<https://inria.hal.science/hal-02545348v1>

Submitted on 17 Apr 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Optimal Time and Space Leader Election in Population Protocols

Petra Berenbrink
Universität Hamburg
Hamburg, Germany
petra.berenbrink@uni-hamburg.de

George Giakkoupis
Inria, Univ Rennes, CNRS, IRISA
Rennes, France
george.giakkoupis@inria.fr

Peter Kling
Universität Hamburg
Hamburg, Germany
peter.kling@uni-hamburg.de

ABSTRACT

Population protocols are a model of distributed computing, where n agents with limited computational power and memory perform randomly scheduled pairwise interactions. A fundamental problem in this setting is that of *leader election*, where all agents start from the same state, and they seek to reach and maintain a global state where exactly one agent is in a dedicated leader state.

A significant amount of work has been devoted to the study of the time and space complexity of this problem. Alistarh et al. (SODA'17) have shown that $\Omega(\log \log n)$ states per agent are needed in order to elect a leader in fewer than $\Theta(n^2)$ expected interactions. Moreover, $\Omega(n \log n)$ expected interactions are required regardless of the number of states (Sudo and Masuzawa, 2019). On the upper bound side, Gasieniec and Stachowiak (SODA'18) have presented the first protocol that uses an optimal, $\Theta(\log \log n)$, number of states and elects a leader in $O(n \log^2 n)$ expected interactions. This running time was subsequently improved to $O(n \log n \log \log n)$ (Gasieniec et al., SPAA'19).

In this paper we provide the first leader election population protocol that is both time and space optimal: it uses $\Theta(\log \log n)$ states per agent, and elects a leader in $O(n \log n)$ interactions in expectation. A key novel component of our approach is a simple protocol that efficiently selects a small set of agents, of $\text{poly}(\log n)$ size, given an initial set of $s = O(n^\epsilon)$ selected agents. Unlike existing approaches, which proceed by shrinking the initial set monotonically over time, our protocol first increases the set in a controlled way to a specific size (which is independent of s), before it shrinks the set to a $\text{poly}(\log n)$ size.

CCS CONCEPTS

• **Mathematics of computing** → **Stochastic processes**; • **Theory of computation** → **Self-organization**.

KEYWORDS

population protocols, leader election, stabilization time

1 INTRODUCTION

Population protocols, originally introduced in a seminal paper by Angluin et al. [6], are a popular model for distributed computing. They are often used for systems where global tasks are performed collectively by a massive number of agents with limited computation and communication power. Examples of such systems include sensor networks, chemical reaction networks [15, 18], animal populations, and gene regulatory networks [14].

A population protocol is typically defined for a system of n agents, which are identical finite state machines. In the standard probabilistic model which we adopt, in each time step a random pair of agents is chosen to interact with each other. In this *interaction*, the two

agents observe each other's state, and update their respective state according to a given *transition function*. The protocol terminates, or *stabilizes*, when it reaches a *configuration* from which the *output* of each agent cannot change. Section 2 gives the formal definitions.

We study the problem of *leader election*, one of the two most intensively studied problems in population protocols. (The other is *majority consensus*, and is discussed in the related work.) In leader election, each agent starts from the same state and the protocol must, eventually, stabilize to a configuration where exactly one agent is in a *leader* state, and all other agents are in *follower* states. As is true for distributed computing in general, leader election is one of the most fundamental primitives, often used as a symmetry breaking mechanism by other protocols. In particular, Angluin et al. [7] have shown that given a leader, population protocols with a constant number of states per agent can efficiently compute any semilinear predicate; it is still an open question whether efficient time protocols exist for this problem if no leader is given.

A lot of work has been devoted recently to the study of the time and space complexity of leader election in population protocols, producing both efficient algorithms and lower bounds. In a breakthrough work, Doty and Soloveichik [21] showed that leader election using a constant number of states requires in expectation $\Omega(n^2)$ interactions.¹ Alistarh et al. [1] generalized this lower bound by showing that protocols with fewer than $0.5 \log \log n$ states need an expected number of $\Omega(n^2 / \text{poly}(\log n))$ interactions to elect a leader. In a recent preprint, Sudo and Masuzawa [29] showed that any leader election population protocol requires $\Omega(n \log n)$ interactions to stabilize, both in expectation and w.h.p.,² regardless of the number of states per agent.

On the upper bound side, Alistarh and Gelashvili [3] presented an elegant, tournament based leader election protocol that stabilizes in $O(n \log^3 n)$ interactions w.h.p., and uses $O(\log^3 n)$ states. Bilke et al. [13] reduced the cubic terms in the time and space complexity to quadratic terms. The space complexity was further improved to $O(\log n)$ states per agent [2, 11]. In a breakthrough result, Gasieniec and Stachowiak [24] introduced the first leader election protocol with sub-logarithmic space complexity and running time $n \cdot \text{poly}(\log n)$. Precisely, their protocol stabilizes w.h.p. after $O(n \log^2 n)$ interactions and uses an asymptotically optimal number of states, $O(\log \log n)$. In a follow-up work, Gasieniec et al. [25] maintained the optimal space complexity while achieving an *expected* stabilization time of $O(n \log n \log \log n)$ interactions. Recently, Sudo et al. [30] proposed a protocol with optimal expected

¹The time complexity is sometimes also expressed in *parallel time*, which corresponds to the number of interactions divided by n , in order to account for the inherent parallelism of the system. To avoid confusion, we use the actual number of interactions throughout this paper.

²*W.h.p.* (with high probability) denotes probabilities of the form $1 - n^{-c}$, for a constant $c > 0$ that can be made arbitrarily large at the cost of the constants involved.

stabilization time, $O(n \log n)$, but at the cost of increasing significantly the space requirements, to $\Theta(\log n)$ states per agent.

Results & Techniques. We propose the first leader election population protocol that both uses an asymptotically optimal number of states and stabilizes in an asymptotically optimal expected number of interactions.

THEOREM 1. *There is a leader election population protocol that uses $\Theta(\log \log n)$ states, and has stabilization time $O(n \log n)$ in expectation and $O(n \log^2 n)$ w.h.p.*

The above protocol matches the lower bound of [29], which states that every leader election protocol needs $\Omega(n \log n)$ expected interactions (regardless of the number of states), and also the lower bound of [1], according to which $\Omega(\log \log n)$ states are needed to stabilize in fewer than $O(n^2 / \text{poly}(\log n))$ interactions. Our protocol also improves over [24, 25], which use $\Theta(\log \log n)$ states but have suboptimal expected stabilization time. Similar to [25], our protocol requires an estimation of $\log \log n$ within a constant additive error.

Our algorithm combines several known techniques, mainly from [24, 25], with some new ideas. A key novel component is a simple protocol that efficiently selects (with probability $1 - 1/\text{poly}(\log n)$) a small ($\text{poly}(\log n)$ size) subset of the agents, starting from an initial set of $s = O(n^{1/2})$ agents. Most existing approaches proceed by shrinking the initial set monotonically over time. In contrast, our protocol first increases the initial set to a particular size, independent of s , before shrinking it to the size of $\text{poly}(\log n)$.

Roughly speaking, this novel component works as follows. We start with $1 \leq s \leq O(\sqrt{n})$ agents in state 1, while all other agents are in state 0. State 1 spreads to agents of state 0 by a “slowed-down” one-way epidemic: whenever an agent in state 0 interacts with an agent in state 1 it changes its state from 0 to 1 with probability $1/4$. As soon as a sufficient number of agents in state 1 exists two of them will interact, and when that happens one of them changes its state from 1 to 2. Similar to state 1, state 2 spreads to agents in state 0 via a one-way epidemic, but this time with probability 1.

It is easy to see that the first agent moves to state 2 when the number of agents in state 1 is roughly $\Theta(\sqrt{n})$. From that point on (and ignoring for simplicity additional interactions between 1s), we have essentially two competing epidemics: one with $\Theta(\sqrt{n})$ initial support which spreads with a slow rate of $1/4$, and another one with initial support of 1 but with a spreading rate of 1. It is not hard to see that this idealized process results in roughly $\Theta(n^{3/4})$ agents which are in state 1s. Now, from a set of $\Theta(n^{3/4})$ agents in state 1, we can select a subset of size $\text{poly}(\log n)$ as follows: when two agents at state 1 interact, one of them switches to state 3, and when two agents at state 3 interact, one switches to state 4. After $\Theta(n \log n)$ interactions, the number of 3s is $\sqrt{n} \cdot \text{poly}(\log n)$ and the number of 4s is $\text{poly}(\log n)$ w.h.p.³

To employ the above protocol, we first need to elect a set (junta) of at most $O(\sqrt{n})$ agents. We use a junta election protocol which is conceptually similar to that from [24, 25]. The idea is that an

³Our actual protocol is bit more involved (and used different state names), but the above description conveys its main ideas. Also note that there are simple variants of this protocol that work equally well, e.g., using a rate other than $1/4$ for the slower one-way epidemic. Such a variant results in a final number of agents in state 1 which is different than $n^{1/4}$, and it has to be combined with an appropriately modified mechanism to select a subset of size $\text{poly}(\log n)$ afterwards.

agent increases its level whenever it meets an agent on at least the same level. Such a protocol can be implemented with $\Theta(\log \log n)$ states and elects a number of agents that is at most $n^{1-\Omega(1)}$ [24, 25]. We use a similar approach, but employ some additional tweaks which ensure that always at least one agent is elected. This was previously achieved only w.h.p, and is critical for the correctness of our protocol. Moreover, it allows agents to reuse the $\Theta(\log \log n)$ states once the protocol has finished (which our agents can realize locally). This part of our protocol elects $O(\sqrt{n})$ agents.⁴

The elected junta is used in the first protocol described above, and is also used to drive a phase clock that synchronizes the agents’ actions. Our phase clock protocol is identical to that in [24]. It consists of two clocks, one with phase interval $\Theta(n \log n)$, and a second with phase interval $\Theta(n \log^2 n)$.

The final part of our protocol reduces the number of selected leader candidates from $\text{poly}(\log n)$ to 1. Thanks to the $\Theta(\log \log n)$ number of states, we can generate coins of success probability $1/\text{poly}(\log n)$, which allow us to reduce the number of leader candidates to expectedly $O(1)$, after $O(1)$ “broadcast rounds,” each taking $\Theta(n \log n)$ interaction. After that, binary coins are used for at most $\Theta(\log n)$ additional broadcast rounds. This component is carefully designed to ensure that not all agents are eliminated, and to guarantee a fast, expected stabilization time. For that we use a “reviving” elimination process that is conceptually similar to ideas from [24, 25], but more efficient, requiring only $O(n \log n)$ expected interactions.

Finally, similar to existing algorithms, we use a fall-back mechanism employing the slow stable elimination protocol from [8]. However, the precise way this protocol is employed is different. E.g., we rely on the fact that there is always at least one agent driving the phase clock. Thus the clocks may get desynchronized but all clocks will eventually reach their maximum value.

Other Related Work. A prominent problem that has been intensively studied in population protocols, besides leader election, is the majority consensus. Each agent starts with one of two (or, in general, k) opinions, and the goal is to reach a configuration where each agent agrees on the correct majority (or, in general, plurality) among the initial opinions. The problem was first studied in [6], and an elegant protocol for approximate majority was proposed in [8]. Algorithms and time-space tradeoffs for exact and approximate majority have been studied more recently in [1, 2, 5, 10, 13]. Other problems have also been considered in population protocols, including approximate counting [12, 19, 20], and controlling the population size [26].

Our protocol uses a junta-driven phase clock introduced in [24]. A leaderless phase clock was proposed in [2], and related oscillators were studied in [16, 22]. Kosowski and Uznanski [28] have used such oscillators to implement efficient leader election and majority protocols with constant number of states, which are not always correct, but may fail with a small probability.

⁴We note that one needs to know the value of n within $\text{poly}(\log n)$ factors to estimate the final junta size within the same $\text{poly}(\log n)$ factor precision, whereas our protocol only knows n within polynomial factors (it knows $\lceil \log \log n \rceil + O(1)$). Moreover, even if we are given the precise value of n , it is not clear how to augment the junta election protocol to elect a junta of a given size, e.g., $\text{poly}(\log n)$ or $n^{3/4} \cdot \text{poly}(\log n)$, under the constraint that the state space size is $\Theta(\log \log n)$, as achieved by the first protocol we presented above.

For a general overview of results for population protocols, see the two recent surveys by Alistarh and Gelashvili [4] and Elsässer and Radzik [23].

2 MODEL & NOTATION

We study population protocols defined on a population of n agents, which are identical finite-state machines. At any time, the global system state is described by a vector c , called a *configuration*, whose i -th component gives the number of agents in the i -th state. We assume the classic model [6, 8], in which a random scheduler sequentially matches agents in independently and uniformly chosen pairs (u, v) . Agent u is the *initiator* and agent v is the *responder*. The two agents perform an *interaction*, in which u observes v 's state and updates its own state according to a deterministic *transition function*, which is given as a set of *transitions*

$$\text{initiatorState} + \text{responderState} \rightarrow \text{newInitiatorState}.$$

For *state values* we use ‘ \perp ’, numbers, and names in typewriter font. To emphasize that a state *variable* s represents the state of agent u , we write $s(u)$. We allow transition rules that use a small amount of randomness (constant many, fair coin tosses). This is w.l.o.g., as such coin tosses can be simulated from the randomness of the scheduler, using so-called synthetic coins [1].

Starting from a set of *initial configurations*, agents try to reach and stay in a set of *correct configurations*. For *leader election*, the set of initial configurations consists of one configuration, with all agents in the same, initial state. The set of correct configurations contains all configurations in which exactly one agent is in one of, possibly, several *leader states*, and all other agents are in non-leader states. Agents are not required to realize when a correct configuration has been reached.

Main Protocol & External Transitions. Our main leader election protocol, called LE, is formed by multiple subprotocols that run in parallel. Those subprotocols are described in the following sections. To facilitate a modular description of the subprotocols and their interaction, we introduce the notion of *external transitions*

$$\text{old} \Rightarrow \text{new} \quad \text{if } \text{condition}.$$

This means that, *after all normal transitions of the interaction are completed*, if the initiator is in state *old* and *condition* is satisfied then the initiator’s state changes to *new*. The *condition* must depend only on the initiator’s own state and typically refers to a state change caused by another subprotocol. In our analysis, we often use the notion of *steps* instead of interactions. A step consists of an interaction (possibly triggering multiple normal transitions) followed by all external transitions triggered by the state changes.

Complexity Measures. Protocol complexity is measured in terms of the *number of states per agent*, and the *stabilization time*. To define stabilization time, we say that a correct configuration c is *stable* if any configuration reachable from c with non-zero probability is also correct. Then the *stabilization time* is the earliest step when a stable correct configuration is reached.⁵

⁵Some previous works, e.g., [1, 13], refer to stabilization time as *convergence time*. We avoid using the latter term, because it is often used to describe a weaker notion of termination, namely, the earliest step after which all future configurations (in the given execution of the protocol) are correct.

3 JUNTA ELECTION

We use two junta election protocols inspired by a similar protocol from [24]. The first one, JE1, elects a junta of size at most $n^{1-\epsilon}$. This junta is used to drive the phase clock LSC, described in Section 4. The second protocol, JE2, further reduces the size of the junta to $O(\sqrt{n \ln n})$, and the resulting junta is used by DES, in Section 5.1.

3.1 Junta Election 1 (JE1)

The state space of JE1 is $\mathcal{S}_{\text{JE1}} := \{-\psi, -\psi + 1, \dots, \varphi_1\} \cup \{\perp\}$, where $\psi := 3 \log \log n$ and $\varphi_1 := \log \log n - \log \log \log n - 3$. We refer to an agent in state $\ell \in \mathcal{S}_{\text{JE1}}$ as being on *level* ℓ . Initially, every agent is on level $-\psi$. First the agents try to reach level 0 by tossing a series of fair coins: on success they increase their level, on failure they reset their level to $-\psi$. Once an agent reached level 0, its level cannot decrease any longer. Now the agent increases its level ℓ whenever it interacts with an agent on a level in $\{\ell, \dots, \varphi_1 - 1\}$. As soon as an agent not yet on level φ_1 interacts with an agent on level φ_1 or in special state \perp , it switches to state \perp . Protocol 1 gives the formal transition rules.

Protocol 1: JE1

$$\begin{array}{ll} \ell + \ell' \rightarrow \begin{cases} \ell + 1 & \text{w.pr. } 1/2 \\ -\psi & \text{w.pr. } 1/2 \end{cases} & \text{if } -\psi \leq \ell < 0 \text{ and } \ell' \notin \{\varphi_1, \perp\} \\ \ell + \ell' \rightarrow \ell + 1 & \text{if } 0 \leq \ell \leq \ell' \text{ and } \ell' \notin \{\varphi_1, \perp\} \\ \ell + \ell' \rightarrow \perp & \text{if } \ell \neq \varphi_1 \text{ and } \ell' \in \{\varphi_1, \perp\} \end{array}$$

We say JE1 is *completed* when every agent is in state φ_1 or \perp . The agent is *elected* in JE1 if it is in φ_1 , and is *rejected* if it is in \perp .

The intuition for the protocol is as follows. The mechanism using the coin tosses achieves that only a fraction $x_0 = 1/\text{poly}(\log n)$ of agents reach level 0 within $O(n \log n)$ interaction. Moreover, the fraction x_ℓ of agents that reach level $\ell > 0$ within $O(n \log n)$ interaction, is roughly $x_{\ell-1}^2$. Thus, x_{φ_1} is roughly $x_0^{2^{\varphi_1}}$, which is $n^{-\epsilon}$ for our choice of ψ and φ_1 .

The next lemma lists the main properties of JE1. Its proof can be found in Appendix B.

LEMMA 2.

- (a) *At least one agent is elected in JE1.*
- (b) *W.h.p. at most $n^{1-\epsilon}$ agents are elected in JE1, for some constant $\epsilon > 0$.*
- (c) *JE1 is completed in $O(n \log n)$ steps w.h.p., and this is true even if all agents start from an arbitrary state.*

3.2 Junta Election 2 (JE2)

The state space of JE2 is $\mathcal{S}_{\text{JE2}} := \{\text{id1}, \text{act}, \text{inact}\} \times \{0, 1, \dots, \varphi_2\}$, for a large enough constant $\varphi_2 \in \mathbb{N}$. We refer to an agent in state $(d, \ell) \in \mathcal{S}_{\text{JE2}}$ as being on *level* ℓ , and as either *idle* ($d = \text{id1}$), *active* ($d = \text{act}$), or *inactive* ($d = \text{inact}$). Initially, all agents are idle on level 0. If an agent is elected in JE1 it becomes active, and if rejected it becomes inactive. Active agents increase their level when they interact with agents on at least the same level, until they reach level φ_2 . If an agent interacts with an agent on a lower level or reaches level φ_2 becomes inactive. Protocol 2 gives the formal transitions.

Protocol 2: JE2

$$\begin{aligned}
 (\text{id1}, 0) &\Rightarrow \begin{cases} (\text{act}, 0) & \text{if elected in JE1} \\ (\text{inact}, 0) & \text{if rejected in JE1} \end{cases} \\
 (\text{act}, \ell) + (\cdot, \ell') &\rightarrow \begin{cases} (\text{act}, \ell + 1) & \text{if } \ell \leq \ell' \text{ and } \ell < \varphi_2 - 1 \\ (\text{inact}, \varphi_2) & \text{if } \ell \leq \ell' \text{ and } \ell = \varphi_2 - 1 \\ (\text{inact}, \ell) & \text{if } \ell > \ell' \end{cases}
 \end{aligned}$$

We combine the above protocol with a one-way epidemic that propagates the maximum level encountered. Formally, in addition to state $(d, \ell) \in \mathcal{S}_{\text{JE2}}$, each agent stores a value $k \in \{0, \dots, \varphi_2\}$, which is initially 0. We will refer to k as the agent's max-level. The combined transition rule is now

$$(d, \ell, k) + (d', \ell', k') \rightarrow (d_{\text{new}}, \ell_{\text{new}}, \max\{k, k', \ell_{\text{new}}\}),$$

where $(d, \ell) + (d', \ell') \rightarrow (d_{\text{new}}, \ell_{\text{new}})$ is the corresponding transition in JE2.

We say that JE2 is *completed* when all agents are inactive, and have the same max-level component. An agent is *rejected* in JE2 if it is inactive and its level ℓ is smaller than its max-level k . We say an agent is *elected* in JE2, if JE2 is completed and the agent is not rejected (i.e., $\ell = k$). Note that an agent that is not yet rejected, cannot locally verify whether it is elected or not.

The protocol assumes that at least one and at most $n^{1-\epsilon}$ agents become active, and parameter φ_2 is a function of ϵ . The intuition for the protocol is the same as for the junta election protocol from [24], namely, that the fraction of agents reaching a level roughly squares with each level. Since only a small fraction becomes active in JE2, a constant number of levels suffices, whereas in [24], a constant fraction of agents is activated and $\Theta(\log \log n)$ levels are needed.

The next lemma lists the main properties of JE2. Its proof can be found in [Appendix C](#).

LEMMA 3.

- (a) *Not all agents are rejected in JE2.*
- (b) *For any constant $0 < \epsilon < 1$, there is a constant $\varphi_2 \in \mathbb{N}$ such that, w.p. $1 - O(1/\log n)$, if at most $n^{1-\epsilon}$ agents are elected in JE1, then at most $O(\sqrt{n \ln n})$ agents are not rejected in JE2.*
- (c) *Suppose JE1 is completed at a given step t_1 . If t_2 denotes the step when JE2 is completed, then $t_2 = t_1 + O(n \log n)$ w.h.p.*

4 PHASE CLOCK (LSC)

We use protocol LSC (for Log-Square Clock) to synchronize our subprotocols. The protocol follows closely the phase clock implementation of [24]. It consists of two clocks. The first one, called *internal*, is a modulo $2m_1 + 1$ clock and ticks every $\Theta(n \log n)$ interactions on average. The second clock, called *external*, stops when it reaches a maximum value of $2m_2$, and ticks every $\Theta(n \log^2 n)$ interactions on average. Parameters m_1, m_2 are large integer constants. The protocol requires a set of $n^{1-\Omega(1)}$ clock agents, which is provided by subprotocol JE1.

Formally, the state space is $\mathcal{S}_{\text{LSC}} := \{\text{clk}, \text{nrm}\} \times \{\text{int}, \text{ext}\} \times \{0, \dots, 2m_1\} \times \{0, \dots, 2m_2\}$. An agent in state $(s, c, t_{\text{int}}, t_{\text{ext}})$ is a *clock agent* if $s = \text{clk}$ and a *normal agent* if $s = \text{nrm}$; c indicates whether the agent updates its internal ($c = \text{int}$) or external ($c = \text{ext}$) clock in its next interaction; t_{int} is the agent's *internal clock*

counter, which increases modulo $2m_1 + 1$; and t_{ext} is the *external clock counter*, which stops when it reaches value $2m_2$. Initially, all agents are in state $(\text{nrm}, \text{int}, 0, 0)$. The transition rules are listed in [Protocol 3](#).

Protocol 3: LSC

$$\begin{aligned}
 (\text{nrm}, \text{int}, 0, 0) &\Rightarrow (\text{clk}, \text{int}, 0, 0) \quad \text{if elected in JE1} \\
 (\text{nrm}, \text{int}, i, j) + (\cdot, \cdot, i', \cdot) &\rightarrow \begin{cases} (\text{nrm}, \text{int}, i', j) & \text{if } 0 < i' - i \leq m_1 \\ (\text{nrm}, \text{ext}, i', j) & \text{if } i - i' > m_1 \end{cases} \quad (*) \\
 (\text{clk}, \text{int}, i, j) + (\cdot, \cdot, i', \cdot) &\rightarrow \begin{cases} (\text{clk}, \text{int}, i' + 1 \bmod (2m_1 + 1), j) & \text{if } 0 \leq i' - i \leq m_1 \\ (\text{clk}, \text{ext}, i' + 1, j) & \text{if } i - i' > m_1 \end{cases} \quad (*) \\
 (\text{nrm}, \text{ext}, i, j) + (\cdot, \cdot, \cdot, j') &\rightarrow (\text{nrm}, \text{int}, i, \max\{j, j'\}) \\
 (\text{clk}, \text{ext}, i, j) + (\cdot, \cdot, \cdot, j') &\rightarrow (\text{clk}, \text{int}, i, \min\{2m_2, \max\{j, j' + 1\}\})
 \end{aligned}$$

Note that as long as no clock agent exists, no normal transitions are triggered. An agent in the initial state becomes a clock agent via an external transition as soon as it is elected in JE1.

We say that an agent is in *internal phase* $\rho \in \mathbb{N}_0$, if its internal clock counter has passed ρ times through zero, i.e., the agent has executed (as an initiator) ρ times the transitions marked by (*). We also say the agent is in *external phase* $\rho' \in \{0, 1, 2\}$ if $\lfloor j/m_2 \rfloor = \rho'$, where j is the node's external clock counter.

In addition to its state in \mathcal{S}_{LSC} , a agent stores *partial* information about its internal phase. Precisely, it maintains two variables. The first is *iphase* $\in \{0, \dots, \nu\}$, and stores the agent's current internal phase, up to a maximum value of $\nu = \Theta(\log \log n)$. Once *iphase* = ν , the variable stops increasing. The second variable is *parity* $\in \{0, 1\}$, and stores the parity of the internal phase. These two variables increase the size of the clock-related state of the agent to $\Theta(\log \log n)$, from $O(1)$. For convenience, we also define variable *xphase*, which stores the current external phase of an agent, and can be computed directly from the external clock counter.

For $\rho \geq 1$, let f_ρ be the step when the first agent reaches internal phase ρ , and let l_ρ be the step when the last agent reaches internal phase ρ . Also, $f_0 = l_0$ is the step when the first clock agent is created. The *length* $L_{\text{int}}(\rho)$ of internal phase ρ is $f_{\rho+1} - l_\rho$, and its *stretch* $S_{\text{int}}(\rho)$ is $f_{\rho+1} - f_\rho$. Similarly, for $\rho' \in \{1, 2\}$, $f'_{\rho'}$ is the step when the first agent reaches external phase ρ' , and $l'_{\rho'}$ is the step when the last agent reaches an external phase greater or equal to ρ' (note that the external phase of an agent may increase from 0 to 2 in a single step). Also, $f'_0 = l'_0 = f_0$. For $\rho' \in \{0, 1\}$, $L_{\text{ext}}(\rho') = f'_{\rho'+1} - l'_{\rho'}$ and $S_{\text{ext}}(\rho') = f'_{\rho'+1} - f'_{\rho'}$.

The next result is essentially a reformulation of [24, Lemma 3.7] to fit our needs. It states that agents stay in sync for at least $n \log^3 n$ steps w.h.p.

LEMMA 4. *Let $0 < \epsilon < 1$ and $c_1, c_2 > 0$ be any constants. There are constants $m_1, m_2 \in \mathbb{N}$, $d_2 \geq d_1 \geq c_1$, and $d_4 \geq d_3 \geq c_2$ such that the following statement holds w.h.p. If at most $n^{1-\epsilon}$ agents are elected in JE1, then*

- (a) $L_{\text{int}}(\rho) \geq d_1 \cdot n \log n$ and $S_{\text{int}}(\rho) \leq d_2 \cdot n \log n$, for all $\rho \in \{0, 1, \dots, \log^2 n\}$;
- (b) $L_{\text{ext}}(\rho') \geq d_3 d_2 \cdot n \log^2 n$ and $S_{\text{ext}}(\rho') \leq d_4 d_2 \cdot n \log^2 n$, for all $\rho' \in \{0, 1\}$.

The following lemma is needed to show that the leader election protocol is correct, even in the unlikely case in which agents are not synchronized.

LEMMA 5. *Suppose that there is at least one clock agent at a given step t . Then $E[L_2^t] \leq t + O(n^2 \log^3 n)$.*

The proofs of Lemmas 4 and 5 are given in Appendix D.

5 EPIDEMIC-BASED SELECTION

The two protocols presented next select a polylogarithmic number of leader candidates. The first protocol, DES, uses the junta of size $O(n^{1-\epsilon})$ elected in JE2, to select $n^{3/4} \cdot \text{poly}(\log n)$ candidates. The second protocol, SRE, reduces the candidates to $\text{poly}(\log n)$.

5.1 Dual Epidemic Selection (DES)

The state space of DES is $\mathcal{S}_{\text{DES}} := \{0, 1, 2\} \cup \{\perp\}$, and all agents are in state 0 initially. An agent switches from state 0 to 1 when its clock in LSC reaches internal phase 1, provided that the agent has not yet been rejected in JE2. Also agents switch from state 0 to 1 with probability 1/4, if they interact with an agent in state 1. When two agents in state 1 interact, one of them switches to state 2. State-2 agents cause state-0 agents to switch to 1 or \perp (each with probability 1/4).⁶ State- \perp agents cause state-0 agents to switch to \perp . Protocol 4 gives the formal transition rules.

Protocol 4: DES

$0 \Rightarrow 1$	if not rejected in JE2 and iphase = 1
$0 + 1 \rightarrow 1$	w.pr. 1/4
$1 + 1 \rightarrow 2$	
$0 + 2 \rightarrow \begin{cases} 1 \\ \perp \end{cases}$	$\begin{cases} \text{w.pr. } 1/4 \\ \text{w.pr. } 1/4 \end{cases}$
$0 + \perp \rightarrow \perp$	

We say DES is *completed* when no agents are left in state 0. An agent is *rejected* in DES if it is in state \perp . Note that after DES is completed, no new agents get rejected. We say an agent is *selected* in DES, if DES is completed and the agent is not rejected (thus, it is in state 1 or 2).

The protocol assumes that at least one and at most $O(\sqrt{n \log n})$ agents are elected in JE2. The intuition is that the first agent reaches state 2 when the number of agents at state 1 is roughly \sqrt{n} , and this takes $O(n \log n)$ steps. It takes $O(n)$ additional steps before the first agent reaches state \perp . From then on, we have essentially two competing one-way epidemics, the one spreading state 1 to state-0 agents with probability 1/4, and the other spreading state \perp to state-0 agents with probability 1. Since roughly \sqrt{n} agents are in state 1 when the first agent reaches state 2, and ignoring additional interactions between two state-2 agents (which has negligible effect), it is not hard to see that the number of agents not in state \perp eventually is on average roughly $n^{3/4}$.

The next lemma provides the main properties of DES. Its proof can be found in Appendix E.

⁶The choice of probability 1/4 is made to simplify some parts of the analysis; the deterministic rule $0 + 2 \rightarrow \perp$ works as well.

LEMMA 6.

- (a) *Not all agents are rejected in DES.*
- (b) *Suppose that JE2 is completed before step f_1 .⁷ Suppose also that at most $O(\sqrt{n \log n})$ agents are elected in JE2. Then w.pr. $1 - O(1/\log n)$, the number of agents that are not rejected in DES is at least $\Omega(n^{3/4}(\log \log n)^{1/4}(\log n)^{-3/4})$ and at most $O(n^{3/4} \log n)$.*
- (c) *Suppose that the first agent reaches state 1 at a given step t_1 . If t_2 denotes the step when DES is completed, then $t_2 = t_1 + O(n \log n)$ w.h.p.*

5.2 Square-Root Elimination (SRE)

The state space of SRE is $\mathcal{S}_{\text{SRE}} := \{o, x, y, z\} \cup \{\perp\}$, and all agents are in state o initially. An agent switches from state o to x when it reaches internal phase 2, if the agent has not yet been rejected in DES. From state x an agent switches to state y when it interacts with an agent in state x or y , and from y switches to z if it interacts with an agent in state y . As soon as some agent reaches state z , status \perp is propagated by a one-way epidemic to all agents not in state z . Protocol 5 gives the formal transition rules.

Protocol 5: SRE

$o \Rightarrow x$	if not rejected in DES and iphase = 2
$x + s \rightarrow y$	if $s \in \{x, y\}$
$y + y \rightarrow z$	
$s + s' \rightarrow \perp$	if $s \neq z$ and $s' \in \{z, \perp\}$

We say SRE is *completed* when every agent is in state z or \perp . The agent *survives* SRE if it is in z , and is *eliminated* if it is in \perp .

The intuition for this protocol is very simple. Starting from close to $n^{3/4}$ agents in state x , after $O(n \log n)$ steps there are roughly \sqrt{n} agents in state y , and after $O(n \log n)$ additional steps there are $\text{poly}(\log n)$ agents in state z .

The next lemma gives the main properties of SRE. Its proof can be found in Appendix F.

LEMMA 7.

- (a) *Not all agents are eliminated in SRE.*
- (b) *Suppose that DES is completed before step f_2 , and at most $O(n^{3/4} \log n)$ agents are selected in DES. W.pr. $1 - O(1/\log n)$ at most $O(\log^7 n)$ agents are not eliminated in SRE.*
- (c) *Suppose that DES is completed before step f_2 , and at least $\Omega(n^{3/4}(\log \log n)^{1/4}(\log n)^{-3/4})$ agents are selected in DES. Fix also l_2 .⁸ If t denotes the step when SRE is completed, then $t = l_2 + O(n \log n)$ w.pr. $1 - O(1/\log n)$.*

6 COIN-BASED ELIMINATION

The next three protocols, LFE, EE1, and EE2, are used to reduce the set of leader candidates provided by DES and SRE, to a single candidate w.h.p. In all three protocols, candidates compete against each other using random coins, eliminating candidates with smaller coin values. Similar protocols are used, e.g., in [24, 25, 30].

⁷Recall that f_1 is the step when the first agent reaches internal phase 1.

⁸Recall that l_2 is the step when the last agent reaches internal phase 2.

Protocol LFE uses $\Theta(\log \log n)$ coins, and reduces the number of candidates by a polylogarithmic factor, during a single internal phase. This yields $O(1)$ expected candidates, if the initial set of candidates is polylogarithmic. Protocol EE1 uses a single coin per candidate per internal phase, reducing the number of candidates by a constant factor per phase, for a total of $O(\log \log n)$ phases. Protocol EE2 is similar to EE1, except that agents can no longer maintain the internal phase value, and use variable parity instead; a mechanism is provided, in SSE, to avoid eliminating all candidates.

6.1 Log-Factors Elimination (LFE)

The state space of LFE is $\mathcal{S}_{\text{LFE}} := \{\text{wait}, \text{in}, \text{out}, \text{toss}\} \times \{0, \dots, \mu\}$, where $\mu := 7 \log \ln n$. We call the second component of an agent's state its *level*. Initially, all agents are in state $(\text{wait}, 0)$. When it reaches internal phase 3, an agent moves from state $(\text{wait}, 0)$ to $(\text{out}, 0)$ if it is eliminated in SRE, or to $(\text{toss}, 0)$ if not eliminated. In the latter case, the agent performs a series of fair coin tosses, increasing its level with each successful coin toss, until the first failure or until it reaches the maximum level μ ; in either case the first component of the agent's state changes from *toss* to *in*. The maximum level reached by any agent is propagated via a one-way epidemic, and any agent with a smaller level changes its first state component from *in* to *out*. Protocol 6 gives the formal transitions.

Protocol 6: LFE

$$\begin{aligned} (\text{wait}, 0) &\Rightarrow \begin{cases} (\text{toss}, 0) & \text{if not eliminated in SRE and iphase} = 3 \\ (\text{out}, 0) & \text{if eliminated in SRE and iphase} = 3 \end{cases} \\ (\text{toss}, \ell) + (\cdot, \cdot) &\rightarrow \begin{cases} (\text{toss}, \ell + 1) & \text{w.pr. } 1/2 \\ (\text{in}, \ell) & \text{w.pr. } 1/2 \end{cases} \quad \text{if } \ell < \mu \\ (\text{toss}, \mu) + (\cdot, \cdot) &\rightarrow (\text{in}, \mu) \\ (s, \ell) + (\cdot, \ell') &\rightarrow (\text{out}, \ell') \quad \text{if } s \in \{\text{in}, \text{out}\} \text{ and } \ell' > \ell \end{aligned}$$

We say LFE is *completed* when every agent is in state (in, x) or (out, x) , where x is the maximal level reached by any agent. An agent is *eliminated* in LFE if it is in a state (out, \cdot) . We say an agent *survives* LFE, if LFE is completed and the agent is not eliminated.

The informal intuition of the protocol is that every agent who survives SRE chooses a random level in $\{0, \dots, \mu\}$ such that level ℓ is chosen with probability $1/2^\ell$, and the agents with the maximal level survive. It follows that an expected constant number of agents survive LFE, if at most 2^μ agents survive SRE.

The next lemma lists the main properties of LFE. Its proof can be found in Appendix G.

LEMMA 8.

- (a) Not all agents are eliminated in LFE.
- (b) Suppose that SRE is completed before step f_3 , and at most $O(2^\mu)$ agents survived SRE. Then in expectation at most $O(1)$ agents are not eliminated in LFE.
- (c) Fix step l_3 . If t denotes the step when LFE is completed, then $t = l_3 + O(n \log n)$ w.h.p.

6.2 Exponential Elimination 1 (EE1)

The state space of EE1 is $\mathcal{S}_{\text{EE1}} := \{\text{in}, \text{out}, \text{toss}\} \times \{0, 1\} \times \{\perp, 4, \dots, v-2\}$, where $v := \Theta(\log \log n)$ is the upper bound of

the range of variable *iphase* (see Section 4). The last component of the state stores the current internal phase ρ of the agent, if $4 \leq \rho \leq v-3$ (i.e., it is equal to *iphase* in this case); the component equals \perp when $\rho < 4$, and $v-2$ when $\rho \geq v-2$.

Initially, all agents are in state $(\text{in}, 0, \perp)$. When it reaches internal phase 4, an agent moves to state $(\text{out}, 0, 4)$ if it is eliminated in LFE, or to $(\text{toss}, 0, 4)$ if not eliminated. In the latter case, the agent performs a single coin toss and stores the outcome, 0 or 1, to the second component of its state, while the first component changes to *in*. The largest coin value is propagated via a one-way epidemic to all agents in internal phase 4, and any agent with a smaller coin value changes its first component to *out*. The same process is repeated for each internal phase $\rho \leq v-2$, and the agents who toss a coin in phase ρ are those who were in state $(\text{in}, \cdot, \rho-1)$ at the end of the previous phase. Protocol 7 gives the formal transitions.

Protocol 7: EE1

$$\begin{aligned} (\text{in}, 0, \perp) &\Rightarrow \begin{cases} (\text{toss}, 0, 4) & \text{if not elim. in LFE \& iphase} = 4 \\ (\text{out}, 0, 4) & \text{if elim. in LFE \& iphase} = 4 \end{cases} \\ (\text{in}, \cdot, \rho) &\Rightarrow (\text{toss}, 0, \rho + 1) \quad \text{if iphase} = \rho + 1 \leq v - 2 \\ (\text{out}, \cdot, \rho) &\Rightarrow (\text{out}, 0, \rho + 1) \quad \text{if iphase} = \rho + 1 \leq v - 2 \\ (\text{toss}, 0, \rho) + (\cdot, \cdot, \cdot) &\rightarrow \begin{cases} (\text{in}, 0, \rho) & \text{w.pr. } 1/2 \\ (\text{in}, 1, \rho) & \text{w.pr. } 1/2 \end{cases} \\ (s, 0, \rho) + (\cdot, 1, \rho) &\rightarrow (\text{out}, 1, \rho) \quad \text{if } s \in \{\text{in}, \text{out}\} \text{ and } \rho \leq v - 2 \end{aligned}$$

We say that an agent is *eliminated* in EE1 if it is in state $(\text{out}, \cdot, \cdot)$. An agent *survives phase ρ* in EE1 if it is not eliminated before reaching internal phase $\rho + 1$.

The intuition for the protocol is that, assuming the clocks of all agents are synchronized, the number of agents that survive roughly halves with each phase, as in expectation half of the coins are 1 and half 0. Thus, if an expected constant number of agents survive LFE, it takes an expected constant number of phases before a single agent is left.

The next lemma lists the main properties of EE1. Its proof can be found in Appendix H.

Recall that $L_{\text{int}}(\rho)$ is the length of internal phase ρ (see Section 4). We define the event

$$\mathcal{W}_{\rho_1, \rho_2} := \bigcap_{\rho_1 \leq r \leq \rho_2} \{L_{\text{int}}(r) \geq (5c_{\mathcal{W}} + 11)n \ln n\}, \quad (1)$$

where $c_{\mathcal{W}} \geq 1$ is a constant to be instantiated in the analysis.

LEMMA 9.

- (a) Not all agents are eliminated in EE1.
- (b) Suppose that LFE is completed before step f_4 , and $k \geq 1$ agents survive LFE. Let s_ρ denote the number of agents that are not eliminated in EE1 before step $f_{\rho+1}$. For any $\rho \in \{4, \dots, v-2\}$, $\mathbb{E}[s_{\rho-1} \cdot \mathbb{1}_{\mathcal{W}_{4, \rho}}] \leq k/2^{\rho-3}$.

6.3 Exponential Elimination 2 (EE2)

The state space of EE2 is $\mathcal{S}_{\text{EE2}} := \{\text{in}, \text{out}, \text{toss}\} \times \{0, 1\} \times \{\perp, 0, 1\}$. The third component of the state stores the parity of the current internal phase ρ of the agent, if $\rho \leq v$, and stores \perp when $\rho < v$. The protocol is essentially the same as EE1, except

that the agents no longer maintain an internal phase counter, and use just the parity instead. [Protocol 8](#) gives the transition rules.

Protocol 8: EE2

$$\begin{aligned}
 (\text{in}, 0, \perp) &\Rightarrow \begin{cases} (\text{toss}, 0, v \bmod 2) & \text{if not elim. in EE1 \& iphase} = v \\ (\text{out}, 0, v \bmod 2) & \text{if elim. in EE1 \& iphase} = v \end{cases} \\
 (\text{in}, \cdot, p) &\Rightarrow (\text{toss}, 0, 1 - p) & \text{if parity} = 1 - p \\
 (\text{out}, \cdot, p) &\Rightarrow (\text{out}, 0, 1 - p) & \text{if parity} = 1 - p \\
 (\text{toss}, 0, p) + (\cdot, \cdot, \cdot) &\rightarrow \begin{cases} (\text{in}, 0, p) & \text{w.pr. } 1/2 \\ (\text{in}, 1, p) & \text{w.pr. } 1/2 \end{cases} \\
 (s, 0, p) + (\cdot, 1, p) &\rightarrow (\text{out}, 1, p) & \text{if } s \in \{\text{in}, \text{out}\}
 \end{aligned}$$

Similar to EE1, we say that an agent is *eliminated* in EE2 if it is in state $(\text{out}, \cdot, \cdot)$, and an agent *survives phase ρ* in EE2 if it is not eliminated before reaching internal phase $\rho + 1$.

The idea is that, as long as the clocks of all agents are synchronized, the internal phases of any two agents differ by at most one, thus, the parity information suffices to tell whether two interacting agents are in the same internal phase. Therefore, EE2 works identically to EE1, in that case. However, if clocks get desynchronized, we cannot ensure the same guarantees; in particular, it is possible that all agents are eliminated. We address this issue in the last subprotocol, SSE, in [Section 7](#).

The next lemma lists the main properties of EE2. Its proof can be found in [Appendix I](#).

LEMMA 10.

- (a) Let $\rho \geq v$. If $L_{\text{int}}(r) > 0$ for all $r \leq \rho + 1$, then some agent survives phase ρ in EE2.
- (b) Let s_ρ denote the number of agents that are not eliminated in EE2 before step $f_{\rho+1}$. For any $v \leq \rho \leq c_W \log n$, it holds $\mathbf{E}[(s_\rho - 1) \cdot \mathbb{1}_{\mathcal{W}_{v-1, \rho+1}}] \leq n/2^{\rho-v+1}$.

7 ENDGAME: SLOW STABLE ELIM'N (SSE)

The last protocol we use, SSE, is responsible for electing a unique leader. Typically, the unique leader is the last surviving candidate in EE1 or EE2. To ensure the protocol is always correct, even if EE1 and EE2 fail to elect a single leader, a basic (slow) mechanism from [\[8\]](#) is employed.

The state space of SSE is $\mathcal{S}_{\text{SSE}} := \{C, E, S, F\}$. (The letters stand for *candidate*, *eliminated*, *survived*, and *failed*, respectively.) Initially, every agent is in state C. Agents eliminated in EE1 switch from state C to E. An agent switches from C to S, if it is not eliminated in EE2 when it reaches external phase 1, or if it is not eliminated in EE1 when it reaches external phase 2. As soon as there is some agent in state S, a one-way epidemic is initiated that spreads state F to agents in states other than S. Moreover, when two agents in state S interact with each other, one of them changes its state to F. [Protocol 9](#) gives the formal transition rules.

We define the set \mathcal{L} of *leader states* of SSE to be the subset $\mathcal{L} := \{C, S\}$. The goal of SSE is to reach a configuration in which exactly one agent is in a leader state. Note that SSE eventually reaches a *final configuration*, where one agent is in state S and all others are in state F. This configurations satisfies the goal stated

Protocol 9: SSE

$$\begin{aligned}
 C &\Rightarrow E & \text{if eliminated in EE1} \\
 C &\Rightarrow S & \text{if (not elim. in EE2 \& xphase} = 1) \text{ or xphase} = 2 \\
 * + S &\rightarrow F \\
 s + F &\rightarrow F & \text{if } s \neq S
 \end{aligned}$$

above, i.e., exactly one agent is in a leader state. However, this goal is typically met much sooner than the final configuration is reached.

The intuition for SSE is as follows. We discuss correctness first (i.e., a single leader is always elected), without worrying about the number of steps. The key observation is that the set of agents whose state is in \mathcal{L} can only shrink over time, but cannot become empty, because not all agents are eliminated in EE1. It follows that some agent will reach state S, and then the protocol transitions ensure that the final configuration is reached eventually.

For the step complexity, assuming clocks are synchronized, we will see that w.pr. $1 - O(1/\log n)$, all but one of the agents are eliminated in EE1, after a constant average number of internal phases, i.e., in $O(n \log n)$ steps. With the remaining probability, $O(1/\log n)$, we rely on EE2 to eliminated all but one agent, which happens in $O(\log n)$ internal phases w.h.p. The single surviving agent moves to state S when it reaches external phase 1, and causes all other agents to switch to state F. In total, this takes $O(n \log^2 n)$ steps, w.h.p. Finally, with *arbitrary* small polynomial probability, clocks get desynchronized, or EE1 and EE2 fail to elect a single leader. In this case, polynomial time is needed: all agents reach the last internal phase in polynomial time, and thus some agent reaches state S, after which SSE guarantees that a single leader is left after polynomial time. Putting all pieces together yields the $O(n \log n)$ bound on the expected step complexity.

The detailed proof of the step complexity is provided in [Section 8](#). The next lemma shows some properties of SSE that are (mostly) independent of the other subprotocols. Its proof can be found in [Appendix J](#).

LEMMA 11. Let L_t denote the set of agents that are either in state C or in state S right after step t .

- (a) For any $t \geq 0$, $L_t \neq \emptyset$ and $L_{t+1} \subseteq L_t$.
- (b) Fix the step l'_1 .⁹ Suppose that exactly one agent is at state S at step l'_1 . Then, w.h.p.,

$$\min \{ t : |L_t| = 1 \} \cup \{ f'_2 \} \leq l'_1 + O(n \log n).$$

- (c) Fix any step $t_2 \geq l'_2$, and suppose that $|L_{t_2}| = \kappa > 1$. Then

$$\mathbf{E}[\min \{ t : |L_t| = 1 \}] \leq t_2 + n^2.$$

8 PROOF OF [Theorem 1](#)

First we prove the correctness, then we show the step complexity bounds, and finally we analyse the space complexity.

8.1 Leader States & Correctness

The state of an agent in LE consists of the agent's state in each of the subprotocols that constitute LE. We define the *leader states* of LE to be all states in which the agent's state in subprotocol SSE is

⁹Recall that for $i \in \{1, 2\}$, f'_i is the step when the first agent reaches external phase i , and l'_i is the step when the last agent reaches external phase at least i .

either C or S, independently of its state in the other subprotocols. LE is a correct leader election protocol if eventually a configuration is reached that is correct, i.e., exactly one agent is in a leader state, and stable, i.e., only correct configurations are reachable from that state. Let L_t denote the set of agents that are in a leader state right after step t . Note that this is the same as the set L_t defined in Lemma 11. From Lemma 11(a), we have that L_t can only shrink over time, but it never becomes empty. Thus it suffice to show that some step t is reached for which $|L_t| = 1$. From Lemma 2(a), at least one agent is elected in JE1, thus at least one agent becomes a clock agent. Then, from Lemma 5, it follows that all agents reach external phase 2, eventually. Once that happens, Lemma 11(c) implies that eventually a step t is reached with $|L_t| = 1$.

8.2 Step Complexity

Let T denote the stabilization time of LE. As discussed above, the set L_t of agents in leader states can only shrink over time t , but it never becomes empty (Lemma 11(a)). Therefore, $T = \min \{ t : |L_t| = 1 \}$.

The proof of the expectation bound proceed by showing that $\mathbb{E}[T \mid E] \cdot \Pr[E] = O(n \log n)$, for a collections of carefully constructed events E , whose union has probability 1 (see (8), (10) and (11)). Then the law of total expectation yields $\mathbb{E}[T] = O(n \log n)$.

Let \mathcal{E}_{JE1} be the event that JE1 is completed in at most $\gamma_1 n \log n$ steps and at most $n^{1-\epsilon}$ agents are elected in JE1. From Lemma 2, for any given $\beta > 0$, there are constants $\epsilon > 0$ and $\gamma_1 > 0$ such that

$$\Pr[\mathcal{E}_{\text{JE1}}] \geq 1 - n^{-\beta}.$$

Let \mathcal{E}_{JE2} be the event that, if at most $n^{1-\epsilon}$ agents are elected in JE1, then JE2 is completed by step $t_{\text{JE1}} + \gamma_2 n \log n$, and $O(\sqrt{n \log n})$ agents are elected in JE2, where t_{JE1} denotes the step when JE1 is completed. From Lemma 3, there is a constant $\gamma_2 > 0$, such that

$$\Pr[\mathcal{E}_{\text{JE2}}] = 1 - O(1/\log n).$$

Let \mathcal{E}_{DES} be the event that, if JE2 is completed before step f_1 and $O(\sqrt{n \log n})$ agents elected in JE2, then DES is completed by step $l_1 + \gamma_3 n \log n$, and the number of agents selected in DES is at least $\Omega(n^{3/4} (\ln \ln n)^{1/4} \cdot (\ln n)^{-3/4})$ and at most $O(n^{3/4} (\ln n)^{3/4})$. From Lemma 6, there is a constant $\gamma_3 > 0$, such that

$$\Pr[\mathcal{E}_{\text{DES}}] = 1 - O(1/\log n).$$

Let \mathcal{E}_{SRE} be the event that, if DES is completed before step f_2 and the number of agents selected in DES is as in \mathcal{E}_{DES} , then SRE is completed by step $l_2 + \gamma_4 n \log n$, and the number of agents that survive SRE is at most $\log^7 n$. From Lemma 7, there is a constant $\gamma_4 > 0$, such that

$$\Pr[\mathcal{E}_{\text{SRE}}] = 1 - O(1/\log n).$$

Let \mathcal{B}_{SRE} be the event that SRE is completed by step f_3 , and the number of agents that survive SRE is $O(\log^7 n)$. Let X be the number of agents not yet eliminated in LFE. Then Lemma 8(b) gives

$$\mathbb{E}[X \mid \mathcal{B}_{\text{SRE}}] = O(1). \quad (2)$$

Let \mathcal{E}_{LFE} be the event that LFE is complete by step $l_3 + \gamma_5 n \log n$. From Lemma 8(c), there is a constant $\gamma_5 > 0$ such that

$$\Pr[\mathcal{E}_{\text{LFE}}] \geq 1 - n^{-\beta}. \quad (3)$$

Let \mathcal{B}_{LFE} be the event that LFE is completed before step f_4 , and let $\mathcal{B} := \mathcal{B}_{\text{SRE}} \cap \mathcal{B}_{\text{LFE}}$. For $\rho \in \{4, \dots, v-2\}$, let $Y_\rho := (s_\rho - 1) \cdot \mathbb{1}_{\mathcal{W}_{4,\rho}}$,

where s_ρ is the number of agents not eliminated in EE1 at $f_{\rho+1}$, and event $\mathcal{W}_{\rho_1, \rho_2}$ is defined in (1). From Lemma 9(b), for any $k \geq 1$,

$$\mathbb{E}[Y_\rho \mid \mathcal{B} \cap \{X = k\}] \leq k/2^{\rho-3}. \quad (4)$$

It follows that

$$\mathbb{E}[Y_\rho \mid \mathcal{B}] \leq \mathbb{E}[X \mid \mathcal{B}]/2^{\rho-3}. \quad (5)$$

Let \mathcal{E}_{SSE} be the event that, if exactly one agent is at state S at step l'_1 , then $\min \{ t : |L_t| = 1 \} \cup \{ f'_2 \} \leq l'_1 + \gamma_6 n \log n$.¹⁰ From Lemma 11(b), there is a constant $\gamma_6 > 0$ such that

$$\Pr[\mathcal{E}_{\text{SSE}}] = 1 - n^{-\beta}.$$

Finally, let \mathcal{E}_{LSC} be the event that if at most $n^{1-\epsilon}$ agents are elected in JE1, then (i) $L_{\text{int}}(\rho) \geq d_1 \cdot n \log n$ and $S_{\text{int}}(\rho) \leq d_2 \cdot n \log n$, for all $\rho \in \{0, 1, \dots, \log^2 n\}$; and (b) $L_{\text{ext}}(\rho') \geq d_3 d_2 \cdot n \log^2 n$ and $S_{\text{ext}}(\rho') \leq d_4 d_2 \cdot n \log^2 n$, for all $\rho' \in \{0, 1\}$. From Lemma 4, there are constants $d_2 \geq d_1 \geq c_1 := (5c_{\mathcal{W}} + 11) + \sum_{1 \leq i \leq 6} \gamma_i$ and $d_4 \geq d_3 \geq c_2 := \beta + 3$, where $c_{\mathcal{W}} := \beta + 2$, such that

$$\Pr[\mathcal{E}_{\text{LSC}}] \geq 1 - n^{-\beta}.$$

By a union bound, the probability of event $\mathcal{E}_{\text{JE1}} \cap \mathcal{E}_{\text{JE2}} \cap \mathcal{E}_{\text{DES}} \cap \mathcal{E}_{\text{SRE}} \cap \mathcal{E}_{\text{LFE}} \cap \mathcal{E}_{\text{LSC}}$ is $1 - O(1/\log n)$. Moreover, this event is a subset of event \mathcal{B} . It follows that the probability of event $\mathcal{E} := \mathcal{E}_{\text{JE1}} \cap \mathcal{E}_{\text{JE2}} \cap \mathcal{E}_{\text{DES}} \cap \mathcal{E}_{\text{SRE}} \cap \mathcal{E}_{\text{LFE}} \cap \mathcal{E}_{\text{LSC}} \cap \mathcal{B}$ is

$$\Pr[\mathcal{E}] = 1 - O(1/\log n).$$

From (2) and the fact that $\mathcal{E} \subseteq \mathcal{B} \subseteq \mathcal{B}_{\text{SRE}}$,

$$\mathbb{E}[X \mid \mathcal{B}] \leq \frac{\mathbb{E}[X \mid \mathcal{B}_{\text{SRE}}]}{\Pr[\mathcal{B} \mid \mathcal{B}_{\text{SRE}}]} \leq \frac{O(1)}{\Pr[\mathcal{B}]} \leq \frac{O(1)}{\Pr[\mathcal{E}]} = O(1). \quad (6)$$

Similarly, from (5) and $\mathcal{E} \subseteq \mathcal{B}$,

$$\mathbb{E}[Y_\rho \mid \mathcal{E}] \leq \frac{\mathbb{E}[Y_\rho \mid \mathcal{B}]}{\Pr[\mathcal{E} \mid \mathcal{B}]} \leq \frac{\mathbb{E}[Y_\rho \mid \mathcal{B}]}{\Pr[\mathcal{E}]} \leq \frac{\mathbb{E}[X \mid \mathcal{B}]}{2^{\rho-3} \Pr[\mathcal{E}]} = O(1/2^{\rho-3}).$$

Then, by Markov's inequality,

$$\Pr[Y_\rho \neq 0 \mid \mathcal{E}] = \Pr[Y_\rho \geq 1 \mid \mathcal{E}] \leq \mathbb{E}[Y_\rho \mid \mathcal{E}] = O(1/2^{\rho-3}). \quad (7)$$

CLAIM 12. $\{Y_\rho = 0\} \cap \mathcal{E} \subseteq \{T \leq d_2(\rho + 2)n \log n\}$.

PROOF. Given \mathcal{E} , it follows from $Y_\rho = 0$ that only one agent is not eliminated in EE1 by step $f_{\rho+1}$. Moreover, from \mathcal{E} it follows that $f_{\rho+1} \leq d_2(\rho + 1)n \log n + \gamma_1 n \log n < d_2(\rho + 2)n \log n < f'_1$. Last, we have that no agent reaches state S in SSE before step f'_1 , and all agents eliminated in EE1 before f'_1 move to state E. Combining all these we obtain that given $\{Y_\rho = 0\} \cap \mathcal{E}$, exactly one agent is in state C in SSE and no agent is in state S at step $d_2(\rho + 2)n \log n$. \square

For $4 < \rho \leq v-2$,

$$\begin{aligned} & \mathbb{E}[T \mid \mathcal{E} \cap \{Y_\rho = 0 \neq Y_{\rho-1}\}] \cdot \Pr[\mathcal{E} \cap \{Y_\rho = 0 \neq Y_{\rho-1}\}] \\ & \leq d_2(\rho + 2)n \log n \cdot \Pr[\mathcal{E}] \cdot \Pr[Y_\rho = 0 \neq Y_{\rho-1} \mid \mathcal{E}] \\ & = O\left(2^{4-\rho} \rho n \log n\right). \end{aligned}$$

where in the second line we used Claim 12, and in the last line we used (7). Similarly, for $\rho = 4$,

$$\mathbb{E}[T \mid \mathcal{E} \cap \{Y_4 = 0\}] \cdot \Pr[\mathcal{E} \cap \{Y_4 = 0\}] = O(n \log n).$$

¹⁰We will only use this event later, but we need to fix γ_6 to define the next event.

Summing the above equations yields

$$\mathbf{E}[T \mid \mathcal{E} \cap \{Y_{v-2} = 0\}] \cdot \Pr[\mathcal{E} \cap \{Y_{v-2} = 0\}] = O(n \log n). \quad (8)$$

For $v \leq \rho \leq c_W \log n$, let $Z_\rho := (s'_\rho - 1) \cdot \mathbb{1}_{\mathcal{W}_{v-1, \rho+1}}$, where s'_ρ is the number of agents not eliminated in EE2 at step $f_{\rho+1}$. From [Lemma 10\(b\)](#),

$$\mathbf{E}[Z_\rho] \leq n/2^{\rho-v+1},$$

and by Markov's inequality,

$$\Pr[Z_\rho \neq 0] = \Pr[Z_\rho \geq 1] \leq \mathbf{E}[Z_\rho] = O(1/2^{\rho-v+1}).$$

Let $\mathcal{E}' := \mathcal{E}_{\text{JE1}} \cap \mathcal{E}_{\text{SSE}} \cap \mathcal{E}_{\text{LSC}}$. By a union bound,

$$\Pr[\mathcal{E}'] \geq 1 - 3n^{-\beta}.$$

Then,

$$\Pr[Z_\rho \neq 0 \mid \mathcal{E}'] \leq \frac{\Pr[Z_\rho \neq 0]}{\Pr[\mathcal{E}']} = O(1/2^{\rho-v+1}). \quad (9)$$

CLAIM 13. For $\rho = (\beta + 2) \log n$,

$$\{Z_\rho = 0\} \cap \mathcal{E}' \subseteq \{T \leq 3d_4 d_2 n \log^2 n\}.$$

PROOF. Given \mathcal{E}' , it follows from $Z_\rho = 0$ that only one agent is not eliminated in EE2 by step $f_{\rho+1}$. (Not all agents are eliminated because of [Lemma 10\(a\)](#).) Moreover, from \mathcal{E}' it follows that $f_{\rho+1} \leq d_2(\rho + 1)n \log n + \gamma_1 n \log n < d_2(\beta + 3)n \log^2 n \leq f'_1$. It also follows that at step l'_1 , there is exactly one agent in state S in SSE (the one not eliminated in EE2). Then, \mathcal{E}_{SSE} yields that $|L_t| = 1$ for $t = l'_1 + \gamma_6 n \log n < f'_2 \leq 2d_4 d_2 n \log^2 n + \gamma_1 n \log n < 3d_4 d_2 n \log^2 n$. \square

Let $\rho^* := (\beta + 2) \log n$. Then

$$\begin{aligned} & \mathbf{E}[T \mid (\mathcal{E}' \cap \{Z_{\rho^*} = 0\}) \setminus (\mathcal{E} \cap \{Y_{v-2} = 0\})] \\ & \quad \cdot \Pr[(\mathcal{E}' \cap \{Z_{\rho^*} = 0\}) \setminus (\mathcal{E} \cap \{Y_{v-2} = 0\})] \\ & \leq 3d_4 d_2 n \log^2 n \cdot (1 - \Pr[\mathcal{E} \cap \{Y_{v-2} = 0\}]) \\ & = 3d_4 d_2 n \log^2 n \cdot (1 - \Pr[\mathcal{E}] \cdot \Pr[Y_{v-2} = 0 \mid \mathcal{E}]) \\ & = O(n \log n), \end{aligned} \quad (10)$$

where in the second line we used [Claim 13](#), and in the last we used $\Pr[\mathcal{E}] = 1 - O(1/\log n)$ and $\Pr[Y_{v-2} = 0 \mid \mathcal{E}] = 1 - O(1/2^v) = 1 - O(1/\log n)$, by [\(7\)](#).

CLAIM 14. Let $\tau := n^2 \log^4 n$. For any integer $k \geq \tau$, $\Pr[T > k] = O(2^{-k/\tau})$.

PROOF. Let t_1 be the step when the first agent is elected in EE1. By repeated application of [Lemma 2\(c\)](#), we obtain that for a large enough constant c , $\Pr[t_1 > i \cdot cn \log n] = O(1/2^i)$. Thus,

$$\Pr[t_1 > k/3] = O(1/2^{(k/3)/(cn \log n)}) = O(1/2^{k/\tau}).$$

Similarly, by repeated application of [Lemma 5](#) and Markov's inequality, we obtain that for a large enough constant c' , $\Pr[l'_2 - t_1 > i \cdot c'n^2 \log^3 n] = O(1/2^i)$. Thus

$$\Pr[l'_2 - t_1 > k/3] = O(1/2^{(k/3)/(c'n^2 \log^3 n)}) = O(1/2^{k/\tau}).$$

Finally, by repeated application of [Lemma 11\(c\)](#) and Markov's inequality, we get $\Pr[T - l'_2 > i \cdot 2n^2] = O(1/2^i)$, thus

$$\Pr[T - l'_2 > (k/3)] = O(1/2^{(k/3)/(2n^2)}) = O(1/2^{k/\tau}).$$

Combining the three equation above, using a union bound, completes the proof. \square

Let C denote the complementary event of

$$(\mathcal{E}' \cap \{Z_{\rho^*} = 0\}) \cup (\mathcal{E} \cap \{Y_{v-2} = 0\}).$$

Then,

$$\Pr[C] \leq 1 - \Pr[\mathcal{E}'] \cdot \Pr[Z_{\rho^*} = 0 \mid \mathcal{E}'] = O(n^{-\beta}),$$

as $\Pr[\mathcal{E}'] \geq 1 - 3n^{-\beta}$, and $\Pr[Z_{\rho^*} = 0 \mid \mathcal{E}'] = 1 - O(n^{-\beta})$ by [\(9\)](#).

Let $\kappa := 1/\Pr[C] = \Omega(n^\beta)$. Choose $\beta \geq 3$, thus $\kappa > \tau$. Then,

$$\begin{aligned} \mathbf{E}[T \mid C] &= \sum_{k \geq 0} \Pr[T > k \mid C] \\ &\leq \kappa + \sum_{k \geq \kappa} \Pr[T > k] / \Pr[C] \\ &\leq \kappa + \sum_{k \geq \kappa} O(2^{-k/\tau}) \cdot \kappa \\ &\leq \kappa + O(2^{-\kappa/\tau} \kappa) \cdot \kappa, \end{aligned}$$

where in the second-last line we used [Claim 14](#), and in the last line we used that $\kappa \geq \tau$ to bound the sum asymptotically by its first κ terms. We then have

$$\mathbf{E}[T \mid C] \cdot \Pr[C] = \mathbf{E}[T \mid C] / \kappa \leq 1 + O(2^{-\kappa/\tau} \kappa).$$

Since $\kappa = \Omega(n^\beta) = \Omega(\tau \log \kappa)$, as $\beta \geq 3$, it follows

$$\mathbf{E}[T \mid C] \cdot \Pr[C] = O(1). \quad (11)$$

Finally, summing [\(8\)](#) [\(10\)](#) [\(11\)](#), and applying the law of total expectation yields $\mathbf{E}[T] = O(n \log n) + O(n \log n) + O(1) = O(n \log n)$.

For the high probability bound, [Claim 13](#) implies

$$\Pr[T \leq 3d_4 d_2 n \log^2 n] \leq \Pr[\{Z_{\rho^*} = 0\} \cap \mathcal{E}'] = 1 - O(n^{-\beta}),$$

as $\Pr[Z_{\rho^*} = 0 \mid \mathcal{E}'] = 1 - O(n^{-\beta})$ by [\(9\)](#), and $\Pr[\mathcal{E}'] \geq 1 - 3n^{-\beta}$.

8.3 Space Complexity

Each individual subprotocol used in protocol LE has a constant-size state space, except for LSC, JE1, LFE, and EE1, which have $\Theta(\log \log n)$ states each. In particular, \mathcal{S}_{LSC} has constant size, but each agent also needs to maintain variable iphase, which takes $v + 1 = \Theta(\log \log n)$ many values. A naive way of combining the states of all subprotocols would be to take the cartesian product of their state spaces, which would yield $\Theta(\log^4 \log n)$ states per agent. A more careful combination of the subprotocols' states allows us to reduce the number of states to $\Theta(\log \log n)$.

First, we make the following observation.

CLAIM 15. If iphase ≥ 1 then the agent's state in JE1 is φ_1 or \perp .

PROOF. Suppose that iphase ≥ 1 for some agent at step t_1 . This is possible only if the agent's internal clock counter in LSC was non-zero at some previous step $t_2 < t_1$, by the definition of the internal phase number. We show by induction that if the internal clock counter of an agent u is non-zero at step t , then u 's state in JE1 at t is either φ_1 or \perp . The claim then follows. The base case of the induction holds vacuously, as all internal clock counters are 0 at $t = 0$. Suppose the statement holds for all steps before step $t \geq 1$, and u 's internal clock counter is non-zero at t . If u is a clock agent at t , then u is elected in JE1 at t , thus its state in JE1 at t is φ_1 . Suppose now that u is not a clock agent at step t . Since its internal clock counter is non-zero, it follows that u interacted (as an initiator) with another agent v at a step $t' \leq t$, and v 's internal

clock counter was non-zero at t' . Since v 's state does not change at t' (as v is the responder), its internal clock counter was non-zero also at $t' - 1$. Applying now the induction hypothesis, which we assume to be true for all steps $t'' < t$ and all agents, we have that v 's state in JE1 at $t' - 1$ (and also at t') is φ_1 or \perp . Thus, when u interacted with v at t' , by the last transition rule of JE1, the status of u in JE1 became \perp (if it was not \perp already). \square

Next we describe a modification to LFE that allows for more efficient space use, without affecting the correctness and step complexity analysis. We essentially stop an agent from executing the protocol as soon as it reaches internal phase 4. Formally, we add to LFE the external transitions

$$\begin{aligned} (s, \cdot) &\Rightarrow (\text{in}, 0) && \text{if } s \in \{\text{in}, \text{toss}\} \text{ and } \text{iphase} = 4 \\ (\text{out}, \cdot) &\Rightarrow (\text{out}, 0) && \text{if } \text{iphase} = 4. \end{aligned}$$

Moreover, we allow the last transition rule in LFE only if $\text{iphase} < 4$, i.e., we replace it by the rule

$$(s, \ell) + (\cdot, \ell') \rightarrow (\text{out}, \ell') \quad \text{if } s \in \{\text{in}, \text{out}\}, \ell' > \ell, \text{ \& iphase} < 4.$$

After these modifications, it is immediate that the following is true.

CLAIM 16. *If $\text{iphase} \geq 4$ then the state of the agent in LFE is either $(\text{in}, 0)$ or $(\text{out}, 0)$.*

Note also that an agent is in the initial state of LFE, $(\text{wait}, 0)$, as long as $\text{iphase} \leq 2$.

We argue now that the correctness and step complexity are not affected. For correctness it suffices to argue that **Lemma 8(a)** still holds: Using a basic induction argument, one can show that for the same sequence of interactions and the same random bits, the set of eliminated agents in the modified LFE protocol is a subset of the corresponding set in the original protocol LFE, at any step t .

Regarding the step complexity analysis, we have that **Lemma 8(b)** and **(c)** no longer hold in general. We observe, however, that the modified LFE protocol is identical to the original one until before step f_4 . Moreover, if LFE is completed before step f_4 , the set of eliminated agents is identical for the two protocols, at all times. Note that **Lemma 8(c)** is only used to show **(3)**. We can replace **(3)** by $\Pr[\mathcal{E}_{\text{LFE}} \cup \bar{\mathcal{B}}_{\text{LFE}}] \geq 1 - n^\beta$, where $\bar{\mathcal{B}}_{\text{LFE}}$ is the complementary of the event \mathcal{B}_{LFE} that LFE is completed before step f_4 . Since **(3)** is used just to compute the probability of event $\mathcal{E} \subseteq \mathcal{B}_{\text{LFE}}$, the above equation suffices. **Lemma 8(b)** is just used to show **(2)**, which is used in **(6)**. In **(6)** the leftmost expectation is conditioned on event $\mathcal{E} \subseteq \mathcal{B}_{\text{LFE}}$. Given \mathcal{B}_{LFE} , variable X is the same in the original and modified process, thus in **(6)** we can assume that the original LFE process is used. For the same reason, **(4)** is not affected either.

For protocol EE1, we recall that the last component of its state can be inferred directly from the value of iphase . Thus we can assume its contribution to the total state space of LE is constant.

We can now count the total number of possible states of an agent as follows. We distinguish three cases based on the value of iphase . If $\text{iphase} = 0$, the agent is in one of $\Theta(\log \log n)$ states in JE1, in one of $\Theta(1)$ states in LSC, and in the initial state of LFE, so $\Theta(\log \log n)$ possible states in total. If $\text{iphase} \in \{1, 2, 3\}$, the agent is in one of two states in JE1, in one of $\Theta(1)$ states in LSC, and in one of $\Theta(\log \log n)$ states in LFE, so again $\Theta(\log \log n)$ states in total. Finally, if $\text{iphase} \in \{4, \dots, \nu\}$, the agent is in one of two states in

JE1, in one of two states in LFE, and in one of $\Theta(\nu)$ many states in LSC, so $\Theta(\nu) = \Theta(\log \log n)$ states in total. Thus, overall we have only $\Theta(\log \log n)$ states per agent.

ACKNOWLEDGMENTS

This research was undertaken, in part, thanks to funding from the ANR Project PAMELA (ANR16-CE23-0016-01).

REFERENCES

- [1] Dan Alistarh, James Aspnes, David Eisenstat, Rati Gelashvili, and Ronald L. Rivest. 2017. Time-space trade-offs in population protocols. In *Proc. Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*. 2560–2579. <https://doi.org/10.1137/1.9781611974782.169>
- [2] Dan Alistarh, James Aspnes, and Rati Gelashvili. 2018. Space-optimal majority in population protocols. In *Proc. Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*. 2221–2239. <https://doi.org/10.1137/1.9781611975031.144>
- [3] Dan Alistarh and Rati Gelashvili. 2015. Polylogarithmic-time leader election in population protocols. In *Proc. 42nd International Colloquium on Automata, Languages, and Programming, ICALP*. 479–491. https://doi.org/10.1007/978-3-662-47666-6_38
- [4] Dan Alistarh and Rati Gelashvili. 2018. Recent algorithmic advances in population protocols. *SIGACT News* 49, 3 (2018), 63–73. <https://doi.org/10.1145/3289137.3289150>
- [5] Dan Alistarh, Rati Gelashvili, and Milan Vojnovic. 2015. Fast and exact majority in population protocols. In *Proc. 2015 ACM Symposium on Principles of Distributed Computing, PODC*. 47–56. <https://doi.org/10.1145/2767386.2767429>
- [6] Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. 2006. Computation in networks of passively mobile finite-state sensors. *Distributed Computing* 18, 4 (2006), 235–253. <https://doi.org/10.1007/s00446-005-0138-3>
- [7] Dana Angluin, James Aspnes, and David Eisenstat. 2008. Fast computation by population protocols with a leader. *Distributed Computing* 21, 3 (2008), 183–199. <https://doi.org/10.1007/s00446-008-0067-z>
- [8] Dana Angluin, James Aspnes, and David Eisenstat. 2008. A simple population protocol for fast robust approximate majority. *Distributed Computing* 21, 2 (2008), 87–102. <https://doi.org/10.1007/s00446-008-0059-z>
- [9] Yossi Azar, Andrei Z. Broder, and Anna R. Karlin. 1994. On-Line Load Balancing. *Theor. Comput. Sci.* 130, 1 (1994), 73–84. [https://doi.org/10.1016/0304-3975\(94\)90153-8](https://doi.org/10.1016/0304-3975(94)90153-8)
- [10] Petra Berenbrink, Robert Elsässer, Tom Friedetzky, Dominik Kaaser, Peter Kling, and Tomasz Radzik. 2018. A population protocol for exact majority with $O(\log^{5/3} n)$ stabilization time and $\Theta(\log n)$ states. In *Proc. 32nd International Symposium on Distributed Computing, DISC*. 10:1–10:18. <https://doi.org/10.4230/LIPIcs.DISC.2018.10>
- [11] Petra Berenbrink, Dominik Kaaser, Peter Kling, and Lena Otterbach. 2018. Simple and efficient leader election. In *Proc. 1st Symposium on Simplicity in Algorithms, SOSA*. 9:1–9:11. <https://doi.org/10.4230/OASiCS.SOSA.2018.9>
- [12] Petra Berenbrink, Dominik Kaaser, and Tomasz Radzik. 2019. On counting the population size. In *Proc. 2019 ACM Symposium on Principles of Distributed Computing, PODC*. 43–52. <https://doi.org/10.1145/3293611.3331631>
- [13] Andreas Bilke, Colin Cooper, Robert Elsässer, and Tomasz Radzik. 2017. Population protocols for leader election and exact majority with $O(\log^2 n)$ states and $O(\log^2 n)$ convergence time. *CoRR* abs/1705.01146 (2017). [arXiv:1705.01146](https://arxiv.org/abs/1705.01146)
- [14] James M. Bower and Hamid Bolouri. 2001. *Computational Modeling of Genetic and Biochemical Networks*. The MIT press.
- [15] Ho-Lin Chen, Rachel Cummings, David Doty, and David Soloveichik. 2017. Speed faults in computation by chemical reaction networks. *Distributed Computing* 30, 5 (2017), 373–390.
- [16] Jurek Czyzowicz, Leszek Gasieniec, Adrian Kosowski, Evangelos Kranakis, Paul G. Spirakis, and Przemyslaw Uznanski. 2015. On convergence and threshold properties of discrete Lotka-Volterra population protocols. In *Proc. 42nd International Colloquium on Automata, Languages, and Programming, ICALP*. 393–405. https://doi.org/10.1007/978-3-662-47672-7_32
- [17] Benjamin Doerr. 2011. Analyzing randomized search heuristics: Tools from probability theory. In *Theory of Randomized Search Heuristics: Foundations and Recent Developments*. World Scientific, 1–20.
- [18] David Doty. 2014. Timing in chemical reaction networks. In *Proc. Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*. 772–784. <https://doi.org/10.1137/1.9781611973402.57>
- [19] David Doty and Mahsa Eftekhari. 2019. Efficient size estimation and impossibility of termination in uniform dense population protocols. In *Proc. 2019*

- ACM Symposium on Principles of Distributed Computing, PODC. 34–42. <https://doi.org/10.1145/3293611.3331627>
- [20] David Doty, Mahsa Eftekhari, Othon Michail, Paul G. Spirakis, and Michail Theofilatos. 2018. Exact size counting in uniform population protocols in nearly logarithmic time. *CoRR abs/1805.04832* (2018). arXiv:1805.04832 <http://arxiv.org/abs/1805.04832>
- [21] David Doty and David Soloveichik. 2018. Stable leader election in population protocols requires linear time. *Distributed Computing* 31, 4 (2018), 257–271. <https://doi.org/10.1007/s00446-016-0281-z>
- [22] Bartłomiej Dudek and Adrian Kosowski. 2018. Universal protocols for information dissemination using emergent signals. In *Proc. 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC*. 87–99. <https://doi.org/10.1145/3188745.3188818>
- [23] Robert Elsässer and Tomasz Radzik. 2018. Recent results in population protocols for exact majority and leader election. *Bulletin of the EATCS* 126 (2018). <http://bulletin.eatcs.org/index.php/beatcs/article/view/549/546>
- [24] Leszek Gasieniec and Grzegorz Stachowiak. 2018. Fast space optimal leader election in population protocols. In *Proc. Twenty-Ninth Annual ACM-SLAM Symposium on Discrete Algorithms, SODA*. 2653–2667. <https://doi.org/10.1137/1.9781611975031.169>
- [25] Leszek Gasieniec, Grzegorz Stachowiak, and Przemyslaw Uznanski. 2019. Almost logarithmic-time space optimal leader election in population protocols. In *Proc. 31st ACM on Symposium on Parallelism in Algorithms and Architectures, SPAA*. 93–102. <https://doi.org/10.1145/3323165.3323178>
- [26] Shafi Goldwasser, Rafail Ostrovsky, Alessandra Scafuro, and Adam Sealfon. 2018. Population stability: Regulating size in the presence of an adversary. In *Proc. 2018 ACM Symposium on Principles of Distributed Computing, PODC*. 397–406. <https://dl.acm.org/citation.cfm?id=3212747>
- [27] Richard M. Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vöcking. 2000. Randomized rumor spreading. In *Proc. 41st Annual Symposium on Foundations of Computer Science, FOCS*. IEEE Computer Society, 565–574. <https://doi.org/10.1109/SFCS.2000.892324>
- [28] Adrian Kosowski and Przemyslaw Uznanski. 2018. Population protocols made easy. *CoRR abs/1802.06872* (2018). arXiv:1802.06872 <http://arxiv.org/abs/1802.06872>
- [29] Yuichi Sudo and Toshimitsu Masuzawa. 2019. Leader election requires logarithmic time in population protocols. *CoRR abs/1906.11121* (2019). arXiv:1906.11121 <http://arxiv.org/abs/1906.11121>
- [30] Yuichi Sudo, Fukuhito Ooshita, Taisuke Izumi, Hirotsugu Kakugawa, and Toshimitsu Masuzawa. 2019. Logarithmic expected-time leader election in population protocol model. In *Proc. 2019 ACM Symposium on Principles of Distributed Computing, PODC*. 60–62. <https://doi.org/10.1145/3293611.3331585>

A PROBABILISTIC TOOLS

A.1 Chernoff Bounds

LEMMA 17. Let X_1, \dots, X_n be independent 0-1 random variables. Let $X := \sum_{1 \leq i \leq n} X_i$ and let $\mu_\ell, \mu_u \geq 0$ be such that $\mu_\ell \leq \mathbf{E}[X] \leq \mu_u$. Then, for any $\delta > 0$,

$$\Pr[X \geq (1 + \delta) \cdot \mu_u] \leq e^{-\frac{\delta^2 \cdot \mu_u}{2 + \delta}}, \quad (12)$$

and for any $0 < \delta < 1$,

$$\Pr[X \leq (1 - \delta) \cdot \mu_\ell] \leq e^{-\frac{\delta^2 \cdot \mu_\ell}{2}}. \quad (13)$$

If the random variables are not independent, then (12) holds, if

$$\sum_{1 \leq i \leq n} \Pr[X_i = 1 \mid X_1, \dots, X_{i-1}] \leq \mu_u,$$

and (13) holds if

$$\sum_{1 \leq i \leq n} \Pr[X_i = 1 \mid X_1, \dots, X_{i-1}] \geq \mu_\ell.$$

A.2 Coupon Collection Bounds

For $0 \leq i < j \leq n$, let $C_{i,j,n}$ be the sum of $j - i$ independent geometric random variables, with expected values $\frac{n}{i+1}, \frac{n}{i+2}, \dots, \frac{n}{j}$. Note that $C_{0,j,n}$ has the same distribution as the number of trials to collect the last j of n coupons in the classic Coupon Collector's

problem. Let $H(k) := \sum_{1 \leq i \leq k} 1/k$ be the k th harmonic number, and recall

$$\ln(k+1) < H(k) \leq \ln k + 1.$$

Let also $H(i,j) := H(j) - H(i)$. Then $\mathbf{E}[C_{i,j,n}] = nH(i,j)$. The next lemma provides some basic tail bounds for $C_{i,j,n}$.

LEMMA 18. For any $0 \leq i < j \leq n$ and $c > 0$,

- (a) $\Pr[|C_{i,j,n} - nH(i,j)| > cn] < \frac{1}{i \cdot c^2}$, if $i \neq 0$.
- (b) $\Pr[C_{i,j,n} > n \cdot \ln\left(\frac{j}{\max\{i,1\}}\right) + cn] < e^{-c}$.
- (c) $\Pr[C_{i,j,n} < n \cdot \ln\left(\frac{j+1}{i+1}\right) - cn] < e^{-c}$.
- (d) $\Pr[C_{0,j,n} < (n-1)\ln(j) - cn] < e^{-e^c}$.

PROOF. Part (a) is obtained using Chebyshev's Inequality We have

$$\text{Var}[C_{i,j,n}] = \sum_{i < k \leq j} \frac{1 - k/n}{(k/n)^2} \leq n^2 \cdot \sum_{i < k \leq j} k^{-2} < n^2/i,$$

and

$$\Pr[|C_{i,j,n} - nH(i,j)| > cn] \leq \frac{\text{Var}[C_{i,j,n}]}{(cn)^2} < \frac{n^2/i}{(cn)^2} = \frac{1}{ic^2}.$$

Part (b), case $i \geq 1$, follows using the standard approach to proving Chernoff-type bounds, based on the moment generating function. For any $s, t > 0$,

$$\Pr[C_{i,j,n} > t] = \Pr[e^{sC_{i,j,n}} > e^{st}] \leq e^{-st} \cdot \mathbf{E}[e^{sC_{i,j,n}}],$$

by Markov's Inequality. Basic calculations yield

$$\mathbf{E}[e^{sC_{i,j,n}}] = \prod_{i < k \leq j} \frac{k/n}{e^{-s} - (1 - k/n)},$$

as long as s is small enough that the denominators in the product are positive. Substituting that above, setting $s := 1/n$ and $t := n \ln \frac{j}{i} + cn$, and using $e^{-1/n} > 1 - 1/n$, gives

$$\begin{aligned} & \Pr\left[C_{i,j,n} > n \cdot \ln\left(\frac{j}{i}\right) + cn\right] \\ & < e^{-\ln(j/i) - c} \cdot \prod_{i < k \leq j} \frac{k/n}{(1 - 1/n) - (1 - k/n)} \\ & = \frac{ie^{-c}}{j} \cdot \prod_{i < k \leq j} \frac{k}{k-1} = e^{-c}. \end{aligned}$$

To show (b), case $i = 0$, we view $C_{0,j,n}$ as the trials to collect the last j of n coupons in the Coupon Collector's problem. By computing the probability that a given coupon is not collected within $n \ln j + cn$ trials, and taking a union bound over the last j coupons, we get

$$\Pr[C_{0,j,n} > n \ln j + cn] \leq j \cdot (1 - 1/n)^{n \ln j + cn} < je^{-\ln j - c} = e^{-c}.$$

For (c), the proof is similar to (b) case $i \geq 1$; we just replace s by $-s$:

$$\Pr[C_{i,j,n} < t] = \Pr[e^{-sC_{i,j,n}} > e^{-st}] \leq e^{st} \cdot \mathbf{E}[e^{-sC_{i,j,n}}],$$

and

$$\mathbf{E}[e^{-sC_{i,j,n}}] = \prod_{i < k \leq j} \frac{k/n}{e^s - (1 - k/n)}.$$

Then, for $s := 1/n$ and $t := n \cdot \ln \left(\frac{j+1}{i+1} \right) - cn$,

$$\begin{aligned} & \Pr \left[C_{i,j,n} > n \cdot \ln \left(\frac{j+1}{i+1} \right) - cn \right] \\ & < e^{\ln \left(\frac{j+1}{i+1} \right) - c} \cdot \prod_{i < k \leq j} \frac{k/n}{(1 + 1/n) - (1 - k/n)} \\ & = \frac{(j+1)e^{-c}}{i+1} \cdot \prod_{i < k \leq j} \frac{k}{k+1} = e^{-c}. \end{aligned}$$

Part (d) is obtained similarly to [17, Theorem 1.24]. We view $C_{0,j,n}$ as the trials to collect the last j of n coupons in the Coupon Collector's problem. Let \mathcal{E}_k be the event that coupon k is collected within $t := (n-1)\ln(j) - cn$ attempts. Then,

$$\Pr[\mathcal{E}_k] = 1 - (1 - 1/n)^t \leq 1 - e^{-t/(n-1)} \leq 1 - e^c/j,$$

where for the first inequality we used the fact $e^{-1} < (1 - 1/n)^{n-1}$. We have

$$\begin{aligned} \Pr[C_{0,j,n} < t] &= \Pr[\mathcal{E}_1 \cap \dots \cap \mathcal{E}_j] \\ &= \Pr[\mathcal{E}_1 \mid \mathcal{E}_2 \cap \dots \cap \mathcal{E}_j] \cdot \Pr[\mathcal{E}_2 \cap \dots \cap \mathcal{E}_j] \\ &\leq \Pr[\mathcal{E}_1] \cdot \Pr[\mathcal{E}_2 \cap \dots \cap \mathcal{E}_j] \\ &\leq \dots \\ &\leq \prod_{1 \leq k \leq j} \Pr[\mathcal{E}_k] \\ &\leq (1 - e^c/j)^j \\ &\leq e^{-e^c}, \end{aligned}$$

where the inequality in the third line holds because the events \mathcal{E}_k are negatively associated. \square

A.3 Runs of a Minimal Length

The following lemma bounds the probability that a run of at least k heads occurs when flipping n independent fair coins.

LEMMA 19. Fix $n, k \in \mathbb{N}$. Let $\mathcal{R}_{n,k}$ denote the event that when flipping n independent fair coins there is a run of at least k consecutive heads. If $n \geq 2k$, then

$$\left(1 - \frac{k+2}{2^{k+1}}\right)^{2^{\lceil n/(2k) \rceil}} \leq \Pr[\overline{\mathcal{R}_{n,k}}] \leq \left(1 - \frac{k+2}{2^{k+1}}\right)^{\lceil n/(2k) \rceil}.$$

PROOF. We first note that for $n = 2k$, we have $\Pr[\mathcal{R}_{n,k}] = (k+2) \cdot 2^{-(k+1)}$. To see this, let $i \in \{0, 1, \dots, k\}$ and let \mathcal{E}_i denote the event that the throws $i+1, i+2, \dots, i+k$ come up head and, if $i \neq 0$, throw i comes up tail. Then $\mathcal{R}_{n,k} = \bigcup_{i=0}^k \mathcal{E}_i$ and for any $i \neq j$ we have $\mathcal{E}_i \cap \mathcal{E}_j = \emptyset$. Thus,

$$\begin{aligned} \Pr[\mathcal{R}_{n,k}] &= \Pr \left[\bigcup_{i=0}^k \mathcal{E}_i \right] = \sum_{i=0}^k \Pr[\mathcal{E}_i] = 2^{-k} + k \cdot 2^{-(k+1)} \\ &= (k+2) \cdot 2^{-(k+1)}. \end{aligned}$$

Now consider $n > 2k$. For $i \in \mathbb{N}_0$ let \mathcal{F}_i denote the event that there is a run of at least k heads during the throws $ik+1, ik+2, \dots, ik+2k$. Then $\Pr[\mathcal{F}_i] = \Pr[\mathcal{R}_{2k,k}] = (k+2) \cdot 2^{-(k+1)}$. Moreover,

\mathcal{F}_i and \mathcal{F}_j are independent if and only if $|i-j| \geq 2$. Since event $\overline{\mathcal{R}_{n,k}}$ implies $\bigcap_{i \text{ even}} \overline{\mathcal{F}_i}$, this yields the upper bound

$$\Pr[\overline{\mathcal{R}_{n,k}}] \leq \Pr \left[\bigcap_{i \text{ even}} \overline{\mathcal{F}_i} \right] \leq \left(1 - \frac{k+2}{2^{k+1}}\right)^{\lceil n/(2k) \rceil}.$$

For the lower bound, we use that $\overline{\mathcal{R}_{n,k}} = (\bigcap_{i \text{ even}} \overline{\mathcal{F}_i}) \cap (\bigcap_{i \text{ odd}} \overline{\mathcal{F}_i})$ and that the events $\overline{\mathcal{F}_i}$ are positively associated to get

$$\begin{aligned} \Pr[\overline{\mathcal{R}_{n,k}}] &= \Pr \left[\bigcap_{i \text{ even}} \overline{\mathcal{F}_i} \cap \bigcap_{i \text{ odd}} \overline{\mathcal{F}_i} \right] \\ &\geq \Pr \left[\bigcap_{i \text{ even}} \overline{\mathcal{F}_i} \right] \cdot \Pr \left[\bigcap_{i \text{ odd}} \overline{\mathcal{F}_i} \right] \\ &\geq \left(1 - \frac{k+2}{2^{k+1}}\right)^{2^{\lceil n/(2k) \rceil}}. \quad \square \end{aligned}$$

A.4 One-way Epidemic

A one-way epidemic for n agents is a population protocol with state space $\{0, 1\}$, and transition rule

$$x + y \rightarrow \max\{x, y\}.$$

Agents in state 1 are called *infected*. Initially, exactly one agent is infected. The number of interactions until all agents are infected is denoted T_{inf} . The next lemma gives bounds on T_{inf} . Similar bounds (without specifying the constants involved) were given in [7].

LEMMA 20. For any $a > 0$ and any n large enough with respect to a , we have that $\Pr[T_{\text{inf}} \leq 4(a+1) \cdot n \ln n] \geq 1 - 2n^{-a}$ and $\Pr[T_{\text{inf}} \geq (n/2) \cdot \ln n] \geq 1 - n^{-a}$.

PROOF. Let N_t denote the number of infected agents after t interactions. Define $T_1 := \min\{t : N_t \geq n/2\}$ and $T_2 := T_{\text{inf}} - T_1$. For $t \geq 0$ and $0 < k < n$,

$$\Pr[N_{t+1} = k+1 \mid N_t = k] = \frac{k \cdot (n-k)}{n \cdot (n-1)} \geq \begin{cases} \frac{k}{2n} & \text{if } k \leq n/2 \\ \frac{n-k}{2n} & \text{if } k > n/2. \end{cases}$$

It follows that both T_1 and T_2 are dominated by $C_{0, \lfloor n/2 \rfloor, 2n}$. Then, for $j \in \{1, 2\}$ and $a > 0$, Lemma 18(b) yields,

$$\begin{aligned} & \Pr[T_j > 2(a+1) \cdot n \ln n] \\ & \leq \Pr[C_{0, \lfloor n/2 \rfloor, 2n} > 2(a+1) \cdot n \ln n] \\ & \leq \Pr[C_{0, \lfloor n/2 \rfloor, 2n} > 2n \ln \lfloor n/2 \rfloor + 2a \cdot n \ln n] < n^{-a}, \end{aligned}$$

The desired upper bound follows by a union bound,

$$\begin{aligned} & \Pr[T_{\text{inf}} > 4(a+1) \cdot n \ln n] \\ & \leq \Pr[\max\{T_1, T_2\} > 4(a+1) \cdot n \ln n] \leq 2n^{-a}. \end{aligned}$$

For the lower bound, we note $\Pr[N_{t+1} = k+1 \mid N_t = k] \leq k/n$. Thus, T_{inf} dominates $C_{0, n-1, n}$. From Lemma 18(d), we get that for any $a > 0$ and n large enough,

$$\begin{aligned} & \Pr[T_{\text{inf}} < (n/2) \cdot \ln n] \\ & \leq \Pr[C_{0, n-1, n} < (n/2) \cdot \ln n] \\ & \leq \Pr[C_{0, n-1, n} < (n-1) \cdot \ln(n-1) - \ln(a \ln n) \cdot n] \\ & < n^{-a}. \quad \square \end{aligned}$$

B ANALYSIS OF JE1

To prove [Lemma 2](#), we consider a variant of protocol JE1 without rejections (i.e., without the transition rule $\ell + \ell' \rightarrow \perp$ from [Protocol 1](#)). It is easy to see that for all $k \in \{-\psi, -\psi + 1, \dots, \varphi_1\}$, the number of agents on level at least k in JE1 is stochastically dominated by the corresponding number of agents in JE1 without rejections. Hence, the number of agents in state φ_1 of JE1 can be upper bounded by the corresponding number of the variant without rejections.

For $k \in \mathcal{S}_{\text{JE1}} \setminus \{\perp\}$ and $t \in \mathbb{N}_0$, let $A_k(t)$ denote the number of agents u with state (level) $\geq k$ after interaction t for JE1 without rejection. $\hat{A}_k(t)$ is defined accordingly for JE1 with rejections. Similarly, $A_\perp(t)$ ($\hat{A}_\perp(t)$) denotes the number of agents u in state \perp after interaction t for JE1 without (with) rejections.

First we show in [Lemma 21](#) that after $\Theta(n \log n)$ interactions the number of agents on level ≥ 0 is w.h.p. $\xi_0 \cdot n$, where $\xi_0 \approx 1/(\log n)^2$. [Lemma 22](#) uses this to derive an upper bound of $n^{1-\Omega(1)}$ on the number of agents in state φ_1 (i.e., on the number of elected agents) after $\Theta(n \log n)$ interactions. The idea of the proof is as follows: When at most $\xi_k \cdot n$ agents are on level $k \in \{0, 1, \dots, \varphi_1 - 1\}$, the probability for a new agent to reach level $k + 1$ is at most ξ_k^2 . Using linearity of expectation and standard tail bounds, this leads to a recursive argument (similar to [\[9\]](#)) showing that after $\Theta(n \log n)$ interactions, w.h.p. not more than $n/(\log n)^{-(2^k+1)}$ agents reach level k . With our choice of the maximum level φ_1 , this results in at most $n^{1-\Omega(1)}$ agents on level φ_1 . [Lemma 23](#) uses a similar recursive approach to prove, w.h.p., a lower bound of at least $n^{1/2}$ agents on level φ_1 after $\Theta(n \log n)$ interactions.

LEMMA 21. *For $a \geq 13/12$ and $c := 12a$ define $\tau := c \cdot n \log n$, $c_1 = c/34$, and $c_2 = 4c$. Consider JE1 without rejections and assume that n is large enough with respect to c .*

$$\Pr \left[A_0(\tau) \in \left[c_1 \cdot \frac{n}{(\log n)^2}, c_2 \cdot \frac{n}{(\log n)^2} \right] \right] \geq 1 - n^{-a}.$$

PROOF. For an agent u let N_u be the random variable counting the number of the first τ interactions which are initiated by u . Since $N_u \sim \text{Bin}(\tau, 1/n)$, standard Chernoff bounds yield $N_u \in [\tau/(2n), 2\tau/n]$ with probability at least $1 - n^{-c/6}$. Let \mathcal{E} denote the event “ $\forall u: N_u \in [\tau/(2n), 2\tau/n]$ ”. A union bound over all n agents gives $\Pr[\bar{\mathcal{E}}] \leq n^{-c/6+1}$.

Let $\mathcal{L}_{u,t}$ denote the event that agent u has level ≥ 0 after agent it initiated t interactions. Furthermore, let L_t denote the number of agents with level ≥ 0 after their respective t -th initiated interaction. We define $\mathcal{R}_{t,\psi}$ as the event that a sequence of ψ consecutive heads occurs when throwing t independent fair coins (see [Lemma 19](#)). By definition of our protocol, we have $\Pr[\mathcal{L}_{u,t}] = \Pr[\mathcal{R}_{t,\psi}]$. As an immediate consequence of [Lemma 19](#) and by our choice of ψ and τ ,

we get

$$\begin{aligned} \Pr[\mathcal{L}_{u, \lceil \tau/(2n) \rceil}] &\geq 1 - \left(1 - \frac{\psi + 2}{2^{\psi+1}} \right)^{\lceil \tau/(4\psi \cdot n) \rceil} \\ &\geq 1 - \exp \left(- \frac{\psi + 2}{2^{\psi+1}} \cdot \left\lfloor \frac{\tau}{4\psi \cdot n} \right\rfloor \right) \\ &\geq 1 - \exp \left(- \frac{\psi + 2}{2^{\psi+1}} \cdot \frac{\tau}{4\psi \cdot n} \cdot (1 - o(1)) \right) \quad (14) \\ &\geq 1 - \exp \left(- \frac{c}{8 \cdot (\log n)^2} \cdot (1 - o(1)) \right) \\ &\geq \frac{c}{16 \cdot (\log n)^2} \cdot (1 - o(1)) \geq \frac{c}{17 \cdot (\log n)^2} \end{aligned}$$

for n large enough with respect to c . Similarly, we get

$$\begin{aligned} \Pr[\mathcal{L}_{u, \lfloor 2\tau/n \rfloor}] &\leq 1 - \left(1 - \frac{\psi + 2}{2^{\psi+1}} \right)^{2\lceil \tau/(\psi \cdot n) \rceil} \\ &\leq \frac{\psi + 2}{2^{\psi}} \cdot \left\lceil \frac{\tau}{\psi \cdot n} \right\rceil \quad (15) \\ &\leq \frac{\psi}{2^{\psi}} \cdot \frac{\tau}{\psi \cdot n} \cdot (1 + o(1)) \\ &= \frac{c}{(\log n)^2} \cdot (1 + o(1)) \leq \frac{2c}{(\log n)^2}. \end{aligned}$$

Note that, for any $t \in \mathbb{N}$, the events $\mathcal{L}_{u,t}$ over all agents u are mutually independent. Thus, [\(14\)](#) and [\(15\)](#) together with standard Chernoff bounds imply

$$\Pr \left[L_{\lceil \tau/(2n) \rceil} \leq \frac{c \cdot n}{34 \cdot (\log n)^2} \right] \leq \exp \left(- \frac{c \cdot n}{136 \cdot (\log n)^2} \right) \quad (16)$$

and

$$\Pr \left[L_{\lfloor 2\tau/n \rfloor} \geq \frac{4c \cdot n}{(\log n)^2} \right] \leq \exp \left(- \frac{2c \cdot n}{3 \cdot (\log n)^2} \right). \quad (17)$$

Note that – since the $\mathcal{L}_{u,t}$ are also independent of \mathcal{E} – both bounds also hold conditioned on \mathcal{E} .

Recall that $A_0(\tau)$ is the number of agents with level ≥ 0 after τ interactions. Conditioned on \mathcal{E} , event $A_0(\tau) \leq x$ implies $L_{\lceil \tau/(2n) \rceil} \leq x$ and $A_0(\tau) \geq x$ implies $L_{\lfloor 2\tau/n \rfloor} \geq x$. The lemma’s statement follows by applying a union bound to [\(16\)](#) and [\(17\)](#) as follows:

$$\begin{aligned} &\Pr \left[A_0(\tau) \notin \left[\frac{c \cdot n}{34 \cdot (\log n)^2}, \frac{4c \cdot n}{(\log n)^2} \right] \right] \\ &\leq \Pr \left[A_0(\tau) \notin \left[\frac{c \cdot n}{34 \cdot (\log n)^2}, \frac{4c \cdot n}{(\log n)^2} \right] \mid \mathcal{E} \right] + \Pr[\bar{\mathcal{E}}] \\ &\leq \Pr \left[L_{\lceil \tau/(2n) \rceil} \leq \frac{c \cdot n}{34 \cdot (\log n)^2} \mid \mathcal{E} \right] \\ &\quad + \Pr \left[L_{\lfloor 2\tau/n \rfloor} \geq \frac{4c \cdot n}{(\log n)^2} \mid \mathcal{E} \right] + n^{-(c-6)/6} \\ &\leq \exp \left(- \frac{c \cdot n}{136 \cdot (\log n)^2} \right) + \exp \left(- \frac{2c \cdot n}{3 \cdot (\log n)^2} \right) + n^{-(c-6)/6} \\ &\leq n^{-c/12}. \end{aligned}$$

The last inequality holds for n large enough with respect to c for $c \geq 13$. \square

LEMMA 22. For $a \geq 1/12$ and $c := 12a + 12$, define $\tau := c \cdot n \log n$ and let $\epsilon := 1/32$. Consider $\mathcal{J}E1$ without rejections and assume that n is large enough with respect to c . Then

$$\Pr[A_{\varphi_1}(\tau) < n^{1-\epsilon}] \geq 1 - n^{-a}.$$

PROOF. In order to prove the lemma's statement, we derive a recursive upper bound on the number of agents on level $\geq k$ as a function of the number of agents on level $\geq k-1$.

For $k \in \{0, 1, \dots, \varphi_1\}$ let \mathcal{B}_k denote the (bad) event " $A_k(\tau) \geq \xi_k \cdot n$ ". The values $\xi_k \in [n^{-1/4}, 1)$ will be fixed during the proof. We will prove that

$$\Pr[\mathcal{B}_k \mid \overline{\mathcal{B}_{k-1}}] \leq \frac{n^{-(a+1)}}{\Pr[\overline{\mathcal{B}_{k-1}}]} \quad (18)$$

for $k \in \{1, 2, \dots, \varphi_1\}$. This implies

$$\begin{aligned} \Pr[\mathcal{B}_k] &\leq \Pr[\mathcal{B}_k \mid \overline{\mathcal{B}_{k-1}}] \cdot \Pr[\overline{\mathcal{B}_{k-1}}] + \Pr[\mathcal{B}_{k-1}] \\ &\leq n^{-(a+1)} + \Pr[\mathcal{B}_{k-1}]. \end{aligned} \quad (19)$$

Applying this recursively and using $\xi_{\varphi_1} \leq n^{-1/32}$ (see below) yields

$$\begin{aligned} \Pr[A_{\varphi_1}(\tau) \geq n^{1-1/32}] &\leq \Pr[\mathcal{B}_{\varphi_1}] \leq \varphi_1 \cdot n^{-(a+1)} + \Pr[\mathcal{B}_0] \\ &\leq (\varphi_1 + 1) \cdot n^{-(a+1)}, \end{aligned} \quad (20)$$

where the last inequality uses Lemma 21 and $\xi_0 = 4c/(\log n)^2$. The lemma's statement follows from this via a union bound since $\varphi_1 = O(\log \log n) = o(n)$.

It remains to prove (18). In the following, let $(\mathcal{F}_t)_{t \in \mathbb{N}_0}$ denote the filtration in which \mathcal{F}_t describes the outcome of the first t interactions. Fix a level $k \in \{1, 2, \dots, \varphi_1\}$ and an interaction $t \in \{1, 2, \dots, \tau\}$. Let the binary random variable X_t be 1 if and only if both a new agent reaches level k during interaction t and $A_{k-1}(t-1) < \xi_{k-1} \cdot n$. Observe that $\Pr[X_t = 1 \mid \mathcal{F}_{t-1}] \leq \frac{\xi_{k-1}^2}{\xi_{k-1}}$ (one of the at most $A_{k-1}(t-1) < \xi_{k-1} \cdot n$ agents on level exactly $k-1$ must initiate an interaction with a responder of the at most $A_{k-1}(t-1) < \xi_{k-1} \cdot n$ agents on level at least $k-1$). Thus, the sum $\sum_{t=1}^{\tau} X_t$ is stochastically dominated by $\text{Bin}(\tau, \frac{\xi_{k-1}^2}{\xi_{k-1}})$. Since $A_{k-1}(t)$ is monotonically non-decreasing in t , $A_{k-1}(\tau) < \xi_{k-1} \cdot n$ (i.e., event $\overline{\mathcal{B}_{k-1}}$) implies $\sum_{t=1}^{\tau} X_t = A_k(\tau)$. We apply this (conditioned) identity and the aforementioned stochastic domination together with a standard Chernoff bound to get

$$\begin{aligned} &\Pr[A_k(\tau) \geq 2\xi_{k-1}^2 \cdot \tau \mid \overline{\mathcal{B}_{k-1}}] \\ &= \Pr\left[\sum_{t=1}^{\tau} X_t \geq 2\xi_{k-1}^2 \cdot \tau \mid \overline{\mathcal{B}_{k-1}}\right] \\ &\leq \Pr\left[\sum_{t=1}^{\tau} X_t \geq 2\xi_{k-1}^2 \cdot \tau \mid \Pr[\overline{\mathcal{B}_{k-1}}]\right] \\ &\leq \exp\left(-\frac{\xi_{k-1}^2 \cdot \tau}{3}\right) / \Pr[\overline{\mathcal{B}_{k-1}}] \\ &\leq \frac{n^{-(a+1)}}{\Pr[\overline{\mathcal{B}_{k-1}}]}, \end{aligned} \quad (21)$$

where the last inequality uses $\xi_{k-1} \geq n^{-1/4}$.

Based on (21) and Lemma 21, we define $\xi_k := 2\xi_{k-1}^2 \cdot \tau/n$ for $k \in \{1, 2, \dots, \varphi_1\}$ and $\xi_0 := 4c/(\log n)^2$. With this, (21) can be written

as (18). It only remains to prove that these ξ_k fulfill the required property $\xi_k \geq n^{-1/4}$ and that $\xi_{\varphi_1} \leq n^{-1/32}$. Both of these follow by our choice of φ_1 . To see this, note that $\xi_k = (2\tau/n)^{2^{k-1}} \cdot \xi_0^{2^k}$ for all $k \in \{0, 1, \dots, \varphi_1\}$ and that ξ_k is monotonously decreasing in k . Using this we see that, for any $\epsilon > 0$, the inequality $\xi_k \leq n^{-\epsilon}$ is equivalent to

$$k \geq \log \log \left[\frac{n^{1+\epsilon}}{2\tau} \right] - \log \log \left[\frac{n}{2\tau \cdot \xi_0} \right].$$

By substituting the definitions of τ and ξ_0 , we see that $k \geq \log \log n - \log \log \log n + \log(2\epsilon)$ is sufficient for $\xi_k \leq n^{-\epsilon}$. In particular, this inequality holds for $\epsilon = 1/32$ and our choice of φ_1 . Similarly we see that $k \leq \log \log n - \log \log \log n + \log(\epsilon/2)$ is sufficient for $\xi_k \geq n^{-\epsilon}$, which holds for $\epsilon = 1/4$ and our choice of φ_1 . \square

LEMMA 23. For $a \geq 1/6$ and $c := 12a + 12$, define $\tau := c \cdot n \log n$. For $\mathcal{J}E1$ without rejections and for n large enough with respect to c ,

$$\Pr[A_{\varphi_1}(\tau) > n^{1/2}] \geq 1 - n^{-a}.$$

PROOF. The outline of the proof is similar to that of Lemma 22. We first derive a recursive lower bound on the number of agents on level $\geq k$ (at a specific point in time) when given a corresponding bound on the number of agents on level $\geq k-1$. More exactly, we define $\tau_0 := (c-1) \cdot n \log n$ and $\tau_k := \tau_{k-1} + \lceil n/2^k \rceil$ for $k \in \{1, 2, \dots, \varphi_1\}$. Note that $\tau_{\varphi_1} \leq \tau_0 + n \leq \tau$. For $k \in \{0, 1, \dots, \varphi_1\}$ let \mathcal{B}_k denote the (bad) event " $A_k(\tau_k) \leq \xi_k \cdot n$ ". The values $\xi_k \in [(\log n)^2/n^{1/2}, 1)$ will be fixed during the proof. We will prove that

$$\Pr[\mathcal{B}_k \mid \overline{\mathcal{B}_{k-1}}] \leq n^{-(a+1)} / \Pr[\overline{\mathcal{B}_{k-1}}] \quad (22)$$

for $k \in \{1, 2, \dots, \varphi_1\}$. This implies

$$\begin{aligned} \Pr[\mathcal{B}_k] &\leq \Pr[\mathcal{B}_k \mid \overline{\mathcal{B}_{k-1}}] \cdot \Pr[\overline{\mathcal{B}_{k-1}}] + \Pr[\mathcal{B}_{k-1}] \\ &\leq n^{-(a+1)} + \Pr[\mathcal{B}_{k-1}]. \end{aligned} \quad (23)$$

Applying this recursively and using $\tau_{\varphi_1} \leq \tau$ as well as $\xi_{\varphi_1} \geq n^{-1/2}$ yields

$$\begin{aligned} \Pr[A_{\varphi_1}(\tau) \leq n^{1/2}] &\leq \Pr[\mathcal{B}_{\varphi_1}] \leq \varphi_1 \cdot n^{-(a+1)} + \Pr[\mathcal{B}_0] \\ &\leq (\varphi_1 + 1) \cdot n^{-(a+1)}, \end{aligned}$$

where the last inequality uses Lemma 21 and $\xi_0 = (c-1)/(34 \cdot (\log n)^2)$. The lemma's statement follows from this via a union bound since $\varphi_1 = O(\log \log n) = o(n)$.

In order to prove (22) we introduce some additional notation. Again, let $(\mathcal{F}_t)_{t \in \mathbb{N}_0}$ denote the filtration in which \mathcal{F}_t describes the outcome of the first t interactions. For $k \in \{1, 2, \dots, \varphi_1\}$ define $I_k := \{\tau_{k-1} + 1, \tau_{k-1} + 2, \dots, \tau_k\}$. Note that $|I_k| = \lceil n/2^k \rceil \geq n/2^{\varphi_1} \geq n/\log n$. Remember the values ξ_k used in the definition of the events \mathcal{B}_k . When we fix ξ_k below, we will ensure that $\xi_k \leq \xi_{k-1}/4$ for $k \in \{1, 2, \dots, \varphi_1\}$ and that $\xi_k \geq (\log n)^2/n^{1/2}$ for $k \in \{0, 1, \dots, \varphi_1\}$.

Now, fix a level $k \in \{1, 2, \dots, \varphi_1\}$ and an interaction $t \in I_k$. Let the binary random variable X_t be 1 if and only if a new agent reaches level k during interaction t or if $A_k(t-1) > \xi_k \cdot n$. Using the basic probability inequality $\Pr[A \vee B \mid C] \geq \Pr[A \mid \overline{B} \wedge C]$ and

$\xi_k \leq \xi_{k-1}/4$, we get

$$\begin{aligned} \Pr[X_t \mid \mathcal{F}_{t-1}, \overline{\mathcal{B}_{k-1}}] &\geq \Pr[X_t \mid \mathcal{F}_{t-1}, \overline{\mathcal{B}_{k-1}}, A_k(t-1) \leq \xi_k \cdot n] \\ &\geq (\xi_{k-1} - \xi_k) \cdot (\xi_{k-1} - 1/n) \geq \xi_{k-1}^2/2. \end{aligned}$$

Thus, conditioned on $\overline{\mathcal{B}_{k-1}}$, the sum $\sum_{t \in I_k} X_t$ stochastically dominates $\text{Bin}(|I_k|, \xi_{k-1}^2/2)$. The definition of the X_t , the stochastic dominance, and standard Chernoff bounds now yield

$$\begin{aligned} &\Pr[A_k(\tau_k) \leq \xi_{k-1}^2 \cdot |I_k|/4 \mid \overline{\mathcal{B}_{k-1}}] \\ &\leq \Pr \left[\sum_{t \in I_k} X_t \leq \xi_{k-1}^2 \cdot |I_k|/4 \mid \overline{\mathcal{B}_{k-1}} \right] \\ &\leq \Pr \left[\sum_{t \in I_k} X_t \leq \xi_{k-1}^2 \cdot |I_k|/4 \mid \Pr[\overline{\mathcal{B}_{k-1}}] \right] \quad (24) \\ &\leq \exp \left(-\frac{\xi_{k-1}^2 \cdot |I_k|}{16} \right) / \Pr[\overline{\mathcal{B}_{k-1}}] \\ &\leq n^{-(a+1)} / \Pr[\overline{\mathcal{B}_{k-1}}], \end{aligned}$$

where the last inequality uses $\xi_{k-1} \geq (\log n)^2/n^{1/2}$.

Based on (24) and Lemma 21, we define $\xi_k := \xi_{k-1}^2 \cdot |I_t|/(4n)$ for $k \in \{1, 2, \dots, \varphi_1\}$ and $\xi_0 := (c-1)/(34 \cdot (\log n)^2)$. With this, (24) can be written as (22). It only remains to prove that these ξ_k fulfill the required properties $\xi_k \leq \xi_{k-1}/4$ and $\xi_k \geq (\log n)^2/n^{1/2}$. The first property is obvious from the recursive definition. The second property follows by our choice of φ_1 . To see this, note that $\xi_k \geq \xi_{k-1}^2 \cdot (n/2^k)/(4n) = \xi_{k-1}^2/2^{l+2}$. Solving the corresponding recursion yields $\xi_k \geq 2^{4-2^{2+l}} \cdot \xi_0^{2^k}$ for all $k \in \{0, 1, \dots, \varphi_1\}$. Using this we see that, for any $\epsilon > 0$, the inequality $\xi_k \geq n^{-\epsilon}$ is implied by

$$k \leq \log \left(\frac{4+k+\epsilon \log n}{4+\log(1/\xi_0)} \right).$$

Since $4+k \geq 0$ and $4 \leq \log(1/\xi_0)/2$, we get that $k \leq \log \log n - \log \log(1/\xi_0) + \log(2\epsilon/3)$ is sufficient for $\xi_k \geq n^{-\epsilon}$. To simplify this further, we use $1/\xi_0 \leq (\log n)^{5/2}$, which yields that $k \leq \log \log n - \log \log \log n + \log(4\epsilon/15)$ is sufficient for $\xi_k \geq n^{-\epsilon}$. In particular, this inequality holds for $\epsilon = 15/32$ and our choice of φ_1 . As a result, we see that $\xi_k \geq n^{-15/32} \geq (\log n)^2/n^{1/2}$. \square

With the above statement, we are now able to prove the main result for JE1.

B.1 Proof of Lemma 2

We first prove statement (a), stating that at least one agent is elected in JE1. Note that here we can assume that all agents start on level $-\psi$. By the definition of the protocol, no agent can enter state \perp before there exists an agent on level φ_1 . Furthermore, any agent that reaches level φ_1 will not be rejected by JE1 during a later interaction. The statement follows from this together with the observation that for any agent on level $< \varphi_1$, the probability to increase its level is positive.

Recall that we use $\tilde{A}_k(t)$ and $A_k(t)$ to denote the number of agents on level $\geq k$ after t interactions of JE1 with and without

rejections, respectively. Define the random variable

$$T_0 := \min \{ t \in \mathbb{N} : |A_{\varphi_1}(t)| + |A_{\perp}(t)| \geq 1 \}$$

as well as the random variable

$$T_1 := \min \{ t \in \mathbb{N} : |A_{\varphi_1}(t)| + |A_{\perp}(t)| = n \}.$$

We use \tilde{T}_0 and \tilde{T}_1 to denote the corresponding variables when running JE1 with rejections. In particular, \tilde{T}_1 and T_1 are the interactions after which JE1 with and without rejections are completed. Note that JE1 with rejections will not reject any agents as long as there is no agent in state φ_1 or state \perp . Hence, JE1 with rejections behaves exactly like JE1 without rejections and T_0 and \tilde{T}_0 both have the same distribution.

We prove statement (c) before statement (b), as the former is useful in the latter's proof. After interaction \tilde{T}_0 , the last transition rule of JE1 distributes \perp in an epidemic fashion. Thus, Lemma 20 implies that, with probability at least $1 - 2n^{-(a+1)}$,

$$\tilde{T}_1 \leq \tilde{T}_0 + 4(a+3) \cdot n \ln n.$$

Thus, it is sufficient to prove that w.h.p. $\tilde{T}_0 = O(n \log n)$. Since T_0 and \tilde{T}_0 both have the same distribution we can also show that w.h.p. $T_0 = O(n \log n)$.

Case 1: All agents start from state $-\psi$. Then Lemma 23 implies that, for any $a > 0$ and $c' := 12a + 24$, with probability at least $1 - n^{-(a+1)}$ we have $T_0 \leq c' \cdot n \log n$, yielding the desired result.

Case 2: All agents start from an arbitrary state. In the following we reduce this case to the case that the agents start from state $-\psi$. If there exists an agent in state φ_1 or \perp it is easy to see that $T_0 = 0$ and we are done. Otherwise, we use a simple identity coupling to couple an execution of JE1 starting from our arbitrary state (called process A in the following) with an execution of JE1 where all agents start from state $-\psi$ (called process B). We define $T_0^{(A)}$ and $T_0^{(B)}$ according to the definition of T_0 above. One can easily check that, for any

$$t < \min \{ T_0^{(A)}, T_0^{(B)} \},$$

the identity coupling maintains the property that $\ell_u^{(A)}(t) \geq \ell_u^{(B)}(t)$, where $\ell_u^{(X)}(t)$ denotes the level of agent u in process X after t interactions. This implies $T_0^{(A)} \leq T_0^{(B)}$. Due to the definition of B (all agents start from state $-\psi$) it follows from Case 1 that w.h.p., $T_0^{(B)} = O(n \log n)$, this yields the desired bound.

Combining the cases above, we get that, with probability at least $1 - n^{-a}$,

$$\tilde{T}_1 \leq \tilde{T}_0 + 4(a+3)n \ln n \leq (16a+27) \cdot n \log n,$$

yielding statement (c).

For statement (b), note that, similarly to above, a simple identity coupling gives that, for all $k \in \mathcal{S}_{\text{JE1}} \setminus \{\perp\}$ and $t \in \mathbb{N}_0$, the random variable $\tilde{A}_k(t)$ is stochastically dominated by $A_k(t)$. Thus, we can apply Lemma 22 to get that, for any $a > 0$, the number of survivors in JE1 after $(12a+24) \cdot n \log n \leq (16a+27) \cdot n \log n$ interactions is, with probability at least $1 - n^{-(a+1)}$, at most $n^{31/32}$. A union bound with statement (c) yields the desired result.

C ANALYSIS OF JE2

For $k \in \{0, \dots, \varphi_2\}$ let the random variable A_k denote the number of agents that reach at least level k . Note, the agents that reach level $\max\{k: A_k > 0\}$ are exactly those that are not rejected in JE2.

LEMMA 24. *Let $k \in \{1, \dots, \varphi_2 - 1\}$, $a > 0$, and $\xi \geq n^{-1/2}$. $\sqrt{32a \cdot \ln n}$. Then*

- (a) $\Pr[\{A_{k+1} \geq 2\xi^2 \cdot n\} \cap \{A_k \leq \xi \cdot n\}] \leq n^{-a}$.
- (b) $\Pr[\{A_{k+1} \leq \xi^2 \cdot n/8\} \cap \{A_k \geq \xi \cdot n\}] \leq n^{-a}$.

PROOF. Let u_i be the i th agent that reaches level k . In the next interaction r_i that u_i initiates, u_i will either reach level $k+1$ or become inactive. The probability that u_i reaches level $k+1$ is at most $(\xi n - 1)/(n - 1) \leq \xi$, if the number of agents on level at least k is at most ξn right before interaction r_i . It follows that, if X_i is the indicator random variable of the event: “ $i \leq A_k \leq \xi n$ and u_i reaches level $k+1$,” then $\Pr[X_i = 1] \leq \xi$; and this holds even conditionally on all $X_{i'}, i' \neq i$. Moreover,

$$\Pr\left[\sum_{1 \leq i \leq \xi n} X_i \geq 2\xi^2 n\right] = \Pr[\{A_{k+1} \geq 2\xi^2 n\} \cap \{A_k \leq \xi n\}].$$

By an application of Chernoff bound (12),

$$\Pr\left[\sum_{1 \leq i \leq \xi n} X_i \geq 2\xi^2 n\right] \leq \exp\left(-\frac{\xi^2 n}{3}\right) \leq n^{-a},$$

where the last inequality uses the lower bound on ξ . Combining the last two equations above proves (a).

The proof of (b) is similar. Assuming again that u_i is the i th agent to reach level k , the probability that u_i reaches level $k+1$ is at least $(i-1)/(n-1)$. It follows that, if Y_i is the indicator of the event: “ $A_k < i$ or u_i reaches level $k+1$,” then $\Pr[Y_i = 1] \geq (i-1)/(n-1)$; and this holds even conditionally on all $Y_{i'}, i' \neq i$. Moreover,

$$\Pr\left[\sum_{1 \leq i \leq \xi n} Y_i \leq \xi^2/8n\right] \geq \Pr[\{A_{k+1} \leq \xi^2 n/8\} \cap \{A_k \geq \xi n\}].$$

By an application of Chernoff bound (13),

$$\Pr\left[\sum_{1 \leq i \leq \xi n} Y_i \leq \xi^2 n/8\right] \leq \exp\left(-\frac{\xi^2 n}{32}\right) \leq n^{-a}.$$

Combining the last two equations above proves (b). \square

C.1 Proof of Lemma 3

We first prove (a). If no agent is activated, no agent can be rejected (since both the agents’ level and max-level cannot increase above 0). Otherwise, at least one agent is activated. Note that levels are monotonically non-decreasing and that agents that never become active remain on level 0. Active agents either increase their level by 1 or become inactive whenever they initiate an interaction. Since the maximum level is bounded by φ_2 , eventually all active agents become inactive. Thus, once all agents that ever become active are inactive, any agent that reached the maximum level is not rejected.

Next we show (b). Note that $A_0 = n$, and that the A_1 agents that reach level at least 1 are exactly those agents that are elected in JE1. By Lemma 2(a), JE1 elects at least one agent, thus $A_1 \geq 1$.

Let $\xi_0 = n^{-1/2} \sqrt{\ln n}$. We apply Lemma 24 for all levels $0 < k < \varphi_2$ and all $\xi = \beta^i \xi_0 \leq 1$ for $i \geq 0$, where $\beta = 1.06$. Then, we apply a

union bound over all those $O(\log n)$ different combinations of k, ξ , to obtain that the event

$$\mathcal{E}' := \bigcap_{1 \leq k < \varphi_2, \xi = \beta^i \xi_0 \leq 1} (\{A_{k+1} < 2\xi^2 n\} \cup \{A_k > \xi n\}) \cap (\{A_{k+1} > \xi^2 n/8\} \cup \{A_k < \xi n\})$$

has probability $1 - o(1/\log n)$. It is easy to see that $\mathcal{E}' \subseteq \mathcal{E}$, where

$$\mathcal{E} := \bigcap_{1 \leq k < \varphi_2, \xi_0 \leq \xi \leq 1} (\{A_{k+1} < 3\xi^2 n\} \cup \{A_k > \xi n\}) \cap (\{A_{k+1} > \xi^2 n/9\} \cup \{A_k < \xi n\}),$$

as for each $\beta^{i-1} \xi_0 < \xi < \beta^i \xi_0$, the events for ξ in the definition of \mathcal{E} follow from the corresponding events for $\beta^{i-1} \xi_0$ or $\beta^i \xi_0$ in the definition of \mathcal{E}' . Thus,

$$\Pr[\mathcal{E}] \geq \Pr[\mathcal{E}'] = 1 - o(1/\log n).$$

Given \mathcal{E} , and the event that $A_1 \leq n^{1-\epsilon}$, we obtain

$$A_{1+i_{\max}} \leq \xi_0 n,$$

where i_{\max} is the number of times we must iteratively apply the function $f(x) = 3x^2$ to $x = n^{-\epsilon}$, before we reach a value less or equal to ξ_0 . Thus i_{\max} is bounded by a constant that depends only on ϵ . Note that we must choose $\varphi_2 \geq i_{\max} + 1$. Moreover, for $i_{\top} := \max\{i: A_i \neq 0\}$, we have

$$A_{i_{\top}} < \xi_0 n,$$

because otherwise \mathcal{E} would imply that $A_{\top} > \xi_0^2 n/9 > 0$, contradicting the definition of i_{\top} . This completes the proof of (b).

We now prove (c). For a constant $c > 0$ define $\tau_1 := c \cdot n \log n$ and let N_u be the number of interactions that agent u initiates during any interaction from $\{t_1 + 1, t_1 + 2, \dots, t_1 + \tau_1\}$. Let \mathcal{D} denote the event “ $\forall u: N_u \in [\tau_1/(2n), 2\tau_1/n]$.” As in the proof of Lemma 21, standard Chernoff bounds yield

$$\Pr[\overline{\mathcal{D}}] \leq n^{-c/6+1}.$$

Thus, for $c \geq 6a + 6$ we have that w.p.r. at least $1 - n^{-a}$, each agent initiates at least $c/2 \cdot \log n$ and at most $2c \cdot \log n$ interactions. Let u be an agent that is still active after interaction t_1 . While u is active, any interaction it initiates causes u to either increase its level or to become inactive. Since any agent reaching level φ_2 becomes inactive, u is guaranteed to be inactive as soon as it initiated at least φ_2 interactions after t_1 . For n large enough (or by choosing c large enough), we have $c/2 \cdot \log n \geq \varphi_2$ (remember that φ_2 is a constant). Thus, as long as \mathcal{D} occurs, all agents are inactive after $t_1 + \tau_1$ steps. By Lemma 20, during another $\tau_2 := 4(a+1) \cdot n \ln n$ steps, with probability at least $1 - 2n^{-a}$, the one-way epidemic for the agents’ max-level is completed and all agents have the same max-level (i.e., JE2 is completed). Combining both results yields that $t_2 \leq t_1 + \tau_1 + \tau_2 = t_1 + O(n \log n)$ w.h.p.

D ANALYSIS OF LSC

D.1 Proof Sketch of Lemma 4

The guarantees of Lemma 4 and their analysis are due to [24]; here we mostly show how their results – which use a slightly different set of notions and wording – map to our formulation.

The following lemma is an adaptation of [24, Lemma 3.7], which assumes that the internal clock counters are unbounded, thus avoids the modulo arithmetic. The lemma assumes that a junta of size $n^{1-\epsilon}$ drives the phase clock, for a constant $0 < \epsilon < 1$. Under this assumption, it states that, w.h.p., the clock counters of the agents stay close together.

LEMMA 25 ([24, LEMMA 3.7]). *For any constant $d > 0$ there is an integer constant $K > 0$ such that the following holds. Let c_{\max} denote the maximum and c_{\min} the minimum unbounded internal clock counters after an interaction $t \in \mathbb{N}$. If $c_{\max} - c_{\min} \leq 2K$, then w.h.p. there is a $t' > t + d \cdot n \log n$ such that:*

- (a) *Interaction t' is the first interaction after which the maximum internal clock counter is $c_{\max} + K$.*
- (b) *After interaction t' , all agents have internal clock counter at least c_{\max} .*

The same result holds for the external clock counters, if we count only those interactions in which an agent updates its external clock, i.e., interactions in which the second component of the initiator's state is ext. (Such interactions were called *meaningful* in [24].)

We are ready to sketch the proof of Lemma 4. We prove only the part for the internal clock. The part for the external clock works analogously by noticing that external clocks are updated exactly once per internal phase. Since we show that an internal phase has, w.h.p., a length of $\Omega(n \log n)$ interactions and a stretch of $O(n \log n)$ interactions, this results in the additional factor $\Theta(\log n)$.

We have to show that, w.h.p., $L_{\text{int}}(\rho) \geq d_1 \cdot n \log n$ and $S_{\text{int}}(\rho) \leq d_2 \cdot n \log n$ for an arbitrary given constant $c_1 > 0$ and some suitable constants $d_2 \geq d_1 \geq c_1$ for all $\rho \in \{0, 1, \dots, \log^2 n\}$. The lower bound on $L_{\text{int}}(\rho)$ follows via an induction over ρ by applying Lemma 25 for $d = d_1$ and by choosing the clock parameter m_1 such that $2m_1 + 1 \geq 6K$. This choice of m_1 and the synchronization guarantee of Lemma 25 imply that, w.h.p., agents can do arithmetic modulo $2m_1 + 1$ and still determine correctly which of two agents' internal clock counters is further ahead. For the upper bound on $S_{\text{int}}(\rho)$, note that the one-way epidemic (cf. Lemma 20) implies that, w.h.p., the maximum clock counter increases within $O(n \log n)$ interactions (when a clock agent finally sees the maximum clock counter). Thus, w.h.p., it takes at most $(2m_1 + 1) \cdot O(n \log n) = O(n \log n)$ interactions for an agent to leave a given phase, such that d_2 is given by the constant hidden in the Landau notation.

D.2 Proof of Lemma 5

Let $(\mathcal{F}_t)_{t \in \mathbb{N}_0}$ denote the filtration in which \mathcal{F}_t describes the outcome of the first t interactions. Let \mathcal{E}_t denote the event that there is at least one clock agent after the first t interactions and note that $\mathcal{E}_t \implies \mathcal{E}_{t'}$ for all $t' \geq t$. We will show that there is a constant $c > 0$ such that

$$\Pr[l'_2 > t + c \cdot n^2 \log^3 n \mid \mathcal{F}_t, \mathcal{E}_t] \leq n^{-1}. \quad (25)$$

Applying (25) repeatedly (via the chain rule), we obtain for $k \geq 1$,

$$\Pr[l'_2 > t + k \cdot c \cdot n^2 \log^3 n \mid \mathcal{E}_t] \leq n^{-k}.$$

From that, it follows

$$\mathbb{E}[l'_2 \mid \mathcal{F}_t, \mathcal{E}_t] \leq t + 2c \cdot n^2 \log^3 n,$$

which implies the lemma.

It remains to show (25). To simplify the exposition, we assume w.l.o.g. that $t = 0$. Define an *epoch* as a sequence of consecutive interactions during which each agent initiates at least one *external clock interaction*, that is, an interaction in which the second component of the initiator's state is ext right before the interaction. Consider a partitioning of the time horizon into epochs of minimal length. We first show that, with probability at least $1 - n^{-2}$, all agents reach external phase 2 within the first $O(\log n)$ epochs. Afterward we show that, with probability at least $1 - n^{-2}$, an epoch requires only $O(n^2 \log^2 n)$ interactions. Using a union bound, these two statements imply (25).

Let p_{\max} denote the maximal external clock counter among all agents after a given interaction and assume $p_{\max} < 2m_2$. Note that the number of epochs until *all* agents have external clock counter at least p_{\max} is stochastically dominated by the runtime of synchronous PULL rumor spreading (the rumor being external clock counter values $\geq p_{\max}$; see [27]). It follows that, w.h.p., all agents have external clock counter at least p_{\max} after $O(\log n)$ epochs. After one more epoch, some clock agent v has increased its external clock counter to at least $p_{\max} + 1$. This and $m_2 = O(1)$ yields that, w.h.p., the maximum external clock counter reaches value $2m_2$ after $O(\log n)$ epochs. Using again the dominance by PULL rumor spreading, after another $O(\log n)$ epochs all agents have external clock counter $2m_2$ (i.e., have reached external phase 2).

It remains to show the second part mentioned above. That is, w.h.p., an epoch ends within $O(n^2 \log^2 n)$ interactions. Before we prove that, it will be useful to consider the event $\mathcal{E}_{u,v}$ that an agent u initiates an interaction before agent v within τ consecutive interactions, where $\tau := d \cdot n \ln n$, and $d > 0$ is a constant to be fixed later. The number of interactions until an agent initiates an interaction is a geometric random variable $\text{Geom}(1/n)$. Thus, the number of interactions until either u or v initiates an interaction is the minimum of two geometric random variables and distributed according to $\text{Geom}(p)$ with $p := 1 - (1 - 1/n)^2$. By symmetry, u and v are equally likely to initiate that interaction. This yields $\Pr[\mathcal{E}_{u,v}] = (1 - (1 - p)^\tau) / 2 \geq (1 - n^{-2d}) / 2 \geq 1/4$.

We now bound the number of interactions required to end an epoch. In the following we say an agent u makes *progress* whenever it increases its internal clock counter (in the modulo arithmetic). Fix an agent u with internal clock counter t_{int} . We first show that the probability for u to make progress within $O(1) \cdot \tau$ interactions is not too small. To this end, consider two cases:

- (i) There is an agent v with internal clock counter t'_{int} such that $0 < t'_{\text{int}} - t_{\text{int}} \leq m_1$ or $t_{\text{int}} - t'_{\text{int}} > m_1$. In this case, if $\mathcal{E}_{u,v}$ occurs and v is chosen as a responder, then u makes progress. As seen above, $\Pr[\mathcal{E}_{u,v}] \geq 1/4$ and, independently, v is chosen as a responder during the interaction initiated by u with probability $\geq 1/n$. Thus, u makes progress within τ interactions with probability at least $1/(4n)$.
- (ii) There is no agent v with internal clock counter t'_{int} such that $0 < t'_{\text{int}} - t_{\text{int}} \leq m_1$ or $t_{\text{int}} - t'_{\text{int}} > m_1$. In this case, the only way that some node progresses to an internal clock counter t'_{int} as above is that some *clock* agent initiates an interaction with an agent that has internal clock counter t_{int} ; in this case, $t'_{\text{int}} = t_{\text{int}} + 1 \pmod{2m_1 + 1}$. The number of interactions for this to occur is stochastically dominated

by a one-way epidemic (the infected agents being those with internal clock counter t_{int}). Thus, if \mathcal{E} denotes the event that some clock agent's internal clock counter reaches value t'_{int} as above within τ interactions, then [Lemma 20](#) implies $\Pr[\mathcal{E}] \geq 1 - 1/n \geq 1/2$ for large enough n and d .

Now assume \mathcal{E} occurs and let v denote the first clock agent with internal clock counter t'_{int} . If $u = v$, agent u made progress. Otherwise, we are in the situation of case (i) and know that u makes progress within τ interactions with probability at least $1/(4n)$.

We conclude that u makes progress within 2τ interactions with probability at least $1/2 \cdot 1/(4n) \geq 1/(8n)$.

Combining both cases, we see that the probability that u makes progress within 2τ interactions is at least $1/(8n)$. Thus, for any $a > 0$ we find a constant $d' > 0$ such that the probability that u makes progress within $(d' \cdot n \ln n) \cdot 2\tau$ interactions is at least $1 - n^{-(a+1)}$. By a union bound, with probability at least $1 - n^{-a}$ each agent makes progress within $O(n^2 \log^2 n)$ interactions. Since $m_1 = O(1)$, we get that, w.h.p., all agents make progress at least $2m_1 + 1$ times (and, thus, initiate an external clock interaction) within $O(n^2 \log^2 n)$ interactions, yielding the desired bound on the epoch length.

E ANALYSIS OF DES

E.1 Proof of [Lemma 6\(a\)](#)

Suppose, for contradiction, that all agents are rejected in DES. An agent gets rejected in DES, i.e., reaches state \perp , only if at some step it is in state 0 and interacts (as an initiator) with an agent in state 2 or \perp . Let t be the step when the first agent reaches state \perp . Then the responder agent v at step t is either in state 2 or in state \perp right before step t . In both cases we have a contradiction: If v is in state \perp before t , this contradicts the definition of t ; if it is in state 2 then v is never rejected in DES.

E.2 Proof of [Lemma 6\(b\)](#): Upper Bound

For any $s \in \mathcal{S}_{\text{DES}}$, let $N_t(s)$ denote the set of agents in state s right after step t , and let $n_t(s) := |N_t(s)|$. We also extend this notation to pairs of states: For distinct $s_1, s_2 \in \mathcal{S}_{\text{DES}}$, we define $N_t(s_1, s_2) := N_t(s_1) \cup N_t(s_2)$, and $n_t(s_1, s_2) := n_t(s_1) + n_t(s_2)$. Let

$$\begin{aligned} t_1 &:= \min \{ t : N_t(1) \neq \emptyset \} \\ t_2 &:= \min \{ t : N_t(2) \neq \emptyset \} \\ t_3 &:= \min \{ t : N_t(\perp) \neq \emptyset \} \\ t_4 &:= \min \{ t : N_t(0) = \emptyset \}. \end{aligned}$$

Note that t_4 is the completion time of DES. Let S denote the set of agents elected in JE2. Since the lemma assumes that JE2 is completed before step f_1 , set S is finalized before f_1 . We also assume $|S| = O((n \log n)^{1/2})$, and from [Lemma 3\(a\)](#), $|S| > 0$.

CLAIM 26. $\Pr[n_{t_3}(1, 2) > c \cdot (n \ln n)^{1/2}] = o(1/\log n)$, for some constant c .

PROOF. Let

$$\begin{aligned} r_1 &:= \min \{ t : n_t(1) \geq 2(n \ln n)^{1/2} \} \cup \{r_2\} \\ r_2 &:= \min \{ t : n_t(2) \geq \ln n \} \cup \{t_3, r_4\} \\ r_4 &:= \min \{ t : n_t(1, 2) \geq n^{3/4} \}. \end{aligned}$$

First we bound the probability that $r_2 > r_1 + 4n$. For any step $r_1 < t \leq r_2$, $n_{t-1}(1) \geq 2(n \ln n)^{1/2} - \ln n$; thus, the probability that $n_t(2) = n_{t-1}(2) + 1$, which is the probability that two agents in state 1 interact at step t , is

$$\frac{n_{t-1}(1)(n_{t-1}(1) - 1)}{n(n-1)} \geq \frac{\ln n}{n}.$$

It follows, by applying Chernoff bound [\(13\)](#), that

$$\Pr[r_2 > r_1 + 4n] = O(1/n).$$

Next we bound the probability that $\min\{t_3, r_4\} > r_2 + 5n$. For any $r_2 < t \leq \min\{t_3, r_4\}$, $n_{t-1}(0) \geq n - n^{3/2}$ and $n_{t-1}(2) \geq \ln n$; thus, the probability that $N_t(\perp) \neq \emptyset$, which is the probability that at step t an agent in state 0 interacts with an agent in state 2 and moves to state \perp , is

$$\frac{(1/4) \cdot n_{t-1}(0) \cdot n_{t-1}(2)}{n(n-1)} \geq \frac{\ln n}{5n}.$$

Then

$$\Pr[\min\{t_3, r_4\} > r_2 + 5n] \leq \left(1 - \frac{\ln n}{5n}\right)^{5n} = O(1/n).$$

From that and the bound we showed earlier on r_2 , it follows

$$\Pr[\min\{t_3, r_4\} > r_1 + 9n] = O(1/n).$$

We now bound $n_{r_1+9n}(1, 2)$. The number of agents in states 1 and 2 can increase in a step for two reasons: either because an agent in state 0 interacts with an agent in state 1, or because an agent in state 0 gets selected in JE2. To avoid having to deal with the timing of the latter transitions we define $n'_t(1, 2) := |N_t(1, 2) \cup S|$. Recall, S is the set of agents elected in JE2, which is fixed before f_1 . For each $t > r_1$, we have $n'_t(1, 2) = n'_{t-1}(1, 2) + 1$ only if at step t an agent in state 0 interacts with an agent in state 1 and its state changes to 1. Thus the probability that $n'_t(1, 2) = n'_{t-1}(1, 2) + 1$ is at most

$$\frac{(1/4) \cdot n_{t-1}(0) \cdot n_{t-1}(1)}{n(n-1)} \leq \frac{n_{t-1}(1)}{4n} \leq \frac{n'_{t-1}(1)}{4n}.$$

Let $k := 2(n \ln n)^{1/2} + |S|$, and note that $k \geq n'_{r_1}(1, 2)$. Then, the probability that $n'_{r_1+9n}(1, 2) > 13k$, i.e., $n'_{r_1+9n}(1, 2)$ increases from at most k to more than $13k$ in the $9n$ steps $r_1 < t \leq r_1 + 9n$, is upper bounded by the probability that the sum $C_{k-1, 13k, 4n}$ of independent geometric random variables with mean values $4n/k, 4n/(k+1), \dots, 4n/(13k)$ is at most $9n$:

$$\begin{aligned} \Pr[n'_{r_1+9n}(1, 2) > 13k] &\leq \Pr[C_{k-1, 13k, 4n} \leq 9n] \\ &\leq \Pr[4nH(k-1, 13k) - C_{k-1, 13k, 4n} > n] \\ &\leq O(1/k) = O((n \log n)^{-1/2}), \end{aligned}$$

where for the second inequality we used that $H(k-1, 13k) \geq \ln(13k/k) > 10/4$, and the last inequality follows from [Lemma 18\(a\)](#). Combining that and the bound on $\min\{t_3, r_4\}$ above, we obtain for the event $\mathcal{E} := \{\min\{t_3, r_4\} \leq r_1 + 9n\} \cap \{n'_{r_1+9n}(1, 2) \leq 13k\}$ that

$$\Pr[\mathcal{E}] = 1 - o(1/\log n).$$

Given \mathcal{E} , we have $n'_{\min\{t_3, r_4\}}(1, 2) \leq n'_{r_1+9n}(1, 2) \leq 13k$, and thus $n_{\min\{t_3, r_4\}}(1, 2) \leq 13k$. From this and the definition of r_4 , it follows that $\min\{t_3, r_4\} = t_3$, thus $n_{t_3}(1, 2) = n_{\min\{t_3, r_4\}}(1, 2) \leq 13k$. Therefore, $\mathcal{E} \subseteq \{n_{t_3}(1, 2) \leq 13k\}$, and

$$\Pr[n_{t_3}(1, 2) \leq 13k] \geq \Pr[\mathcal{E}] = 1 - o(1/\log n).$$

Since $k = O((n \log n)^{1/2})$, the claim follows. \square

Let us now fix step t_3 . We call step $t > 0$ a *zero-step* if the initiator agent of the interaction in step t is in state 0 before the step. Let i^* be the total number of zero-steps $t > t_3$. For $1 \leq i \leq i^*$, let z_i be the i th zero-step $t > t_3$, and let $z_0 := t_3$. For $0 \leq i \leq i^*$, let

$$a_i := n_{z_i}(1, 2), \quad b_i := n_{z_i}(\perp),$$

and for $i > i^*$, let $a_i := a_{i^*}$ and $b_i := b_{i^*}$. Note that $a_0 = n_{t_3}(1, 2)$, $b_0 = 1$, $z_{i^*} = t_4$, and $a_{i^*} + b_{i^*} = n$.

CLAIM 27. $\Pr[i^* > \sigma] = O(1/\log n)$, for $\sigma := n \ln n + n \ln \ln n$.

PROOF. For each $0 < i \leq i^*$ and $0 < k < n$,

$$\Pr[b_i = b_{i-1} + 1 \mid b_{i-1} = k] \geq k/(n-1),$$

because the probability that the responder in step z_i is in state \perp , and thus the initiator's state 0 changes to \perp , is $n_{z_i-1}(\perp)/(n-1) = b_{i-1}/(n-1)$. Above we have inequality instead of equality because the initiator's state may change to \perp also when the responder is in state 2. It follows

$$\Pr[i^* > \sigma] \leq \Pr[C_{0, n-1, n-1} > \sigma] = O(1/\log n),$$

by Lemma 18(b). \square

CLAIM 28. If $a_0 \leq c \cdot (n \ln n)^{1/2}$ for some constant $c > 0$, then for $\sigma := n \ln n + n \ln \ln n$, $\Pr[a_\sigma > n^{3/4} \ln n] = o(1/\log n)$.

PROOF. Let $a'_i := |N_{z_i}(1, 2) \cup S|$, and note that $a_i \leq a'_i < a_i + |S|$. As before, we work with a'_i instead of a_i to avoid dealing explicitly with the timing of external transitions $0 \Rightarrow 1$. For each $0 < i \leq i^*$ and $a'_0 < k < n-1$,

$$\Pr[a'_i = a'_{i-1} + 1 \mid a'_{i-1} = k] \leq k/(4(n-1)),$$

because the probability that the responder in step z_i is in state 1 or 2 is $n_{z_i-1}(1, 2)/(n-1) \leq a'_{i-1}/(n-1)$, and if the initiator is not in S , its state changes from 0 to 1 with probability $1/4$. The above inequality holds also trivially if $i > i^*$, since then $a_i = a_{i^*}$. It follows

$$\Pr[a'_\sigma > n^{3/4} \ln n] \leq \Pr[C_{k_1-1, k_2, 4(n-1)} \leq \sigma],$$

where $k_1 := c \cdot (n \ln n)^{1/2} + |S|$ and $k_2 := n^{3/4} \ln n$. We can bound the right side by using Lemma 18(a): We have $H(k_1 - 1, k_2) > \ln(k_2/k_1) > (1/4) \ln n + (1/3) \ln \ln n$, thus,

$$\sigma = n \ln n + n \ln \ln n < 4(n-1)H(k_1 - 1, k_2) - (1/4)(n-1) \ln \ln n.$$

From that and Lemma 18(a), it follows

$$\Pr[C_{k_1-1, k_2, 4(n-1)} \leq \sigma] = O(1/k_1) = o(\log n).$$

Therefore, $\Pr[a'_\sigma > n^{3/4} \ln n] = o(1/\log n)$. Since $a_\sigma \leq a'_\sigma$, the claim follows. \square

From Claims 27 and 28, and a union bound, it follows that if $a_0 \leq c \cdot (n \ln n)^{1/2}$, then $\Pr[a_{i^*} > n^{3/4} \ln n] = O(1/\log n)$. Also, from Claim 26, $\Pr[a_0 \leq c \cdot (n \ln n)^{1/2}] = 1 - o(1/\log n)$. Therefore, $\Pr[a_{i^*} > n^{3/4} \ln n] = O(1/\log n)$, which implies the lemma.

E.3 Proof of Lemma 6(b): Lower Bound

We will use the notation from Appendix E.2. The next claim implies that to compute a lower bound on the number of agents selected in DES, it suffices to assume that the set S of agents elected in JE2 has size one. Recall that S is fixed before step f_1 , thus before any agent reaches state 1. Recall also that $t_4 := \min\{t : N_t(0) = \emptyset\}$.

CLAIM 29. For all $k \geq 0$,

$$\Pr[n_{t_4}(1, 2) \geq k \mid |S| = 1] \leq \Pr[n_{t_4}(1, 2) \geq k \mid |S| > 1].$$

PROOF. Fix a $k \geq 0$. We will write $x = (x_0, x_1, x_2, x_\perp)$ to denote a configuration with x_s agents in state $s \in \mathcal{S}_{\text{DES}}$. Let E_x be an execution of DES starting from configuration x , in which no external transitions are allowed. Let $p(x)$ be the probability that $n_{t_4}(1, 2) \geq k$ in E_x . We will often write $p(x_0, x_1, x_2, x_\perp)$ instead of $p(x)$.

For two configuration x, y , we will prove that if the following conditions hold,

$$x_1 \geq y_1, \quad x_\perp \leq y_\perp, \quad x_1 - y_1 \geq \max\{y_0 - x_0, y_2 - x_2\},$$

then $p(x) \geq p(y)$. By transitivity, it suffices to prove the following: (i) $p(x) \geq p(x_0, x_1 - 1, x_2 + 1, x_\perp)$, (ii) $p(x) \geq p(x_0, x_1, x_2 - 1, x_\perp + 1)$, (iii) $p(x) \geq p(x_0 + 1, x_1 - 1, x_2, x_\perp)$, (vi) $p(x) \leq p(x_0 + 1, x_1, x_2, x_\perp - 1)$. The proof is by induction on $x_0 \in \{0, \dots, n-1\}$.

We denote by y the second configuration in each of (i)–(iv).

For $x_0 = 0$, the final number k_x of agents in states 1 and 2 in E_x is $k_x = x_1 + x_2$. For the corresponding number k_y in E_y , depending on the case, we have: (i) $k_y = x_1 + x_2$; (ii) $k_y = x_1 + x_2 - 1$; (iii) $k_y \leq x_1 + x_2$, as the remaining agent in state 0 may or may not reach states 1 and 2; and, for the same reason, (iv) $k_y \geq x_1 + x_2$. This proves the base case.

For the inductive step, suppose that statements (i)–(iv) hold for $x_0 = i - 1$, for some $1 \leq i < n - 1$; we show they hold for $x_0 = i$.

We consider the cases (i) and (ii) first. In both cases, we couple E_x and E_y such that the same pairs of agents interact in the two executions, and the same random coins are used. Also, every agent starts in the same state in the two executions, except for agent u . In (i), u starts in state 1 in E_x and in state 2 in E_y ; in (ii), u starts in state 2 in E_x and in state \perp in E_y . In both cases we show that after the same number of steps, configurations x' and y' are reached in E_x and E_y , respectively, such that either $x' = y'$, or the induction hypothesis implies $p(x') \geq p(y')$. In either case we conclude that $p(x) \geq p(y)$.

Let t be the first step in which u interacts (as an initiator or responder), and in addition the state of the initiator changes at that step in at least one of the two executions. It is easy to see that until right before that point the state of any agent $v \neq u$ is the same in the two executions. If no step t as above exists then the number of agents in state 0 must become $k - 1$ at some point t' , and the induction hypothesis implies $p(x') \geq p(y')$, for the configurations x', y' reached in E_x, E_y , respectively, right after step t' .

Suppose now that t exists, and consider case (i) first. In this case, the following transitions are possible at step t : (1) Agent u interacts (as an initiator or responder) with an agent in state 1, and in E_x the initiator moves to state 2, while in E_y there is no transition. Then both executions are in the same configuration after step t . (2) Agent u is the responder, and the initiator v is in state 0. Then with probability $1/4$, v moves to state 1 in both

executions; with probability $1/2$, v stays in state 0 in both; and with probability $1/4$, v moves to state \perp in E_y and stays in state 0 in E_x . We can now apply the induction hypothesis to obtain $p(x') \geq p(y')$ for the configurations reached after the step. E.g., in the very last subcase, $x' = x$ and $y' = (x_0 - 1, x_1 - 1, x_2 + 1, x_\perp + 1)$, thus $p(x') = p(x) \geq p(x_0 - 1, x_1, x_2, x_\perp + 1) \geq p(y')$, from induction hypothesis (iv) and (i).

In case (ii), u must be the responder, and the initiator v must be in state 0 before t . In E_x , v 's state changes to 1, or to \perp , or does not change. In E_y , v 's state changes to \perp . After the step, we can again apply the induction hypothesis to obtain $p(x') \geq p(y')$ for the configurations reached.

For case (iii), we have that from configuration $y = (x_0 + 1, x_1 - 1, x_2, x_\perp)$ the only possible transitions are to one of the three configuration $x, z = (x_0, x_1 - 1, x_2 + 1, x_\perp)$, or $w = (x_0, x_1 - 1, x_2, x_\perp + 1)$, and we have just shown in the induction step that $p(x) \geq p(z)$ and $p(x) \geq p(w)$. It follows that $p(x) \geq p(y)$. The proof for case (iv) is similar. This completes the induction proof.

Finally, we observe that an external transitions $0 \Rightarrow 1$ just changes a configuration x to $y = (x_0 - 1, x_1 + 1, x_2, x_3)$, and we have shown that $p(y) \geq p(x)$. The claim then follows by taking any execution, with multiple transitions $0 \Rightarrow 1$, and remove them one by one (starting from the last one), keeping just the first one, and applying the above observation. \square

Claim 29 allows us to assume $|S| = 1$ for the rest of the proof.

Recall that $t_2 := \min \{t : N_t(2) \neq \emptyset\}$. The next statement is the lower-bound version of **Claim 26**.

CLAIM 30. $\Pr[n_{t_2}(1, 2) \leq (n/\ln n)^{1/2}] = O(1/\log n)$.

PROOF. Recall that $t_1 := \min \{t : N_t(1) \neq \emptyset\}$. We call step $t > 0$ a *transition-step* if the state of the initiator agent changes in that step. Let r_i , for $i \geq 1$, be the i th transition-step $t > t_1$, and let $r_0 = t_1$. Let \tilde{i} be number of transition-steps $t > t_1$ until the first agent reaches state 2, i.e., $r_{\tilde{i}} = t_2$. For any $0 < i \leq \tilde{i}$,

$$\Pr[n_{r_i}(2) = 1 \mid n_{r_{i-1}}(2) = 0] = \frac{i(i-1)}{i(i-1) + (n-i)i/4} \leq \frac{4i}{n+3i},$$

because $n_{r_{i-1}}(1) = i$, and in step r_i the number of agents in state 2 increases when two agents in state 1 interact, while the number of agents in state 1 increases when an agent in state 0 interacts with an agent in state 1 and the coin with success probability $1/4$ is favorable. From the bound above, the probability that no agent reaches state 2 in the first $(n/\ln n)^{1/2}$ transition-steps is at least

$$\prod_{1 \leq i \leq (n/\ln n)^{1/2}} \left(1 - \frac{4i}{n+3i}\right) \geq 1 - \sum_{1 \leq i \leq (n/\ln n)^{1/2}} \frac{4i}{n+3i} \geq 1 - \frac{2}{\ln n}.$$

The claim then follows \square

Let

$$\tau_1 := \min \{t : n_t(1, 2) = (n/\ln n)^{1/2}\} \cup \{t_4\}$$

$$\tau_2 := \min \{t : n_t(0) \leq n/2\}.$$

Recall that step $t > 0$ is a *zero-step* if the initiator agent in that step is in state 0 before the step. Let j^* be the total number of zero-steps $t > \tau_1$. For $1 \leq j \leq j^*$, let ϱ_j be the j th zero-step $t > \tau_1$, and let $\varrho_0 := \tau_1$ (for $j > j^*$, $\varrho_j = \infty$). For $0 \leq j \leq j^*$, let

$$d_j := n_{\varrho_j}(1, 2), \quad e_j := n_{\varrho_j}(2), \quad f_j := n_{\varrho_j}(\perp);$$

and for $j > j^*$, let $d_j := d_{j^*}$, $e_j := e_{j^*}$, and $f_j := f_{j^*}$. Note that $d_0 = n_{\tau_1}(1, 2)$, $\varrho_{j^*} = t_4$, and $d_{j^*} + f_{j^*} = n$. Also, if $t_2 > \tau_1$ then $d_0 = (n/\ln n)^{1/2}$ and $e_0 = f_0 = 0$. Let

$$\lambda := n^{3/4}(\ln \ln n)^{1/4}(\ln n)^{-3/4}$$

$$\sigma := n \ln n - n \ln \ln n + n \ln \ln \ln n + n.$$

In the next claims, we compute upper and lower bounds for the quantities d_j , e_j , and f_j .

CLAIM 31. *If $t_2 > \tau_1$ then*

$$\Pr[\{d_\sigma < \lambda\} \cap \{j^* \geq \sigma\}] = o(1/\log n).$$

PROOF. For any $0 < j \leq j^*$ and $0 < k < n$,

$$\Pr[d_j = d_{j-1} + 1 \mid d_{j-1} = k] = k/(4(n-1)), \quad (26)$$

because the probability that the responder in step ϱ_j is in state 1 or 2 is $n_{\varrho_{j-1}}(1, 2)/(n-1) = d_{j-1}/(n-1)$, and in this case the initiator's state changes from 0 to 1 with probability $1/4$. It follows

$$\Pr[\{d_\sigma < \lambda\} \cap \{j^* \geq \sigma\}] \leq \Pr[C_{d_0-1, \lambda, 4(n-1)} > \sigma].$$

We can bound the right side using **Lemma 18(a)**: Since $t_2 > \tau_1$, we have $d_0 = (n/\ln n)^{1/2}$. Then,

$$\begin{aligned} H(d_0 - 1, \lambda) &\leq \ln(\lambda/(d_0 - 1)) \\ &\leq (1/4) \ln n + (1/4) \ln \ln \ln n - (1/4) \ln \ln n + o(1), \end{aligned}$$

thus, $\sigma \geq 4(n-1)H(d_0 - 1, \lambda) + n/2$. From **Lemma 18(a)**, then $\Pr[C_{d_0-1, \lambda, 4(n-1)} > \sigma] = O(1/d_0) = o(1/\log n)$. Substituting this above we obtain $\Pr[\{d_\sigma < \lambda\} \cap \{j^* \geq \sigma\}] = o(1/\log n)$. \square

CLAIM 32. *For any $j \geq 0$ and for $k_1 := 2 \cdot e^{j/(4(n-1))}d_0$,*

$$\Pr[d_j > k_1] = O(1/d_0).$$

PROOF. From (26) and the fact that $d_j = d_{j-1}$ for $j > j^*$, it follows

$$\Pr[d_j > k_1] \leq \Pr[C_{d_0-1, k_1, 4(n-1)} \leq j].$$

The right side is at most

$$\Pr[C_{d_0-1, k_1, 4(n-1)} < 4(n-1)H(d_0 - 1, k_1) - n],$$

since $H(d_0 - 1, k_1) \geq \ln(k_1/d_0) = j/(4(n-1)) + \ln 2$. The probability above is $O(1/d_0)$, from **Lemma 18(a)**. \square

CLAIM 33. *For any $j \geq 0$ and for $k_2 := 32k_1^2/n + 1$, where k_1 is defined as in **Claim 32**,*

$$\Pr[\{e_j > k_2\} \cap \{\varrho_j \leq \tau_2\}] = O(1/d_0) + e^{-16k_1^2/(3n)}.$$

PROOF. We call step $t > 0$ a *good-step* if the state of the initiator changes to either 1 or 2 in that step. Let r_i , for $i \geq 1$, be the i th good-step $t > t_1$, or $r_i = \infty$ if no such step exists; let also $r_0 = t_1$. For any $i > 0$ and $j \geq 0$, if $r_i < \infty$ and the number of agents in state 0 before the step is at least $n/2$, the probability some agent reaches state 2 in this step is

$$\begin{aligned} &\Pr[n_{r_i}(2) = j + 1 \mid n_{r_{i-1}}(2) = j] \\ &= \frac{n_{r_{i-1}}(1) \cdot (n_{r_{i-1}}(1) - 1)}{(1/4)n_{r_{i-1}}(0) \cdot n_{r_{i-1}}(1, 2) + n_{r_{i-1}}(1) \cdot (n_{r_{i-1}}(1) - 1)} \\ &\leq \frac{n_{r_{i-1}}(1) - 1}{n_{r_{i-1}}(0)/4} \leq \frac{i - 2j - 1}{n/8} \leq \frac{i - 1}{n/8}, \end{aligned}$$

because $n_{r_{i-1}}(1) = i - 2j$, $n_{r_{i-1}}(0) \geq n/2$, and in step r_i the number of agents in state 2 increases when two agents in state 1 interact, while the number of agents in state 1 increases when an agent in state 0 interacts with an agent in state 1 or 2 and the coin with success probability $1/4$ is favorable. It follows that for $\tilde{i} := \max\{i' : r_{i'} \leq \tau_2\}$, and any $i > 0$,

$$\mathbb{E}[n_{r_{\min\{i,\tilde{i}\}}}(2)] \leq \sum_{1 \leq j \leq i} \frac{j-1}{n/8} \leq 4i^2/n.$$

Since the probability bound above on $n_{r_i}(2)$ holds independently of $n_{r_{i'}}(2)$, for $i' < i$, we can apply Chernoff bound (12) to obtain

$$\Pr[n_{r_{\min\{i,\tilde{i}\}}}(2) > 2 \cdot 4i^2/n] < e^{-4i^2/(3n)}.$$

Setting $i = 2k_1$, where $k_1 = 2e^{j/(4(n-1))}d_0$, gives

$$\Pr[n_{r_{\min\{2k_1,\tilde{i}\}}}(2) > 32k_1^2/n] < e^{-16k_1^2/(3n)}. \quad (27)$$

From Claim 32, we have

$$\Pr[d_j + e_j < 2k_1 - 1] = 1 - O(1/d_0),$$

because $e_j \leq d_j - 1$. If $d_j + e_j < 2k_1 - 1$, $\varrho_j \leq \tau_2$, and $r_{2k_1} < \infty$, then

$$n_{r_{2k_1}}(1, 2) + n_{r_{2k_1}}(2) = 2k_1 - 1 > d_j + e_j = n_{\varrho_j}(1, 2) + n_{\varrho_j}(2),$$

thus $r_{2k_1} > \varrho_j$, as $n_t(1, 2)$ and $n_t(2)$ are non-decreasing. It follows

$$\Pr[\{r_{2k_1} > \varrho_j\} \cup \{\varrho_j > \tau_2\} \cup \{r_{2k_1} > \tau_2\}] = 1 - O(1/d_0).$$

This implies

$$\Pr[\{r_{\min\{2k_1,\tilde{i}\}+1} > \varrho_j\} \cup \{\varrho_j > \tau_2\}] = 1 - O(1/d_0),$$

because if $\varrho_j \leq \tau_2$ and $r_{\min\{2k_1,\tilde{i}\}+1} \leq \varrho_j$ then $r_{\min\{2k_1,\tilde{i}\}+1} \leq \tau_2$, and from that, $\min\{2k_1,\tilde{i}\} \neq \tilde{i}$, since $r_{\tilde{i}+1} > \tau_2$ by the definition of \tilde{i} , and thus $r_{2k_1} < r_{2k_1+1} = r_{\min\{2k_1,\tilde{i}\}+1} \leq \tau_2$. From (27),

$$\Pr[n_{r_{\min\{2k_1,\tilde{i}\}+1}}(2) > 1 + 32k_1^2/n] < e^{-16k_1^2/(3n)}.$$

Combining the last two equations above, and using the fact that, if $r_{\min\{2k_1,\tilde{i}\}+1} \leq \varrho_j$, then $e_j = n_{r_{\varrho_j}}(2) \geq n_{r_{\min\{2k_1,\tilde{i}\}+1}}(2)$, we obtain

$$\Pr[\{e_j > 1 + 32k_1^2/n\} \cap \{\varrho_j \leq \tau_2\}] < e^{-16k_1^2/(3n)} + O(1/d_0). \quad \square$$

The next statement is immediate from Claim 33, and the fact that $d_0 = (n/\ln n)^{1/2}$ if $t_2 > \tau_1$.

COROLLARY 34. *If $t_2 > \tau_1$ then for $k_3 = \frac{2^6 e^{\frac{j}{2(n-1)}}}{\ln n}$ and any $j \geq 0$,*

$$\Pr[\{e_j > 2k_3 + 1\} \cap \{\varrho_j \leq \tau_2\}] = O((n/\log n)^{-1/2}) + e^{-k_3/3}.$$

To bound f_j , we first define a random sequence $g = (g_0, g_1, \dots)$, which we compare to the sequence of f_j . We have $g_0 := 1$, and for $j > 0$, either $g_j = g_{j-1}$ or $g_j = g_{j-1} + 1$, such that

$$\Pr[g_j = g_{j-1} + 1 \mid g_{j-1}, \dots, g_1] = \frac{g_{j-1} + c_g(\ln n)^{-1} e^{\frac{j}{2(n-1)}} - 3/4}{n-1},$$

where $c_g := 2^6$. If the right side above exceeds 1, we implicitly assume it is equal to 1.

LEMMA 35. *If $t_2 > \tau_1$ then there is a coupling of sequence g with the sequence of f_i such that $\Pr[\bigcap_{j \geq 0} \{g_j \geq f_j\} \cup \{\varrho_j > \tau_2\}] = 1 - O(1/\log n)$.*

PROOF. Suppose that $t_2 > \tau_1$. We will couple the two sequences above for all $j \geq 0$ for which $\varrho_j \leq \tau_2$. The coupling consists of two parts. The first part is for $0 \leq j \leq n(\ln \ln n)^2$, and the second for $j \geq n(\ln \ln n)^2$. For $j \geq 1$, let

$$g_j^* := c_g(\ln n)^{-1} e^{\frac{j}{2(n-1)}},$$

and note that $\Pr[g_j = g_{j-1} + 1 \mid g_{j-1}] = (g_{j-1} - 3/4 + g_j^*)/(n-1)$. Let $j_{\text{mid}} := n(\ln \ln n)^2$.

Coupling for $0 \leq j \leq j_{\text{mid}}$. The coupling is as follows. For $j \geq 0$, we define the next event under the coupling,

$$\begin{aligned} \mathcal{E}_j^1 := & \{ \varrho_j < \tau_2 \} \cap \{ g_j - 3/4 + g_{j+1}^* < n-1 \} \\ & \cap \{ g_j \geq f_j + e_j/4 + 3/4 \} \cap \{ d_j^2/n \leq g_{j+1}^* \}. \end{aligned}$$

For each $1 \leq j \leq j_{\text{mid}}$, if $\varrho_{j-1} < \tau_2$, we first execute the steps $\varrho_{j-1} < t \leq \varrho_j$ of the protocol. After that, if event \mathcal{E}_{j-1}^1 holds, we select g_j according to one of the cases below.

- (i) If $f_j + e_j = f_{j-1} + e_{j-1} + 1$, then we set $g_j = g_{j-1} + 1$.
- (ii) If $f_j + e_j = f_{j-1} + e_{j-1}$, then we set $g_j = g_{j-1} + 1$ with probability

$$q_j := \frac{g_{j-1} - 3/4 + g_j^*}{n-1} \cdot \left(1 + \frac{p_j^1}{p_j^0} \right) - \frac{p_j^1}{p_j^0},$$

where for $i \in \{0, 1\}$,

$$p_j^i := \Pr[f_j + e_j = f_{j-1} + e_{j-1} + i \mid \mathcal{E}_{j-1}^1].$$

With the remaining probability, $1 - q_j$, we set $g_j = g_{j-1}$.

- (iii) If $f_j + e_j \geq f_{j-1} + e_{j-1} + 2$, then we choose the value of g_j independently (according to its distribution).

If \mathcal{E}_{j-1}^1 does not hold, again we choose g_j 's value independently.

We must show that the marginal distribution of g in the coupling above is the correct one. In Claim 36 we prove that if \mathcal{E}_{j-1}^1 holds, then $0 \leq q_j \leq 1$. Given that, we have that if \mathcal{E}_{j-1}^1 occurs, then the probability that $g_j = g_{j-1} + 1$ is

$$p_j^1 + p_j^0 q_j + (1 - p_j^1 - p_j^0) \cdot \frac{g_{j-1} - 3/4 + g_j^*}{n-1} = \frac{g_{j-1} - 3/4 + g_j^*}{n-1},$$

which is the right probability.

CLAIM 36. *For any $1 \leq j \leq j_{\text{mid}}$, if \mathcal{E}_{j-1}^1 occurs then $0 \leq q_j \leq 1$.*

PROOF. Suppose that \mathcal{E}_{j-1}^1 occurs. The probability that $e_j > e_{j-1}$, i.e., between zero-steps ϱ_{j-1} and ϱ_j there is at least one interaction between two agents in state 1, is

$$\begin{aligned} \Pr[e_j > e_{j-1}] & \leq \mathbb{E}[e_j - e_{j-1}] \\ & \leq \sum_{i \geq 1} 2^{-i-1} \frac{d_{j-1}(d_{j-1} - 1)}{n(n-1)} \leq \frac{d_{j-1}^2}{2n^2} \leq \frac{g_j^*}{2n}, \end{aligned} \quad (28)$$

where the second inequality holds because the probability there are exactly i non zero-steps between two zero-steps is at most 2^{-i-1} , as the number of agents at state 0 is at least $n/2$ before $\varrho_j \leq \tau_2$; and

the last inequality above follows from \mathcal{E}_{j-1}^1 . The probability that $f_j > f_{j-1}$, i.e., at zero-step q_j the initiator moves to state \perp , is

$$\begin{aligned} \Pr[f_j > f_{j-1}] &= \frac{f_{j-1} + e_{j-1}/4}{n-1} + \frac{\mathbb{E}[e_j - e_{j-1}]/4}{n-1} \\ &\leq \frac{g_{j-1} - 3/4}{n-1} + \frac{g_j^*/8}{n-1}, \end{aligned}$$

where the last inequality follows from \mathcal{E}_{j-1}^1 and the bound on $\mathbb{E}[e_j - e_{j-1}]$ we computed above. Then,

$$p_j^0 \geq 1 - \Pr[e_j > e_{j-1}] - \Pr[f_j > f_{j-1}] \geq 1 - \frac{g_{j-1} - 3/4 + g_j^*}{n-1}.$$

Let $x := \frac{g_{j-1} - 3/4 + g_j^*}{n-1}$, and note that $0 < x < 1$, where the right inequality follows from \mathcal{E}_{j-1}^1 . To prove the claim we must show that $x \cdot (1 + p_j^1/p_j^0) - p_j^1/p_j^0 \in [0, 1]$. We showed above that $p_j^0 \geq 1 - x$, and we have $p_j^1 \leq 1 - p_j^0 \leq x$. Then

$$\begin{aligned} x \cdot (1 + p_j^1/p_j^0) - p_j^1/p_j^0 &= [xp_j^0 - (1-x)p_j^1]/p_j^0 \\ &\geq [x(1-x) - (1-x)x]/p_j^0 = 0. \end{aligned}$$

And

$$x \cdot (1 + p_j^1/p_j^0) - p_j^1/p_j^0 = x - (1-x)p_j^1/p_j^0 \leq x < 1. \quad \square$$

CLAIM 37.

$$\Pr\left[\bigcup_{1 \leq j \leq j_{\text{mid}}} (\mathcal{E}_{j-1}^1 \cap \{f_j + e_j \geq f_{j-1} + e_{j-1} + 2\})\right] = o(1/\log n).$$

PROOF. Let $1 \leq j \leq j_{\text{mid}}$. Suppose \mathcal{E}_{j-1}^1 occurs and $g_{j-1} \leq n^{1/3}$. Similarly to (28), we get

$$\begin{aligned} \Pr[e_j \geq e_{j-1} + 2] &\leq \sum_{i \geq 2} 2^{-i-1} \binom{i}{2} \left(\frac{d_{j-1}(d_{j-1} - 1)}{n(n-1)}\right)^2 \\ &\leq d_{j-1}^4/n^4 \leq (g_j^*/n)^2 = o(n^{-3/2}), \end{aligned}$$

as $j \leq j_{\text{mid}}$. Also,

$$\Pr[f_j > f_{j-1} \mid e_j = e_{j-1} + 1] = \frac{f_{j-1} + e_{j-1}/4 + 1/4}{n-1} \leq \frac{g_{j-1}}{n-1}.$$

From that and (28),

$$\Pr[\{f_j > f_{j-1}\} \cap \{e_j = e_{j-1} + 1\}] \leq \frac{g_{j-1}}{n-1} \cdot \frac{g_j^*}{2n} = o(n^{-3/2}),$$

as we have assumed that $g_{j-1} \leq n^{1/3}$. Combining the above, yields

$$\Pr[\mathcal{E}_{j-1}^1 \cap \{g_{j-1} \leq n^{1/3}\} \cap \{f_i + e_j \geq f_{j-1} + e_{j-1} + 2\}] = o(n^{-3/2}).$$

Taking the union over all j , gives that the probability of event

$$\mathcal{B} := \bigcup_{1 \leq j \leq j_{\text{mid}}} (\mathcal{E}_{j-1}^1 \cap \{f_i + e_j \geq f_{j-1} + e_{j-1} + 2\}) \cap \{g_{j_{\text{mid}}} \leq n^{1/3}\}$$

is $\Pr[\mathcal{B}] = o(j_{\text{mid}} n^{-3/2}) = o(1/\log n)$. Moreover,

$$\Pr[g_{j_{\text{mid}}} > n^{1/3}] \leq \Pr[\mathcal{C}_{g_{j_{\text{mid}}}^*, n^{1/3-1}, n-1} < j_{\text{mid}}] = o(1/\log n), \quad (29)$$

by Lemma 18(c). It follows that $\Pr[\mathcal{B} \cup \{g_{j_{\text{mid}}} > n^{1/3}\}] = o(1/\log n)$, which implies the claim. \square

$$\text{CLAIM 38. } \Pr[\bigcup_{1 \leq j \leq j_{\text{mid}}} \{d_{j-1}^2/n > g_j^*\}] = o(1/\log n).$$

PROOF. Since $t_2 > \tau_1$, we have $d_0 = (n/\ln n)^{1/2}$. Then, from Claim 32, for any $j \geq 0$,

$$\Pr[d_j^2/n > 4 \cdot e^{j/(2(n-1))} (\ln n)^{-1}] = O((n/\ln n)^{-1/2}).$$

Applying this for all $j \leq j_{\text{mid}}$ that are multiples of n , taking the union bound over those $\Theta((\ln \ln n)^2)$ different j , and using the fact that $d_{j'}^2/n \leq d_j^2/n$ if $j' < j$, we obtain

$$\Pr\left[\bigcup_{0 \leq j < j_{\text{mid}}} \{d_j^2/n > 8 \cdot e^{j/(2(n-1))} (\ln n)^{-1}\}\right] = O\left(\frac{(\ln \ln n)^2}{(n/\ln n)^{1/2}}\right).$$

The claim then follows. \square

We now prove the main property of the coupling.

CLAIM 39. Under the coupling described above,

$$\Pr\left[\bigcap_{0 \leq j \leq j_{\text{mid}}} (\{g_j \geq f_j + e_j/4 + 3/4\} \cup \{q_j > \tau_2\})\right] = 1 - o(1/\log n).$$

PROOF. We define the events

$$\begin{aligned} \mathcal{D}_1 &:= \bigcap_{1 \leq j \leq j_{\text{mid}}} (\bar{\mathcal{E}}_{j-1}^1 \cup \{f_j + e_j \leq f_{j-1} + e_{j-1} + 1\}) \\ \mathcal{D}_2 &:= \bigcap_{1 \leq j \leq j_{\text{mid}}} \{d_{j-1}^2/n \leq g_j^*\}, \quad \mathcal{D}_3 := \{g_{j_{\text{mid}}} \leq n^{1/3}\}, \end{aligned}$$

where $\bar{\mathcal{E}}_{j-1}^1$ is the complement of \mathcal{E}_{j-1}^1 . Let $\mathcal{D} := \mathcal{D}_1 \cap \mathcal{D}_2 \cap \mathcal{D}_3$. From Claims 37 and 38, and equation (29), it follows

$$\Pr[\mathcal{D}] = 1 - o(1/\log n).$$

We prove below that if \mathcal{D} holds then, for every $0 \leq j \leq j_{\text{mid}}$, $g_j \geq f_j + e_j/4 + 3/4$ or $q_j > \tau_2$. The claim then follows.

The proof is by induction on j . We have

$$g_0 = 1 > 3/4 = f_0 + e_0/4 + 3/4,$$

since $t_2 > \tau_1$ and thus $f_0 = e_0 = 0$; hence, the base case holds. Suppose now that $g_{j-1} \geq f_{j-1} + e_{j-1}/4 + 3/4$ or $q_{j-1} > \tau_2$, for some $1 \leq j \leq j_{\text{mid}}$. We must prove that if $q_j \leq \tau_2$, then $g_j \geq f_j + e_j/4 + 3/4$. From $q_j \leq \tau_2$, it follows $q_{j-1} < \tau_2$. From that and the induction hypothesis, we get $g_{j-1} \geq f_{j-1} + e_{j-1}/4 + 3/4$. Also, from \mathcal{D}_2 , $d_{j-1}^2/n \leq g_j^*/2$, and from \mathcal{D}_3 , $g_{j-1} - 3/4 + g_j^* < g_{j_{\text{mid}}} - 3/4 + g_{j_{\text{mid}}}^* < n-1$. Therefore, all four events whose intersection defines \mathcal{E}_{j-1}^1 hold, and thus \mathcal{E}_{j-1}^1 holds. Since \mathcal{E}_{j-1}^1 holds, the coupling stipulates that one of the cases (i)–(iii) of the coupling applies. Moreover, from \mathcal{D}_1 , we have $f_j + e_j \leq f_{j-1} + e_{j-1} + 1$, thus case (iii) does not apply. In both cases (i) and (ii), we have that $g_j - g_{j-1} \geq (f_j + e_j) - (f_{j-1} + e_{j-1})$. From that and fact $g_{j-1} \geq f_{j-1} + e_{j-1}/4 + 3/4$ (from \mathcal{E}_{j-1}^1), it follows

$$g_j \geq f_j + e_j/4 + 3/4.$$

This completes the induction proof. \square

Coupling for $j \geq j_{\text{mid}}$. The coupling is as follows. For $j \geq 0$, we define the next event under the coupling,

$$\begin{aligned} \mathcal{E}_j^2 &:= \{q_j < \tau_2\} \cap \{g_j - 3/4 + g_{j+1}^* < n-1\} \\ &\quad \cap \{f_j \leq g_j\} \cap \{e_{j+1}/4 \leq g_{j+1}^* - 3/4\}. \end{aligned}$$

For each $j > j_{\text{mid}}$, if $q_{j-1} < \tau_2$, then we first execute the steps $q_{j-1} < t \leq q_j$ of the protocol. After that, if event \mathcal{E}_{j-1}^2 holds, we select g_j according to one of the cases below.

- (i) If $f_j = f_{j-1} + 1$, then we set $g_j = g_{j-1} + 1$.
 (ii) If $f_j = f_{j-1}$, then we set $g_j = g_{j-1} + 1$ with probability

$$p_j := \left(\frac{g_{j-1} - 3/4 + g_j^*}{n-1} - \frac{f_{j-1} + e_j/4}{n-1} \right) \cdot \left(1 - \frac{f_{j-1} + e_j/4}{n-1} \right)^{-1};$$

and set $g_j = g_{j-1}$ with the remaining probability, $1 - p_j$.

If \mathcal{E}_{j-1}^2 does not hold, we choose g_j 's value independently of f_j .

It is not hard to see that the marginal distribution of g in the coupling above is the correct one. In particular, if \mathcal{E}_{j-1}^2 holds then the probability of $f_j = f_j + 1$ is

$$\frac{f_{j-1} + e_j/4}{n-1} \leq \frac{g_{j-1} + g_j^* - 3/4}{n-1} < 1,$$

where the two inequalities follow from \mathcal{E}_{j-1}^2 . Thus, if \mathcal{E}_{j-1}^2 holds, the probability of $g_j = g_j + 1$ is

$$\frac{f_{j-1} + e_j/4}{n-1} + \left(1 - \frac{f_{j-1} + e_j/4}{n-1} \right) \cdot p_j = \frac{g_{j-1} + g_j^* - 3/4}{n-1},$$

which is the correct probability.

CLAIM 40.

$$\Pr \left[\bigcup_{j > j_{\text{mid}}} (\{e_j/4 > g_j^* - 3/4\} \cap \{\varrho_j \leq \tau_2\}) \right] = o(1/\log n).$$

PROOF. From Corollary 34, for any $j > j_{\text{mid}}$,

$$\Pr[\{e_j > 2^7 \cdot e^{j/(2(n-1))} (\ln n)^{-1} + 1\} \cap \{\varrho_j \leq \tau_2\}] = O((n/\log n)^{-1/2}).$$

Applying this for all $j_{\text{mid}} < j \leq n \ln^2 n$ that are multiples of n , taking the union bound over those $\Theta(\ln^2 n)$ different j , and using the fact that $e_{j'} \leq e_j$ if $j' < j$, we obtain

$$\Pr \left[\bigcup_{j_{\text{mid}} < j \leq n \ln^2 n} \{e_j > 2^8 \cdot e^{j/(2(n-1))} (\ln n)^{-1} - 3\} \cap \{\varrho_j \leq \tau_2\} \right] = o(1/\log n).$$

Since $\Pr[\varrho_{n \ln^2 n} < \infty] = o(1/n)$, we can extend the range of the union above to $j_{\text{mid}} < j < \infty$, without increasing the right side. The claim then follows by applying the definition of g_j^* . \square

We now prove the main property of the coupling.

CLAIM 41. *Under the coupling described above,*

$$\Pr \left[\bigcap_{j \geq j_{\text{mid}}} (\{g_j \geq f_j\} \cup \{\varrho_j > \tau_2\}) \cup \{g_{j_{\text{mid}}} < f_{j_{\text{mid}}}\} \right] = 1 - o(1/\log n).$$

PROOF. We define the events $\mathcal{D}_1 := \{g_{j_{\text{mid}}} \geq f_{j_{\text{mid}}}\}$, and

$$\mathcal{D}_2 := \bigcap_{j > j_{\text{mid}}} (\{e_j/4 \leq g_j^* - 3/4\} \cup \{\varrho_j > \tau_2\}).$$

We show by induction that if $\mathcal{D}_1 \cap \mathcal{D}_2$ holds, then for every $j \geq j_{\text{mid}}$, $g_j \geq f_j$ or $\varrho_j > \tau_2$. The base case holds because of \mathcal{D}_1 . Suppose now that $g_{j-1} \geq f_{j-1}$ or $\varrho_{j-1} > \tau_2$, for some $j > j_{\text{mid}}$. We must prove that if $\varrho_j \leq \tau_2$ then $g_j \geq f_j$. From $\varrho_j \leq \tau_2$, it follows $\varrho_{j-1} < \tau_2$. From that and the induction hypothesis, we get $g_{j-1} \geq f_{j-1}$. Also, from \mathcal{D}_2 , $e_j/4 \leq g_j^* - 3/4$. Therefore, three of the four events whose intersection defines \mathcal{E}_{j-1}^2 hold. We consider two cases depending on whether the fourth event, $g_{j-1} - 3/4 + g_j^* < n - 1$ holds.

If $g_{j-1} - 3/4 + g_j^* < n - 1$, then \mathcal{E}_{j-1}^2 holds. The coupling then stipulates that one of the cases (i) or (ii) applies. In both cases, we have $g_j - g_{j-1} \geq f_j - f_{j-1}$. From that and the induction hypothesis $g_{j-1} \geq f_{j-1}$, we get $g_j \geq f_j$.

If $g_{j-1} - 3/4 + g_j^* \geq n - 1$, then $g_j = g_{j-1} + 1$. From that and the induction hypothesis $g_{j-1} \geq f_{j-1}$, we get $g_j \geq f_j$.

This completes the induction proof. We have thus shown that $\mathcal{D}_1 \cap \mathcal{D}_2 \subseteq \bigcap_{j \geq j_{\text{mid}}} (\{g_j \geq f_j\} \cup \{\varrho_j > \tau_2\})$. The claim now follows by taking the union of each side with event \mathcal{D}_1 , and using that

$$\Pr[(\mathcal{D}_1 \cap \mathcal{D}_2) \cup \mathcal{D}_1] \geq \Pr[\mathcal{D}_2] = 1 - o(1/\log n),$$

by Claim 40. \square

From Claims 39 and 41, and a union bound, we obtain the statement of Lemma 35. \square

Next we define a second random sequence, $h = (h_0, h_1 \dots)$, which is easier to bound, and compare it with g . We have $h_0 := 1$, and for $j > 0$, either $h_j = h_{j-1}$ or $h_j = h_{j-1} + 1$, such that¹¹

$$\Pr[h_j = h_{j-1} + 1 \mid h_{j-1}, \dots, h_1] = \frac{h_{j-1}}{n'},$$

where

$$n' := (n-1) \cdot \left(1 + \frac{3 \ln \ln n}{4 \ln n} \right)^{-1}.$$

LEMMA 42. *There is coupling of the sequences h and g such that $\Pr[\bigcap_{j \geq 0} \{h_j \geq g_j\}] = 1 - O(1/\log n)$.*

PROOF. For $j \geq 1$, let

$$g_j^+ := c_g (\ln n)^{-1} e^{\frac{j}{2(n-1)}} - 3/4.$$

Then, $\Pr[g_j = g_{j-1} + 1 \mid g_{j-1}] = (g_{j-1} + g_j^+)/ (n-1)$. Let also

$$h_j^+ := h_{j-1} \cdot (3 \ln \ln n) / (4 \ln n).$$

Then, $\Pr[h_j = h_{j-1} + 1 \mid h_{j-1}] = (h_{j-1} + h_j^+)/ (n-1)$.

We describe now the coupling of h and g . For each $j \geq 1$:

- (i) If $h_{j-1} \geq g_{j-1}$, $h_j^+ \geq g_j^+$, and $h_{j-1} + h_j^+ < n - 1$, then we first choose the value of h_j according to the law of h , i.e., $h_j = h_{j-1} + 1$ with probability $(h_{j-1} + h_j^+)/ (n-1)$, and $h_j = h_{j-1}$ with the remaining probability. After that, if $h_j = h_{j-1} + 1$, we set $g_j = g_{j-1} + 1$ with probability $(g_{j-1} + g_j^+)/ (h_{j-1} + h_j^+)$; with the remaining probability, or if $h_j = h_{j-1}$, we let $g_j = g_{j-1}$.
 (ii) If any of the three conditions in (i) are not met, then we choose the values of g_j and h_j independently, from their respective distributions.

Clearly the marginal distributions of h and g in the coupling above are the correct ones. In particular, in case (i), the probability that $h_j = h_{j-1} + 1$ is $[(h_{j-1} + h_j^+)/ (n-1)] \cdot [(g_{j-1} + g_j^+)/ (h_{j-1} + h_j^+)] = (g_{j-1} + g_j^+)/ (n-1)$.

Let

$$\zeta := \min \{j : h_{j-1} + h_j^+ \geq n - 1\}.$$

We observe that if $h_j^+ \geq g_j^+$ for all $0 < j < \zeta$, then $h_j \geq g_j$ for all $j \geq 0$. The reason is that $h_0 = g_0 = 1$, and for all $0 < j < \zeta$ a simple inductive argument shows that case (i) above applies, while

¹¹As before, if the right side exceeds 1 we implicitly assume it is 1.

for $j \geq \zeta$, h_j increases by one in each step, and thus $h_{\zeta-1} \geq g_{\zeta-1}$ implies $h_j \geq g_j$. Therefore, to prove the lemma it suffices to show

$$\Pr\left[\bigcap_{0 < j < \zeta} \{h_j^+ \geq g_j^+\}\right] = 1 - O(1/\log n). \quad (30)$$

The remaining of the proof is devoted to the proof of (30).

For $k \geq 0$, let

$$j_k := 2(n-1) \ln \ln n + 2(n-1)k \ln \ln \ln n - 2(n-1) \ln(4c_g/3).$$

Note that

$$g_{j_k}^+ = (3/4) \cdot (\ln \ln n)^k - 3/4.$$

CLAIM 43. Let $m := n' \cdot (3 \ln \ln n)/(4 \ln n)$. For all $k \geq 0$,

$$\Pr[h_{j_k}^+ < \min\{g_{j_{k+1}}^+, m\}] = O((\ln n)^{-1} \cdot (\ln \ln n)^{-k}).$$

PROOF. Substituting the definition of h_j^+ , and the values of $g_{j_{k+1}}^+$ and m , gives

$$\begin{aligned} \Pr[h_{j_k}^+ < \min\{g_{j_{k+1}}^+, m\}] \\ = \Pr[h_{j_k-1} < \min\{(\ln \ln n)^k \ln n - \ln n, n'\}]. \end{aligned}$$

The right side is at most

$$\Pr[C_{0, \min\{(\ln \ln n)^k \ln n - \ln n, n'\}, n'} > j_k - 1].$$

We bound this probability by applying Lemma 18(b): For $x = \ln \ln n + k \ln \ln \ln n - 2 \ln(4c_g/3)$,

$$\begin{aligned} n' \ln((\ln \ln n)^k \ln n - \ln n) + xn' < \\ (n-1) \ln((\ln \ln n)^k \ln n) + x(n-1) = j_k - 1. \end{aligned}$$

Then, from Lemma 18(b), the probability above is at most

$$e^{-x} = O((\ln n)^{-1} \cdot (\ln \ln n)^{-k}). \quad \square$$

From Claim 43 and a union bound over all k , we obtain

$$\Pr\left[\bigcap_{k \geq 0} \{h_{j_k}^+ \geq \min\{g_{j_{k+1}}^+, m\}\}\right] = 1 - O(1/\log n).$$

We observe that if $h_{j_k}^+ \geq \min\{g_{j_{k+1}}^+, m\}$ for some $k \geq 0$, then for any $j_k \leq j \leq j_{k+1}$, $h_j^+ \geq h_{j_k}^+ \geq \min\{g_{j_{k+1}}^+, m\} \geq \min\{g_j^+, m\}$. Also, for $0 < j < j_0$, we have $h_j^+ > 0 > g_j^+$. It follows

$$\bigcap_{k \geq 0} \{h_{j_k}^+ \geq \min\{g_{j_{k+1}}^+, m\}\} \subseteq \bigcap_{j > 0} \{h_j^+ \geq \min\{g_j^+, m\}\}.$$

Moreover, the event on the right is a subset of $\bigcap_{0 < j < \zeta} \{h_j^+ \geq g_j^+\}$, because if $j < \zeta$ then $h_j^+ + h_{j-1} < n-1$, which implies $h_j^+ < m$. It follows that the probability of event $\bigcap_{0 < j < \zeta} \{h_j^+ \geq g_j^+\}$ is at least $1 - O(1/\log n)$. This proves (30), and completes the proof of Lemma 42. \square

Next we compute a bound for the sequence h . Recall that $\sigma = n \ln n - n \ln \ln n + n \ln \ln \ln n$.

CLAIM 44. $\Pr[h_\sigma > n/4] = O(1/\log n)$.

PROOF. We use Lemma 18(d). We have

$$\Pr[h_\sigma > n/4] \leq \Pr[C_{0, n/4, n'} \leq \sigma],$$

and

$$\begin{aligned} (n'-1) \ln(n/4) - n' \ln \ln n = \\ n \ln n - (3/4)n \ln \ln n - n \ln \ln \ln n - O(n) > \sigma. \end{aligned}$$

Substituting this upper bound of σ to the right side of the previous equation, and applying Lemma 18(d), yields

$$\Pr[h_\sigma > n/4] \leq e^{-e^{\ln \ln n}} = 1/\ln n. \quad \square$$

The next statement follows from Lemmas 35 and 42 and Claim 44.

CLAIM 45. If $t_2 > \tau_1$ then

$$\Pr[\{f_\sigma \leq n/4\} \cup \{d_\sigma \geq n/4\}] = 1 - O(1/\log n).$$

PROOF. Note that step τ_2 is a zero-step, and define j to be the index of that zero-step, i.e., $\varrho_j = \tau_2$. From Lemma 35,

$$\Pr\left[\bigcap_{0 \leq j \leq j} \{g_j \geq f_j\}\right] = 1 - O(1/\log n).$$

From that, Lemma 42, and Claim 44,

$$\Pr\left[\bigcap_{0 \leq j \leq j} \{h_j \geq f_j\} \cap \{h_\sigma \leq n/4\}\right] = 1 - O(1/\log n).$$

If the even above occurs, i.e., $h_j \geq f_j$ for all $j \leq j$, and $h_\sigma \leq n/4$, then: (i) if $j > \sigma$, we have $f_\sigma \leq h_\sigma \leq n/4$; and (ii) if $j \leq \sigma$, we have $f_j \leq f_\sigma \leq h_\sigma \leq n/4$, and since $d_j + f_j \geq n/2$, we have $d_j \geq n/4$ and thus $d_\sigma \geq n/4$. Therefore, the probability of event $\{f_\sigma \leq n/4\} \cup \{d_\sigma \geq n/4\}$ is at least as large as the event's above. \square

We can now conclude the proof of the lower bound of Lemma 6(b) as follows. If $t_2 > \tau_1$ then from Claims 31 and 45, the event $\mathcal{E} := (\{d_\sigma \geq \lambda\} \cup \{j^* < \sigma\}) \cap (\{f_\sigma \leq n/4\} \cup \{d_\sigma \geq n/4\})$ has probability

$$\Pr[\mathcal{E}] = 1 - O(1/\log n).$$

We argue that $\mathcal{E} \subseteq \{d_\sigma \geq \lambda\}$: If $d_\sigma < \lambda$ then for \mathcal{E} to hold it must be the case that both $j^* < \sigma$ and $f_\sigma \leq n/4$. However, $j^* < \sigma$ implies $d_\sigma + f_\sigma = n$, thus $d_\sigma = n - f_\sigma > n - \lambda > n/2$, which contradicts $f_\sigma \leq n/4$. Therefore, if $t_2 > \tau_1$,

$$\Pr[d_\sigma \geq \lambda] \geq \Pr[\mathcal{E}] = 1 - O(1/\log n).$$

Finally,

$$\Pr[t_2 > \tau_1] = \Pr[n_{t_2} > (n/\ln n)^{1/2}] = 1 - O(1/\log n),$$

from Claim 30. Combining the last two results above, gives

$$\Pr[d_\sigma \geq \lambda] = 1 - O(1/\log n).$$

Since the total number of agents in states 1 and 2 eventually is $d_{j^*} \geq d_\sigma$, the claim follows.

E.4 Proof of Lemma 6(c)

Recall, $t_1 := \min \{t : n_t(0) < n\}$ and $t_4 := \min \{t : n_t(0) = 0\}$, and let $h := \min \{t : n_t(0) \leq n/2\}$. When an agent u in state 0 interacts with an agent in a non-0 state, u 's state changes to non-0 with probability at least $1/4$. Also, no transition is possible from a non-0 state to 0. Therefore, for any $t > t_1$ and $0 < k < n$,

$$\Pr[n_t(0) = k - 1 \mid n_{t-1}(0) = k] \geq \frac{k(n-k)}{4n(n-1)} \geq \frac{\min\{k, n-k\}}{8n}.$$

It follows that $h - t_1$ and $t_4 - h$ are both dominated by $C_{0, n/2, 8n}$, thus, from Lemma 18(b), they are both at most $O(n \log n)$, w.h.p. Then, by union bound, $t_4 - t_1 = O(n \log n)$ w.h.p.

F ANALYSIS OF SRE

F.1 Proof of Lemma 7(a)

The proof is very similar to that of Lemma 6(a). Suppose, for contradiction, that all agents are eliminated in SRE. An agent gets eliminated in SRE, i.e., reaches state \perp , if at some step it is in a state $s \notin \{z, \perp\}$ and interacts with an agent in state z or \perp . Let t be the step when the first agent reaches state \perp . Then the responder agent v in that step is either in state z or in state \perp right before step t . In both cases we have a contradiction: If v is in state \perp before t , this contradicts the definition of t ; if it is in state z then v is never eliminated in SRE.

F.2 Proof of Lemma 7(b)

Let X_t denote the set of agents that are in state x right after step t , and let $x_t := |X_t|$. Define similarly Y_t, y_t, Z_t, z_t , and B_t, b_t , for states y, z , and \perp , respectively. Let

$$\begin{aligned} t_1 &:= \min \{t : y_t \geq n^{1/2}\} \cup \{t_2\} \\ t_2 &:= \min \{t : z_t > 0\} \\ t_3 &:= \min \{t : z_t + b_t = n\} \\ t_4 &:= t_1 + 9n \ln n. \end{aligned}$$

Note that t_3 is the completion time of SRE.

CLAIM 46. $\Pr[t_3 > t_4] = O(1/n)$.

PROOF. Observe that $t_3 - t_2$ is upper bounded by the completion time of a one-way epidemic originated at the first agent that reaches state z . Thus, Lemma 20 gives

$$\Pr[t_3 - t_2 > 8n \ln n] = O(1/n). \quad (31)$$

If $z_{t_1} > 0$ then $t_2 - t_1 = 0$. If $z_{t_1} = 0$ then $y_{t_1} \geq n^{1/2}$ and $t_2 - t_1$ is upper bounded by the number of steps $t > t_1$, until two agents from Y_{t_1} interact. It follows

$$\Pr[t_2 - t_1 > n \ln n] \leq \left(1 - \frac{n^{1/2}(n^{1/2} - 1)}{n(n-1)}\right)^{n \ln n} = O(1/n),$$

where the last equation is obtained using the fact $(1 - \epsilon) \leq e^{-\epsilon}$. By a union bound, $\Pr[t_3 - t_1 > 9n \ln n] = O(1/n)$, and substituting $t_1 + 9n \ln n = t_4$ yields the claim. \square

Let S be the set of agents selected in DES. Since we have assumed DES is completed before f_2 , set S is finalized before f_2 . Let $k := |S|$ and suppose that $k = O(n^{3/4} \log n)$. Also, $k \geq 1$, from Lemma 6(a).

Let $Y_t^+ := \bigcup_{t' \leq t} Y_{t'}$ be the set of agents that reach state y before or at step t , and let $y_t^+ := |Y_t^+|$.

CLAIM 47. *There is a constant $c > 0$ such that for $\hat{y} := cn^{1/2} \log^3 n$, $\Pr[y_{t_4}^+ > \hat{y}] = o(1/n)$,*

PROOF. The difference $y_{t_4}^+ - y_{t_1}^+$ is upper bounded by the number of interactions between agents in S , during the steps $t_1 < t \leq t_4$. In each such step, the probability q that both agents that interact belong to S is independent for different steps and

$$q = \frac{k(k-1)}{n(n-1)} = \frac{O(n^{3/4} \log n)^2}{n(n-1)} = O(n^{-1/2} \log^2 n).$$

Then,

$$\mathbb{E}[y_{t_4}^+ - y_{t_1}^+] \leq q 9n \ln n \leq c_1 n^{1/2} \log^3 n,$$

for some constant c_1 . Chernoff bound (12) then gives

$$\Pr[y_{t_4}^+ - y_{t_1}^+ > 2c_1 n^{1/2} \log^3 n] = o(1/n).$$

Since $y_{t_1}^+ \leq n^{1/2} + 1$, the claim follows. \square

CLAIM 48. $\Pr[\{z_{t_4} \leq c' \log^7 n\} \cup \{y_{t_4}^+ > \hat{y}\}] = 1 - o(1/n)$, for some constant $c' > 0$.

PROOF. Let a_t , for $t_1 < t \leq t_4$, be 1 if two agents from set Y_{t-1} interact in step t and $y_{t-1}^+ \leq \hat{y}$; let $a_t = 0$ otherwise. Observe that $z_{t_4} - 1$ is bounded by $A := \sum_{t_1 < t \leq t_4} a_t$, if $y_{t_4}^+ \leq \hat{y}$. Thus, for any ℓ ,

$$\{A \leq \ell\} \subseteq \{z_{t_4} - 1 \leq \ell\} \cup \{y_{t_4}^+ > \hat{y}\}. \quad (32)$$

For each $t_1 < t \leq t_4$,

$$\Pr[a_t = 1] \leq \frac{\hat{y}(\hat{y} - 1)}{n(n-1)} \leq \frac{(cn^{1/2} \log^3 n)^2}{n^2} = c^2 n^{-1} \log^6 n.$$

Then,

$$\mathbb{E}[A] \leq 9n \ln n \cdot c^2 n^{-1} \log^6 n \leq c_2 \log^7 n,$$

for some constant c_2 . The bound on $\Pr[a_t = 1]$ above is independent of all $a_{t'}, t' \neq t$, thus we can apply Chernoff bound (12) to obtain

$$\Pr[A > 2c_2 \log^7 n] = o(1/n).$$

From that and (32), we get $\Pr[\{z_{t_4} - 1 \leq 2c_2 \log^7 n\} \cup \{y_{t_4}^+ > \hat{y}\}] \geq \Pr[A \leq 2c_2 \log^7 n] = 1 - o(1/n)$. \square

By union bound, we have $\Pr[\{z_{t_4} \leq c' \log^7 n\} \cup \{y_{t_4}^+ > \hat{y}\}] \leq \Pr[z_{t_4} \leq c' \log^7 n] + \Pr[y_{t_4}^+ > \hat{y}] = \Pr[z_{t_4} \leq c' \log^7 n] + o(1/n)$, by Claim 47. Substituting that to the equation of Claim 48 yields

$$\Pr[z_{t_4} \leq c' \log^7 n] = 1 - o(1/n).$$

From that and Claim 46,

$$\Pr[z_{t_3} \leq c' \log^7 n] \geq \Pr[\{z_{t_4} \leq c' \log^7 n\} \cap \{t_3 \leq t_4\}] = 1 - O(1/n).$$

This completes the proof of Lemma 7(b).

F.3 Proof of Lemma 7(c)

We will use the steps t_2, t_3 defined in part (b). We also define

$$t'_1 := \min \left\{ t: y_t \geq \lambda^{1/2} \right\} \cup \{t_2\}, \quad \lambda := \frac{n \ln \ln n}{\ln n}.$$

The next result is similar to Claim 46.

CLAIM 49. $\Pr[t_3 > t'_1 + 9n \ln n] = O(1/\log n)$.

PROOF. If $z_{t'_1} > 0$ then $t_2 - t'_1 = 0$. If $z_{t'_1} = 0$ then $y_{t'_1} \geq \lambda^{1/2}$ and $t_2 - t'_1$ is upper bounded by the number of steps $t > t'_1$, until two agents from $Y_{t'_1}$ interact. It follows

$$\Pr[t_2 - t'_1 > n \ln n] \leq \left(1 - \frac{\lambda^{1/2}(\lambda^{1/2} - 1)}{n(n-1)} \right)^{n \ln n} = O(1/\ln n).$$

Combining that and (31), yields the claim. \square

It remains to bound $t'_1 - l_2$. Let S be the set of agents selected in DES. Since DES is completed before f_2 , S is finalized before f_2 . Let $k := |S|$ and suppose that $k = \Omega(n^{3/4}(\log \log n)^{1/4}(\log n)^{-3/4})$.

CLAIM 50. $\Pr[t'_1 > l_2 + \hat{c}n \ln n] = O(1/n)$, for a constant $\hat{c} > 0$.

PROOF. Let $t_5 := l_2 + \hat{c}n \log n$, where constant \hat{c} is fixed later. Let d_t , for $l_1 < t \leq t_5$, be 1 if two agents from set X_{t-1} interact in step t or if $x_{t-1} < k/2$; let $d_t = 0$ otherwise. Let $D := \sum_{l_1 < t \leq t_5} d_t$. We argue that if $D \geq \lambda^{1/2}$ then $t'_1 \leq t_5$: Suppose that $t'_1 > t_5$. Then $y_{t_5} < \lambda^{1/2}$ and for each $l_1 < t \leq t_5$, $x_{t-1} > k - \lambda^{1/2} > k/2$. Hence, for each $l_1 < t \leq t_5$, $d_t = 1$ only if two agents from set X_{t-1} interact in step t , i.e., if $y_t = y_{t-1} + 1$. Note also that $y_t \geq y_{t-1}$, for $t < t_2$, and thus for $t < t'_1$. Therefore, $D = y_{t_5} - y_{l_2} < \lambda^{1/2}$, which is a contradiction. We have thus shown that $D \geq \lambda^{1/2}$ implies $t'_1 \leq t_5$. It follows

$$\Pr[t'_1 \leq t_5] \geq \Pr[D \geq \lambda^{1/2}]. \quad (33)$$

For each $l_1 < t \leq t_5$,

$$\Pr[d_t = 1] \geq \frac{(k/2)((k/2) - 1)}{n(n-1)} \geq c_3 n^{-1/2} \frac{(\log \log n)^{1/2}}{(\log n)^{3/2}},$$

for a constant $c_3 > 0$, as $k = \Omega(n^{3/4}(\log \log n)^{1/4}/(\log n)^{3/4})$. Then,

$$\begin{aligned} \mathbb{E}[D] &\geq \hat{c}n \log n \cdot c_3 n^{-1/2} \frac{(\log \log n)^{1/2}}{(\log n)^{3/2}} \\ &\geq 2n^{1/2} \frac{(\log \log n)^{1/2}}{(\log n)^{1/2}} = 2\lambda^{1/2}, \end{aligned}$$

where the second inequality holds if we choose constant \hat{c} large enough. The bound on $\Pr[d_t = 1]$ above is independent of all $d_{t'}, t' \neq t$, thus we can apply Chernoff bound (13) to obtain

$$\Pr[D \geq \lambda^{1/2}] = 1 - o(1/n).$$

From that and (33), the claim follows. \square

Finally, combining Claims 49 and 50 using a union bound, completes the proof of Lemma 7(c).

G ANALYSIS OF LFE

G.1 Proof of Lemma 8(a)

The proof is by contradiction. Suppose that all agents are eventually eliminated in LFE. An agent can get eliminated in LFE only if it is eliminated in SRE or interacts with an agent in a state (\cdot, ℓ) , where $\ell > 0$. From Lemma 7(a), not all agents are eliminated in SRE, thus at least one agent reaches a state (\cdot, ℓ) , with $\ell > 0$. Let ℓ^* be the maximum level reached by any agent. Thus, $\ell^* \geq \ell$ for any state (\cdot, ℓ) reached by an agent. Then, $\ell^* > 0$. Let u be an agent that reaches level ℓ^* , i.e., reaches state (in, ℓ^*) or (toss, ℓ^*) . Then u was not eliminated when it reached internal phase 3. Since we have assumed that all agents are eliminated in LFE eventually, u must interact at some step with an agent in a state (\cdot, ℓ) , where $\ell > \ell^*$. However, this contradicts that $\ell^* \geq \ell$ for any state (\cdot, ℓ) reached by an agent.

G.2 Proof of Lemma 8(b)

Let S be the set of agents that survive SRE. Since we assume SRE is completed before step f_3 , the set S is finalized before f_3 . Suppose that $|S| = k$ (we have $k > 0$ by Lemma 7(a)). For any agent $u \in S$, u is not eliminated in LFE precisely if the largest level ℓ it reaches is greater or equal to the levels reached by the other $k-1$ agents in S . Hence, the probability that u is not eliminated in LFE is

$$\sum_{\ell=0}^{\mu-1} \left(2^{-\ell-1} (1 - 2^{-\ell-1})^{k-1} \right) + 2^{-\mu} = \sum_{\ell=1}^{\mu} \left(2^{-\ell} (1 - 2^{-\ell})^{k-1} \right) + 2^{-\mu}.$$

We have

$$\sum_{\log k \leq \ell \leq \mu} \left(2^{-\ell} (1 - 2^{-\ell})^{k-1} \right) \leq \sum_{\log k \leq \ell \leq \mu} 2^{-\ell} \leq 2/k,$$

and

$$\sum_{1 \leq \ell \leq \log k} \left(2^{-\ell} (1 - 2^{-\ell})^{k-1} \right) \leq \sum_{1 \leq \ell \leq \log k} \left(2^{-\ell} e^{-2^{-\ell}(k-1)} \right) = O(1/k).$$

Substituting these above, yields that the probability of u not being eliminated in LFE is $O(1/k)$, thus, the expected total number of agents that are not eliminated in LFE is $O(|S|/k) = O(1)$.

G.3 Proof of Lemma 8(c)

Observe that after step l_3 , there are no agents left in the initial state, $(\text{wait}, 0)$. Moreover, each agent in state (toss, \cdot) participates (as an initiator) in at most $\mu + 1 = O(\log \log n)$ interactions before it reaches state (in, \cdot) . By Chernoff bound (12), an agent that is in state (toss, \cdot) right after l_3 , will reach a state (in, \cdot) after $O(n \log n)$ steps w.h.p. And by a union bound, w.h.p. every agent is in some state (in, \cdot) or (out, \cdot) at the end of step $t_1 = l_3 + cn \log n$, for a large enough constant $c > 0$. From then on, the additional number of step until the max-level propagates to all agents is bounded by the completion time of a one-way epidemic, which is $O(n \log n)$ w.h.p. (Lemma 20). By a union bound, the completion time of LFE is $t_1 + O(n \log n) = l_3 + O(n \log n)$ w.h.p.

H ANALYSIS OF EE1

H.1 Proof of Lemma 9(a)

The proof is by contradiction. Suppose that all agents are eventually eliminated in EE1. An agent u can get eliminated in EE1 only if it is eliminated in LFE, or if it is in state $(\text{in}, 0, \rho)$ and interacts with an agent in a state $(\cdot, 1, \rho)$. From Lemma 8(a), not all agents are eliminated in LFE. Thus, there is a step t in which some agent in state $(\text{in}, 0, \rho)$ interacts with an agent in state $(\cdot, 1, \rho)$. Let ρ^* be the largest ρ for which such a step t exists. Then the responder agent in step t is in state $(\cdot, 1, \rho^*)$. Let w denote the first agent that reaches a state $(\cdot, 1, \rho^*)$; then the state reached by w must be $(\text{in}, 1, \rho^*)$. It follows that w is not eliminated in EE1 before reaching internal phase $\rho^* + 1$. Since w must be eliminated eventually, as we have assumed that all agents are, there is some $\rho' > \rho^*$ and a step t' , such that w is in state $(\text{in}, 0, \rho)$ right before t' , and at t' w interacts with an agent in state $(\cdot, 1, \rho)$. This, however, contradicts the definition of ρ^* .

H.2 Proof of Lemma 9(b)

We start by analyzing a simple game.

CLAIM 51. *Consider the following game. We start with $k_0 = k$ fair coins. In each round $r \geq 1$ of the game, we toss all coins that are still in the game. A coin is removed from the game in round r if its outcome is tails and the outcome of at least one other coin in the round is heads. Let k_r denote the number of coins that are still in the game after r rounds. Then $\mathbb{E}[k_r - 1] \leq (k - 1)/2^r$.*

PROOF. The proof is by induction on r . For $r = 0$, $k_r - 1 = k - 1 = (k - 1)/2^0$, thus the claim holds. Let $r \geq 0$ and suppose that $\mathbb{E}[k_r - 1] \leq (k - 1)/2^r$. If $k_r = s$, i.e., there are s coins still in the game at the beginning of round $r + 1$, the probability a given coin is not removed in this round is $1/2 + 1/2^s$. It follows

$$\mathbb{E}[k_{r+1} \mid k_r = s] = s(1/2 + 1/2^s).$$

Then

$$\mathbb{E}[k_{r+1} - 1 \mid k_r = s] = s(1/2 + 1/2^s) - 1 \leq (s - 1)/2.$$

It follows

$$\mathbb{E}[k_{r+1} - 1] \leq \mathbb{E}[k_r - 1]/2 \leq (k - 1)/2^{r+1},$$

by the induction hypothesis. This completes the inductive proof. \square

Let S be the set of agents that survive LFE. Since we assume LFE is completed before step f_4 , the set S is finalized before f_4 . Suppose that $|S| = k$. W.l.o.g., we assume that for every agent $u \in S$, we toss *in advance* μ coins, and we use the outcome of these coins to determine the outcome of any interaction $(\text{toss}, 0, \cdot) + (\cdot, \cdot, \cdot)$ in which u participates as an initiator; let R_u denote this sequence of coin tosses. Let E_ρ , for $4 \leq \rho \leq v - 2$ denote the event that right before step $f_{\rho+1}$, every agent is either in state (in, x, ρ) or (out, x, ρ) , where $x \in \{0, 1\}$ is the maximum value of any coin in phase ρ . Note that if E_ρ occurs, then the agents in state (in, x, ρ) right before $f_{\rho+1}$ are precisely those that survive phase ρ in EE1. Let $\mathcal{E}_\rho := \bigcap_{4 \leq i \leq \rho} E_i$. It is now easy to see that, given \mathcal{E}_ρ , s_ρ is equal to k_r , for $r = \rho - 3$, if the same coins are used in the game as those in the protocol, i.e., the sequence of outcomes of the i th coin,

among the k coins that start the game, is a prefix of the sequence R_u of coin values generated for the i th agent $u \in S$. It follows

$$\mathbb{E}[k_{\rho-3} - 1] \geq \mathbb{E}[(s_\rho - 1) \cdot \mathbb{1}_{\mathcal{E}_\rho}].$$

Using Claim 51 to bound the left side, yields

$$\mathbb{E}[(s_\rho - 1) \cdot \mathbb{1}_{\mathcal{E}_\rho}] \leq (k - 1)/2^{\rho-3}. \quad (34)$$

CLAIM 52. $\Pr[\neg \mathcal{W}_{4,\rho} \setminus \mathcal{E}_\rho] = O(\rho/n^{c_W+1})$.

PROOF. Let t'_ρ be the first step such that each agent participates as an initiator in at least one interaction in steps $l_\rho < t \leq t'_\rho$. Let also t_ρ be the first step that satisfies the following condition. Consider the schedule of interactions in steps $t'_\rho < t \leq t_\rho$, i.e., the sequence of the agent pairs that interact in those steps. The condition that t_ρ must satisfy is that a one-way epidemic originated at any agent is completed after at most $t_\rho - t'_\rho$ steps, if the schedule of interactions is the same as the one above. We have

$$\Pr[t'_\rho > l_\rho + (c_W + 2)n \ln n] \leq n(1 - 1/n)^{(c_W+2)n \ln n} \leq n^{-c_W-1}.$$

Also, by Lemma 20,

$$\Pr[t_\rho > t'_\rho + (4c_W + 8)n \ln n] \leq 2n^{-c_W-1}.$$

It follows

$$\Pr[t_\rho \geq l_\rho + (5c_W + 11)n \ln n] \leq 3n^{-c_W-1}.$$

Observe now that $\{t_\rho < f_{\rho+1}\} \subseteq E_\rho$: As soon as every agent has participated in at least one interaction after step l_ρ , no agent is still in state $(\text{toss}, 0, \rho)$; and once that happens, completing the one-way epidemic dissemination of the largest coin value before step $f_{\rho+1}$ ensures that every agent is in state (in, x, ρ) or (out, x, ρ) before $f_{\rho+1}$. It follows that $\{t_\rho < l_\rho + (5c_W + 11)n \ln n\} \subseteq \{t_\rho < f_{\rho+1}\} \cup \{f_{\rho+1} < l_\rho + (5c_W + 11)n \ln n\} \subseteq E_\rho \cup \bar{\mathcal{W}}_{\rho,\rho} \subseteq E_\rho \cup \bar{\mathcal{W}}_{4,\rho}$. Then

$$\begin{aligned} \Pr[E_\rho \cup \bar{\mathcal{W}}_{4,\rho}] &\geq \Pr[t_\rho < l_\rho + (5c_W + 11)n \ln n] \\ &\geq 1 - 3n^{-c_W-1}, \end{aligned}$$

and by union bound,

$$\Pr[\mathcal{E}_\rho \cup \bar{\mathcal{W}}_{4,\rho}] \geq 1 - 3(\rho - 3)n^{-c_W-1}.$$

This implies the claim. \square

We have

$$\begin{aligned} \mathbb{E}[(s_\rho - 1) \cdot \mathbb{1}_{\mathcal{W}_{4,\rho}}] &\leq \mathbb{E}[(s_\rho - 1) \cdot \mathbb{1}_{\mathcal{E}_\rho}] + \mathbb{E}[(s_\rho - 1) \mathbb{1}_{\mathcal{W}_{4,\rho} \setminus \mathcal{E}_\rho}] \\ &\leq (k - 1)/2^{\rho-3} + (k - 1) \cdot \Pr[\bar{\mathcal{W}}_{4,\rho} \setminus \mathcal{E}_\rho] \\ &\leq (k - 1)/2^{\rho-3} + O((k - 1)\rho/n^{c_W+1}) \\ &\leq k/2^{\rho-3}, \end{aligned}$$

where in the second line we used (34), in the third we used Claim 52, and the last line holds for all large enough n since $c_W \geq 1$.

I ANALYSIS OF EE2

I.1 Proof of Lemma 10(a)

We will use the next claim, which follows from a simple inductive argument.

CLAIM 53. *Suppose that $L_{\text{int}}(r) > 0$, for all $\rho_1 \leq r \leq \rho_2$. For any $l_{\rho_1} \leq t < f_{\rho_2}$, if two agents have the same value in their variable parity right after step t , then they are in the same internal phase.*

Suppose that $L_{\text{int}}(r) > 0$, for all $r \leq \rho + 1$. We must show that not all agents are eliminated in EE2 before reaching internal phase $\rho + 1$. The proof is similar to that of Lemma 9(a). Suppose for contradiction that every agent is eliminated in EE2 before reaching internal phase $\rho + 1$. An agent u can get eliminated in EE2 only if it is eliminated in EE1, or if it is in state $(\text{in}, 0, p)$ and interacts with an agent in a state $(\cdot, 1, p)$. From Lemma 9(a), not all agents are eliminated in EE1. Thus, there is a step t in which some agent in state $(\text{in}, 0, p)$ and internal phase r , interacts with an agent in state $(\cdot, 1, p)$. Let ρ^* be the largest r for which such a step t exists. Then $\rho^* \leq \rho$ since we have assumed that all agents are eliminated in EE2 before reaching internal phase $\rho + 1$. For the same reason, $t < l_{\rho+1}$, and thus $t < f_{\rho+2}$, because of the assumption that $L_{\text{int}}(\rho + 1) > 0$. Since the responder agent v in step t is in state $(\cdot, 1, p)$, its variable parity is the same as u 's before the step, and Claim 53 implies that v 's internal phase is ρ^* at that point. Let w denote the first agent that reaches state $(\cdot, 1, p)$ while being in internal phase ρ^* ; let t' be the step when that happens. Then the state of w right after step t' must be $(\text{in}, 1, p)$, otherwise w interacts at t' with an agent w' who is already in state $(\cdot, 1, p)$ and, by Claim 53, also in internal phase ρ^* , contrary to the definition of w . It follows that w is not eliminated in EE2 before reaching internal phase $\rho^* + 1$. Since w must be eliminated before reaching internal phase $\rho + 1$, as we have assumed that all agents are, there is some $\rho' > \rho^*$ and a step t'' , such that w is in state $(\text{in}, 0, p')$ and internal phase ρ' right before t'' , and at t'' w interacts with an agent in $(\cdot, 1, p')$. This, however, contradicts the definition of ρ^* .

I.2 Proof of Lemma 10(b)

The proof is similar to that of Lemma 9(b). Suppose that $L_{\text{int}}(v - 1) > 0$. Let S be the set of agents not eliminated in EE1. Since $L_{\text{int}}(v - 1) > 0$, S is finalized before step $f_v > l_{v-1}$. Suppose that $|S| = k$. W.l.o.g., we assume that for every agent $u \in S$, we toss in advance $c_W \log n$ coins, and we use the outcome of these coins to determine the outcome of any interaction $(\text{toss}, 0, \cdot) + (\cdot, \cdot, \cdot)$ in which u participates as an initiator; let R_u denote this sequence of coin tosses. Let E'_ρ be the event $E'_\rho := \bigcup_{\rho-1 \leq r \leq \rho+1} \{L_{\text{int}}(r) > 0\}$. Note that given E'_ρ , Claim 53 implies that if an agent u in internal phase ρ has the same parity value as another agent v , then v is also in internal phase ρ . Let E_ρ , for $v \leq \rho \leq c_W \log n$, be the event that, if E'_ρ occurs, then right before step $f_{\rho+1}$, every agent is either in state (in, x, p) or (out, x, p) , and in internal phase ρ , where $x \in \{0, 1\}$ is the maximum value of any coin in phase ρ , and $p = \rho \bmod 2$. Define $\mathcal{E}_\rho := \bigcap_{v \leq i \leq \rho} (E_i \cap E'_i)$. Given \mathcal{E}_ρ , s_ρ is equal to k_r , for $r = \rho - v + 1$, if the same coins are used in the game as those in the protocol, similarly to the proof of Lemma 9(b). Then

$$\mathbb{E}[k_{\rho-v+1} - 1] \geq \mathbb{E}[(s_\rho - 1) \cdot \mathbb{1}_{\mathcal{E}_\rho}].$$

and from Claim 51,

$$\mathbb{E}[(s_\rho - 1) \cdot \mathbb{1}_{\mathcal{E}_\rho}] \leq (k - 1)/2^{\rho-v+1}. \quad (35)$$

Also, similarly to Claim 52, we have:

$$\text{CLAIM 54. } \Pr[\mathcal{W}_{v-1, \rho+1} \setminus \mathcal{E}_\rho] = O(\rho/n^{c_W+1}).$$

PROOF. Let t'_ρ and t_ρ be defined as in the proof of Claim 52. We showed there that

$$\Pr[t_\rho \geq l_\rho + (5c_W + 1)n \ln n] \leq 3n^{-c_W-1}.$$

Also, similarly to the proof of Claim 52, $\{t_\rho < f_{\rho+1}\} \subseteq E_\rho$. It follows $\{t_\rho < l_\rho + (5c_W + 1)n \ln n\} \subseteq \{t_\rho < f_{\rho+1}\} \cup \{f_{\rho+1} < l_\rho + (5c_W + 1)n \ln n\} \subseteq E_\rho \cup \mathcal{W}_{v-1, \rho+1}$. Note also $E_\rho \cup \mathcal{W}_{v-1, \rho+1} = (E_\rho \cap E'_\rho) \cup \mathcal{W}_{v-1, \rho+1}$, as $\mathcal{W}_{v-1, \rho+1} \subseteq E'_\rho$. Then

$$\begin{aligned} \Pr[(E_\rho \cap E'_\rho) \cup \mathcal{W}_{v-1, \rho+1}] &\geq \Pr[t_\rho < l_\rho + (5c_W + 1)n \ln n] \\ &\geq 1 - 3n^{-c_W-1}, \end{aligned}$$

and by union bound,

$$\Pr[\mathcal{E}_\rho \cup \mathcal{W}_{v-1, \rho+1}] \geq 1 - 3(\rho - v + 1)n^{-c_W-1}.$$

This implies the claim. \square

We have

$$\begin{aligned} &\mathbb{E}[(s_\rho - 1) \cdot \mathbb{1}_{\mathcal{W}_{v-1, \rho+1}}] \\ &\leq \mathbb{E}[(s_\rho - 1) \cdot \mathbb{1}_{\mathcal{E}_\rho}] + \mathbb{E}[(s_\rho - 1) \mathbb{1}_{\mathcal{W}_{v-1, \rho+1} \setminus \mathcal{E}_\rho}] \\ &\leq (k - 1)/2^{\rho-v+1} + (k - 1) \cdot \Pr[\mathcal{W}_{v-1, \rho+1} \setminus \mathcal{E}_\rho] \\ &\leq (k - 1)/2^{\rho-v+1} + O((k - 1)\rho/n^{c_W+1}) \\ &\leq k/2^{\rho-v+1}, \end{aligned}$$

where in the third line we used (35), in the fourth we used Claim 54, and the last line holds for all large enough n , as $\rho \leq c_W \log n$.

J ANALYSIS OF SSE

J.1 Proof of Lemma 11(a)

It is immediate from the SSE protocol that $L_{t+1} \subseteq L_t$, for any $t \geq 0$, because there is no transition that changes the state of an agent from non-C to C, and the only possible transition to state S is from state C.

Let $S_t \subseteq L_t$ be the set of agents in state S right after step t .

CLAIM 55. *If $S_t \neq \emptyset$ then $S_{t'} \neq \emptyset$ for all $t' > t$.*

PROOF. Suppose for contradiction this is not true, and let $t' > t$ be the first step for which $S_{t'} = \emptyset$. Then in step t' , an agent's state changed from S to non-S. This is possible only if it interacted with another agent v in state S at step t' . But then v 's state is still S right after step t' , which contradicts $S_{t'} = \emptyset$. \square

We now show that $L_t \neq \emptyset$, for all $t \geq 0$. Suppose for contradiction this is not true, and let t' be the first step for which $L_{t'} = \emptyset$. Clearly $t' > 0$, as L_0 contains all agents. Also $S_t = \emptyset$ for all $0 \leq t \leq t'$, otherwise Claim 55 would imply that $S_{t'} \neq \emptyset$ and thus $L_{t'} \neq \emptyset$. It follows that no agent reached state S in any of the first t' steps, and thus no agent reached state F either. Since we have assumed that $L_{t'} = \emptyset$, and all agents are initially in state C, they must all have their state changed to E at some point during the first t' steps.

This means that all agents are eliminated in EE1, which contradicts Lemma 9(a). This completes the proof of Lemma 11(a).

J.2 Proof of Lemma 11(b)

CLAIM 56. *If $S_{l'_1} = \{u\}$ then $v \notin S_t$ for all $v \neq u$ and $l'_1 \leq t < f'_2$.*

PROOF. Suppose for contradiction that $v \in S_t$, for some v, t as above, and consider the smallest such t . Clearly $t > l'_1$, as we have assumed $S_{l'_1} = \{u\}$. Since $t < f'_2$, agent v can only move to state S at step t , if it is true at t that v is not eliminated in EE2 and $x_{\text{phase}} = 1$, and it is true right before step t that v is in state C. If this is the case, however, the same conditions are also true for the step $t' \leq l'_1$ at which v reached external phase 1. It follows that $v \in S_{t'}$, which contradicts that v is in state C right before step t , as an agent cannot return to state C once it leaves that state. \square

Suppose that $S_{l'_1} = \{u\}$. Let F_t be the set of agents in state F right after step t , and let $t_1 := \min\{t : |S_t \cup F_t| = n\}$. Then $t_1 - l'_1$ is upper bounded by the completion time of a one-way epidemic started at u , because when an agent u' interacts with an agent v who is in state S or F, agent u' is also in state S or F after the interaction. Since the completion time of a one-way epidemic is $O(n \log n)$ w.h.p.

(Lemma 20), we have w.h.p.

$$\min\{t : |S_t \cup F_t| = n\} \leq l'_1 + O(n \log n).$$

From Claims 55 and 56, it follows $|S_t| = 1$ for all $l'_1 \leq t < f'_2$. Thus $\min\{t : |S_t \cup F_t| = n\} \geq \min\{t : |F_t| = n - 1\} \cup \{f'_2\} \geq \min\{t : |L_t| = 1\} \cup \{f'_2\}$. Combining that with the bound above completes the proof of Lemma 11(b).

J.3 Proof of Lemma 11(c)

Suppose that $|L_{t_2}| = \kappa > 1$. Since $t_2 \geq l'_2$, no agent is in state C after t_2 , and thus no new agents reach state S at any later step. Therefore, $\min\{t : |L_t| = 1\} = \min\{t : |S_t| = 1\}$. To bound the right side, it suffices to consider just interactions in which both agents are in state S. Each such interaction reduces the size of S_t by one. If $|S_t| = k$, the expected number of steps before an interaction between agents in S occurs is $\frac{n(n-1)}{k(k-1)}$. It follows

$$\begin{aligned} \mathbf{E}[\min\{t : |L_t| = 1\}] &= \mathbf{E}[\min\{t : |S_t| = 1\}] \\ &= t_2 + \sum_{2 \leq k \leq \kappa} \frac{n(n-1)}{k(k-1)} \leq t_2 + \sum_{k \geq 2} \frac{n(n-1)}{k(k-1)} \\ &= t_2 + n(n-1). \end{aligned}$$