

A time and space optimal stable population protocol solving exact majority

David Doty¹, Mahsa Eftekhari¹, Leszek Gąsieniec², Eric Severson¹, Grzegorz Stachowiak³,
and Przemysław Uznański³

¹University of California, Davis, {doty,mhseftekhari,eseverson}@ucdavis.edu

²University of Liverpool, l.a.gasieniec@liverpool.ac.uk

³University of Wrocław, {gst,puznanski}@cs.uni.wroc.pl

Abstract

We study population protocols, a model of distributed computing appropriate for modeling well-mixed chemical reaction networks and other physical systems where agents exchange information in pairwise interactions, but have no control over their schedule of interaction partners. The well-studied *majority* problem is that of determining in an initial population of n agents, each with one of two opinions A or B , whether there are more A , more B , or a tie. A *stable* protocol solves this problem with probability 1 by eventually entering a configuration in which all agents agree on a correct consensus decision of A , B , or T , from which the consensus cannot change. We describe a protocol that solves this problem using $O(\log n)$ states ($\log \log n + O(1)$ bits of memory) and optimal expected time $O(\log n)$. The number of states $O(\log n)$ is known to be optimal for the class of polylogarithmic time stable protocols that are “output dominant” and “monotone” [5]. These are two natural constraints satisfied by our protocol, making it simultaneously time- and state-optimal for that class. We introduce a key technique called a “fixed resolution clock” to achieve partial synchronization.

Our protocol is *nonuniform*: the transition function has the value $\lceil \log n \rceil$ encoded in it. We show that the protocol can be modified to be uniform, while increasing the state complexity to $\Theta(\log n \log \log n)$.

Acknowledgements. Doty, Eftekhari, and Severson were supported by NSF award 1900931 and CAREER award 1844976.

1 Introduction

Population protocols [9] are asynchronous, complete networks that consist of computational entities called *agents* with no control over the schedule of interactions with other agents. In a population of n agents, repeatedly a random pair of agents is chosen to interact, each observing the state of the other agent before updating its own state.¹ They are an appropriate model for electronic computing scenarios such as sensor networks and for “fast-mixing” physical systems such as animal populations [48], gene regulatory networks [23], and chemical reactions [46], the latter increasingly regarded as an implementable “programming language” for molecular engineering, due to recent experimental breakthroughs in DNA nanotechnology [27, 47].

Time complexity in a population protocol is defined by *parallel time*: the total number of interactions divided by the population size n , henceforth called simply *time*. This captures the natural timescale in which each individual agent experiences expected $O(1)$ interactions per unit time. All problems solvable with zero error probability by a constant-state population protocol are solvable in $O(n)$ time [11, 31]. The benchmark for “efficient” computation is thus sublinear time, ideally $\text{polylog}(n)$, with $\Omega(\log n)$ time a lower bound on most nontrivial computation, since a simple coupon collector argument shows that is the time required for each agent to have at least one interaction.

As a simple example of time complexity, suppose we want to design a protocol to decide whether at least one x exists in an initial population of x ’s and q ’s. The single transition $x, q \rightarrow x, x$ indicates that if agents in states x and q interact, the q agent changes state to x . If x outputs “yes” and q outputs “no”, this takes expected time $O(\log n)$ to reach a consensus of all x ’s (i.e., $O(n \log n)$ total interactions, including null interactions between two x ’s or between two q ’s). However, the transitions $x, x \rightarrow y, y$; $y, x \rightarrow y, y$; $y, q \rightarrow y, y$, where x, q output “no” and y outputs “yes”, which computes whether at least *two* x ’s exist, is exponentially slower: expected time $O(n)$. The worst-case input is exactly 2 x ’s and $n - 2$ q ’s, where the first interaction between the x ’s takes expected $\binom{n}{2} = \Theta(n^2)$ interactions, i.e., $\Theta(n)$ time.

To have probability 0 of error, a protocol must eventually *stabilize*: reach a configuration where all agents agree on the correct output, which is *stable*, meaning no subsequently reachable configuration can change the output.² The original model [9] assumed states and transitions are constant with respect to n . However, for important problems such as leader election [32], majority computation [4], and computation of other functions and predicates [13], no constant-state protocol can stabilize in sublinear time with probability 1.³ This has motivated the study of population protocols whose number of states is allowed to grow with n , and as a result they can solve such problems in polylogarithmic time [4–6, 8, 14, 16–21, 24, 29, 30, 35, 37, 38].

1.1 The majority problem in population protocols

Angluin, Aspnes, and Eisenstat [12] showed a protocol they called *approximate majority*, which means that starting from an initial population of n agents with opinions A or B , if $|A - B| = \omega(\sqrt{n \log n})$ (i.e., the *gap* between the initial majority and minority counts is greater than roughly \sqrt{n}), then with high probability the algorithm stabilizes to all agents adopting the majority opinion in $O(\log n)$ time. A tighter analysis by Condon, Hajiaghayi, Kirkpatrick, and Mañuch [28] reduced the required gap to $\Omega(\sqrt{n \log n})$.

Mertzios, Nikolettseas, Raptopoulos, and Spirakis [41], and independently Draief and Vojnović [33], showed a 4-state protocol that solves *exact* majority problem, i.e., it identifies the majority correctly, no matter how small the initial gap.⁴ We henceforth refer to this simply as the *majority* problem. The protocol of [33, 41] is also *stable* in the sense that it has probability 1 of getting the correct answer. However, this protocol takes $\Omega(n)$ time in the worst case: when the gap is $O(1)$. Known work on the stable majority problem is summarized in Table 1. Gašieniec, Hamilton, Martin, Spirakis, and Stachowiak [36] investigated $\Omega(n)$ time

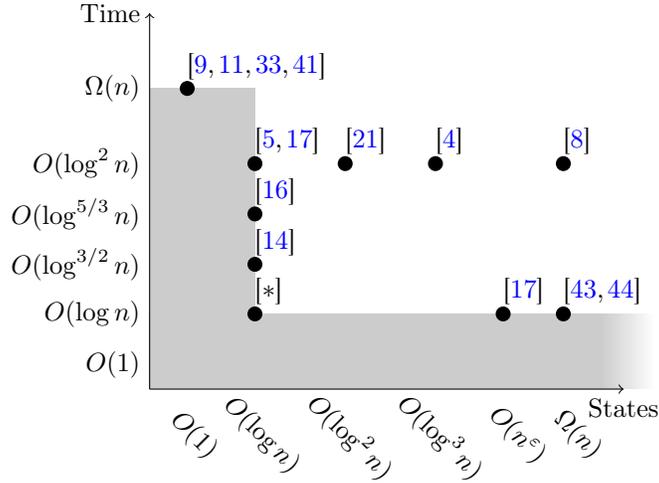
¹Using message-passing terminology, each agent sends its entire state of memory as the message.

²Technically this connection between probability 1 correctness and reachability requires the number of producible states for any fixed population size n to be finite, which is the case for our protocol.

³These problems have $O(1)$ state, sublinear time *converging* protocols [40]. A protocol *converges* when it reaches the correct output without subsequently changing it—though it may remain *changeable* for some time after converging—whereas it *stabilizes* when the output becomes unchangeable. See [17, 32] for a discussion of the distinction between stabilization and convergence. In this paper we consider only stabilization time.

⁴The 4-state protocol doesn’t identify ties, (gap = 0), but this can be handled with 2 more states; see [Stable Backup](#).

Table 1: Summary of results on the stable exact majority problem in population protocols, including this paper [*]. Gray regions are provably impossible: $o(\log \log n)$ state, $o(n)$ time unconditionally [4], $o(\log n)$ state, $O(n^{1-\varepsilon})$ time for monotone, output-dominant protocols [5], and $o(\log n)$ time unconditionally.



protocols for majority and the more general “plurality consensus” problem. Blondin, Esparza, Jaax, and Kučera [22] show a similar stable (also $\Omega(n)$ time) majority protocol that also reports if there is a tie.

Alistarh, Gelashvili, and Vojnović [8] showed the first stable majority protocol with worst-case poly-logarithmic expected time, requiring $\Omega(n)$ states. A series of positive results reduced the state and time complexity for stable majority protocols [4, 5, 14, 16, 17, 21, 43, 44]. Ben-Nun, Kopelowitz, Kraus, and Porat showed the current fastest stable sublinear-state protocol [14] using $O(\log^{3/2} n)$ time and $O(\log n)$ states. The current state-of-the-art protocols use alternating phases of *cancelling* (two biased agents with opposite opinions both become unbiased, preserving the difference between the majority and minority counts) and *splitting* (a.k.a. *doubling*): a biased agent converts an unbiased agent to its opinion; if all biased agents that didn’t cancel can successfully split in that phase, then the count difference doubles. The goal is to increase the count difference until it is n ; i.e., all agents have the majority opinion. See [7, 35] for relevant surveys.

Some non-stable protocols solve exact majority with high probability but have a small positive probability of incorrectness. Berenbrink, Elsässer, Friedetzky, Kaaser, Kling, and Radzik [17] showed a protocol that with initial gap α uses $O(s + \log \log n)$ states and WHP converges in $O(\log n \log_s(\frac{n}{\alpha}))$ time.⁵ Setting $\alpha = 1$ and $s = O(1)$, their protocol uses $O(\log \log n)$ states and converges in $O(\log^2 n)$ time. Kosowski and Uznański [40] showed a protocol using $O(1)$ states and converging in poly-logarithmic time with high probability.

On the negative side, Alistarh, Aspnes, and Gelashvili [5] showed that any stable majority protocol taking (roughly) less than linear time requires $\Omega(\log n)$ states if it also satisfies two conditions (satisfied by all known stable majority protocols, including ours): *monotonicity* and *output dominance*. These concepts are discussed in Section 9. In particular, the $\Omega(\log n)$ state bound of [5] applies only to *stable* (probability 1) protocols; the high probability protocol of [17], for example, uses $O(\log \log n)$ states and $O(\log^2 n)$ time.

1.2 Our contribution

We show a stable population protocol solving the exact majority problem in optimal $O(\log n)$ time (in expectation and with high probability) that uses $O(\log n)$ states. Our protocol is both monotone and output dominant (see Section 9 or [5] for discussion of these definitions), so by the $\Omega(\log n)$ state lower bound of [5], our protocol is both time and space optimal for the class of monotone, output-dominant stable protocols.

A high-level overview of the algorithm is given in Sections 3.1 and 3.2, with a full formal description given in Section 3.4. Like most known majority protocols using more than constant space (the only exceptions

⁵This protocol is SIMPLEMAJORITY in [17], which they then build on to achieve multiple stable protocols. The stable protocols require either $\Omega(n)$ stabilization time or $\Omega(\log n)$ states to achieve sublinear stabilization time.

being in [17]), our protocol is *nonuniform*: agents have an estimate of the value $\lceil \log n \rceil$ embedded in the transition function and state space. Section 8 describes how to modify our main protocol to make it uniform, retaining the $O(\log n)$ time bound, but increasing the state complexity to $O(\log n \log \log n)$ in expectation and with high probability. That section discusses challenges in creating a uniform $O(\log n)$ state protocol.

2 Preliminaries

We write $\log n$ to denote $\log_2 n$, and $\ln n$ to denote the natural logarithm. We write $x \sim y$ to denote that x and y are asymptotically equivalent (implicitly in the population size n), meaning $\lim_{n \rightarrow \infty} \frac{x(n)}{y(n)} = 1$.

2.1 Population protocols

A *population protocol* is a pair $\mathcal{P} = (\Lambda, \delta)$, where Λ is a finite set of *states*, and $\delta : \Lambda \times \Lambda \rightarrow \Lambda \times \Lambda$ is the *transition function*.⁶ In this paper we deal with *nonuniform* protocols in which a different Λ and δ are allowed for different population sizes n (one for each possible value of $\lceil \log n \rceil$), but we abuse terminology and refer to the whole family as a single protocol. In all cases (as with similar nonuniform protocols), the nonuniformity is used to embed the value $\lceil \log n \rceil$ into each agent; the transitions are otherwise “uniformly specified”. See Section 8 for more discussion of uniform protocols.

A *configuration* \mathbf{c} of a population protocol is a multiset over Λ of size n , giving the states of the n agents in the population. For a state $s \in \Lambda$, we write $\mathbf{c}(s)$ to denote the count of agents in state s . A *transition* (a.k.a., *reaction*) is a 4-tuple $\alpha = (r_1, r_2, p_1, p_2)$, written $\alpha : r_1, r_2 \rightarrow p_1, p_2$, such that $\delta(r_1, r_2) = (p_1, p_2)$. If an agent in state r_1 interacts with an agent in state r_2 , then they change states to p_1 and p_2 . For every pair of states r_1, r_2 without an explicitly listed transition $r_1, r_2 \rightarrow p_1, p_2$, there is an implicit *null* transition $r_1, r_2 \rightarrow r_1, r_2$ in which the agents interact but do not change state. For our main protocol, we specify transitions formally with pseudocode that indicate how agents alter each independent field in their state.

More formally, given a configuration \mathbf{c} and transition $\alpha : r_1, r_2 \rightarrow p_1, p_2$, we say that α is *applicable* to \mathbf{c} if $\{r_1, r_2\} \subseteq \mathbf{c}$, i.e., \mathbf{c} contains 2 agents, one in state r_1 and one in state r_2 . If α is applicable to \mathbf{c} , then we write $\mathbf{c} \Rightarrow_{\alpha} \mathbf{c}'$, where $\mathbf{c}' = \mathbf{c} - \{r_1, r_2\} + \{p_1, p_2\}$ is the configuration that results from applying α to \mathbf{c} . We write $\mathbf{c} \Rightarrow \mathbf{c}'$, and we say that \mathbf{c}' is *reachable* from \mathbf{c} , if there is a sequence $(\alpha_0, \alpha_1, \dots, \alpha_k)$ of transitions such that $\mathbf{c} \Rightarrow_{\alpha_0} \mathbf{c}_1 \Rightarrow_{\alpha_1} \dots \Rightarrow_{\alpha_k} \mathbf{c}'$. This notation omits mention of \mathcal{P} ; we always deal with a single protocol at a time, so it is clear from context which protocol is defining the transitions.

2.2 Stable majority computation with population protocols

There are many modes of computation considered in population protocols: computing integer-valued functions [13, 26, 31] where the number of agents in a particular state is the output, Boolean-valued predicates [10, 11] where each agent outputs a Boolean value as a function of its state and the goal is for all agents eventually to have the correct output, problems such as leader election [4, 6, 18, 19, 32, 37, 38], and generalizations of predicate computation, where each agent individually outputs a value from a larger range, such as reporting the population size [20, 29, 30]. Majority computation is Boolean-valued if computing the predicate “ $A \geq B$?”, where A and B represent the initial counts of two opinions A and B . We define the slightly generalized problem that requires recognizing when there is a tie, so the range of outputs is $\{A, B, T\}$.

Formally, if the set of states is Λ , the protocol defines a disjoint partition of $\Lambda = \Lambda_A \cup \Lambda_B \cup \Lambda_T$. For $u \in \{A, B, T\}$, if $a \in \Lambda_u$ for all $a \in \mathbf{c}$, we define *output* $\phi(\mathbf{c}) = u$ (i.e., all agents in \mathbf{c} agree on the output u). Otherwise $\phi(\mathbf{c})$ is undefined (i.e., agents disagree on the output). We say \mathbf{o} is *stable* if $\phi(\mathbf{o})$ is defined and, for all \mathbf{o}_2 such that $\mathbf{o} \Rightarrow \mathbf{o}_2$, $\phi(\mathbf{o}) = \phi(\mathbf{o}_2)$, i.e., the output cannot change.

The protocol identifies two special *input states* $A, B \in \Lambda$. A *valid* initial configuration \mathbf{i} satisfies $a \in \{A, B\}$ for all $a \in \mathbf{i}$. We say the *majority opinion* of \mathbf{i} is $M(\mathbf{i}) = A$ if $\mathbf{i}(A) > \mathbf{i}(B)$, $M(\mathbf{i}) = B$ if $\mathbf{i}(A) < \mathbf{i}(B)$,

⁶To understand the full generality of our main protocol, we include randomized transitions in our model. However, there is only one type of randomized transition in the protocol (the “drip reactions” of Phase 3 described in Section 3.2), parameterized by probability p , and in fact we prove the protocol works even when these transitions are deterministic, i.e., when $p = 1$.

and $M(\mathbf{i}) = \top$ if $\mathbf{i}(A) = \mathbf{i}(B)$. The protocol *stably computes majority* if, for any valid initial configuration \mathbf{i} , for all \mathbf{c} such that $\mathbf{i} \Rightarrow \mathbf{c}$, there is a stable \mathbf{o} such that $\mathbf{c} \Rightarrow \mathbf{o}$ and $\phi(\mathbf{o}) = M(\mathbf{i})$. Let $O_{\mathbf{i}} = \{\mathbf{o} : \phi(\mathbf{o}) = M(\mathbf{i})\}$ be the set of all correct, stable configurations. In other words, for any reachable configuration, it is possible to reach a correct, stable configuration, or equivalently reach a strongly connected component in $O_{\mathbf{i}}$.

2.3 Time complexity

In any configuration the next interaction is chosen by selecting a pair of agents uniformly at random and applying an applicable transition, with appropriate probabilities for any randomized transitions. Thus the sequence of transitions and configurations they reach are random variables. To measure time we count the total number of interactions (including null transitions such as $a, b \rightarrow a, b$ in which the agents interact but do not change state), and divide by the number of agents n . In the population protocols literature, this is often called “parallel time”: n interactions among a population of n agents equals one unit of time.

If the protocol stably computes majority, then for any valid initial configuration \mathbf{i} , the probability of reaching a stable, correct configuration, $\mathbb{P}[\mathbf{i} \Rightarrow O_{\mathbf{i}}] = 1$.⁷ We define the *stabilization time* S to be the random variable giving the time to reach a configuration $\mathbf{o} \in O_{\mathbf{i}}$.

When discussing random events in a protocol of population size n , we say event E happens *with high probability* if $\mathbb{P}[\neg E] = O(n^{-c})$, where c is a constant that depends on our choice of parameters in the protocol, where c can be made arbitrarily large by changing the parameters. In other words, the probability of failure can be made an arbitrarily small polynomial. For concreteness, we will write a particular polynomial probability such as $O(n^{-2})$, but in each case we could tune some parameter (say, increasing the time complexity by a constant factor) to increase the polynomial’s exponent. We say event E happens *with very high probability* if $\mathbb{P}[\neg E] = O(n^{-\omega(1)})$, i.e., if its probability of failure is smaller than any polynomial probability.

3 Nonuniform majority algorithm description

The goal of Sections 3 through 7 is to show the following theorem:

Theorem 3.1. *There is a nonuniform population protocol **Nonuniform Majority**, using $O(\log n)$ states, that stably computes majority in $O(\log n)$ stabilization time, both in expectation and with high probability.*

3.1 High-level overview of algorithm

In this overview we use “pseudo-transitions” such as $A, B \rightarrow \mathcal{O}, \mathcal{O}$ to describe agents updating a portion of their states, while ignoring other parts of the state space.

Each agent initially has a **bias**: $+1$ for opinion A and -1 for opinion B , so the population-wide sum $g = \sum_v v \cdot \text{bias}$ gives the *initial gap* between opinions. The majority problem is equivalent to determining $\text{sign}(g)$. Transitions redistribute biases among agents but, to ensure correctness, maintain the population-wide g as an invariant. Biases change through *cancel reactions* $+\frac{1}{2^i}, -\frac{1}{2^i} \rightarrow 0, 0$ and *split reactions* $\pm\frac{1}{2^i}, 0 \rightarrow \pm\frac{1}{2^{i+1}}, \pm\frac{1}{2^{i+1}}$, down to a minimum $\pm\frac{1}{2^L}$. The constant $L = \lceil \log_2(n) \rceil$ ensures $\Theta(\log n)$ possible states. The *gap* is defined to be $\sum_v \text{sign}(v \cdot \text{bias})$, the difference in counts between majority and minority biases. Note the gap should grow over time to spread the correct majority opinion to the whole population, while the invariant g should ensure correctness of the final opinion.

The cancel and split reactions average the bias value between both agents, but only when the average is also a power of 2, or 0. If we had averaging reactions between all pairs of biases (also allowing, e.g., $\frac{1}{2}, \frac{1}{4} \rightarrow \frac{3}{8}, \frac{3}{8}$), then all biases would converge to $\frac{g}{n}$, but this would use too many states.⁸ With our limited set $\{0, \pm\frac{1}{2}, \pm\frac{1}{4}, \dots, \pm\frac{1}{2^L}\}$ of possible biases, allowing all cancel and split reactions simultaneously does not work. Most biases appear simultaneously across the population, reducing the count of each bias, which slows

⁷Since population protocols have a finite reachable configuration space, this is equivalent to the stable computation definition that for all \mathbf{c} reachable from \mathbf{i} , there is a $\mathbf{o}' \in O_{\mathbf{i}}$ reachable from \mathbf{c} .

⁸This was effectively the approach used for majority in [8, 44], for an $O(n)$ state, $O(\log n)$ time protocol.

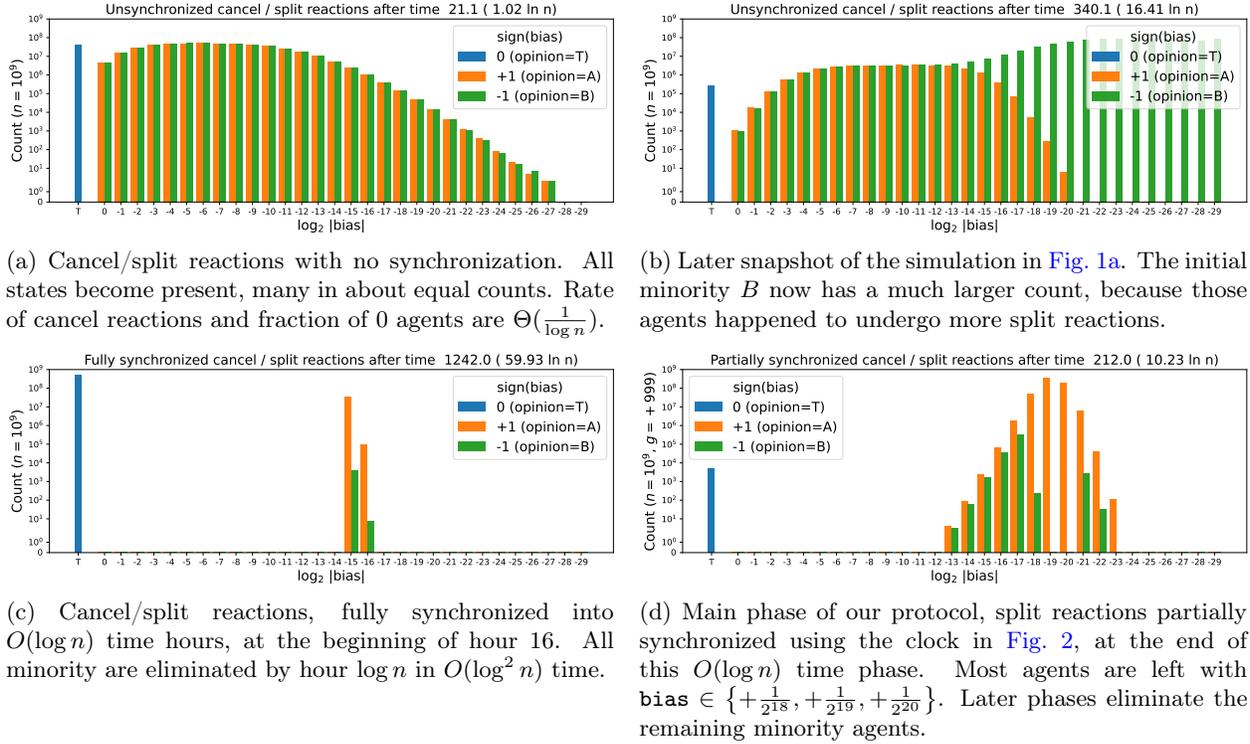


Figure 1: Cancel / split reactions with no synchronization (1a,1b), perfect synchronization (1c), and partial synchronization (1d) via the fixed-resolution phase clock of our main protocol. Plots generated from [1].

the rate of cancel reactions. Then the count of unbiased 0 agents is reduced, which slows the rate of split reactions, see Fig. 1a. Also, there is a non-negligible probability for the initial minority opinion to reach a much greater count, if those agents happen to do more split reactions, see Fig. 1b. Thus using only the count of positive versus negative biases will not work to solve majority even with high probability.

To solve this problem, we partially synchronize the unbiased agents with a field hour, adding $\log n$ states $0_0, 0_1, 0_2, \dots, 0_L$. The new split reactions

$$\pm \frac{1}{2^i}, 0_h \rightarrow \pm \frac{1}{2^{i+1}}, \pm \frac{1}{2^{i+1}} \quad \text{if } h > i$$

will wait until $\text{hour} \geq h$ before doing splits down to $\text{bias} = \pm \frac{1}{2^h}$. We could use existing phase clocks to perfectly synchronize hour , by making each hour use $\Theta(\log n)$ time, enough time for all opinionated agents to split. Then WHP all agents would be in states $\{0_h, +\frac{1}{2^h}, -\frac{1}{2^h}\}$ by the end of hour h , see Fig. 1c. The invariant $g = \sum_v v \cdot \text{bias}$ implies that all minority opinions would be eliminated by hour $\lceil \log_2 \frac{1}{g} \rceil \leq L$. This would give an $O(\log n)$ -state, $O(\log^2 n)$ -time majority algorithm, essentially equivalent to [5, 17].

The main idea of our algorithm is to use these rules with a faster clock using only $O(1)$ time per hour. The hour field of unbiased agents is synchronized to a separate subpopulation of clock agents, who use a field minute , with k consecutive minutes per hour. Minutes advance by *drip reactions* $C_i, C_i \rightarrow C_i, C_{i+1}$, and catch up by *epidemic reactions* $C_i, C_j \rightarrow C_{\max(i,j)}, C_{\max(i,j)}$. See Fig. 2 for an illustration of the clock minute and hour dynamics.

Since $O(1)$ time per hour is not sufficient to bring all agents up to the current hour before advancing to the next, we now have only a large constant fraction of agents, rather than all agents, synchronized in the current hour. Still, we prove this looser synchronization keeps the values of hour and bias relatively concentrated, so by the end of this phase, we reach a configuration as shown in Fig. 1d. Most agents have

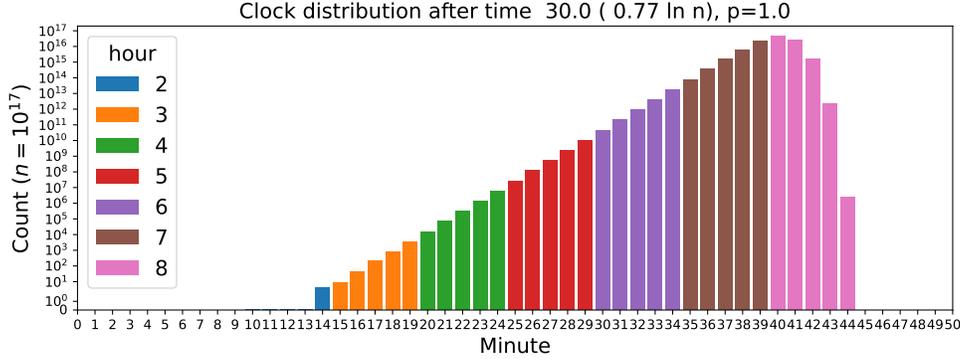


Figure 2: Clock rules of our protocol, showing a travelling wave distribution over minutes, on a larger population of size $n = 10^{17}$ to emphasize the distribution. The distribution’s back tail decays exponentially, and its front tail decays doubly exponentially. A large constant fraction of agents are in the same two consecutive `hour`’s (here 7 and 8). Plot generated from [1].

the majority opinion (WLOG positive), with three consecutive biases $+\frac{1}{2^t}, +\frac{1}{2^{t+1}}, +\frac{1}{2^{t+2}}$.

Detecting ties. This algorithm gives an elegant way to detect a tie with high probability. In this case, $g = 0$, and with high probability, all agents will finish the phase with `bias` $\in \{0, \pm\frac{1}{2^L}\}$. Checking this condition *stably* detects a tie (i.e., with probability 1, if this condition is true, then there is a tie), because for any nonzero value of g , there must be some agent with $|\text{bias}| > \frac{1}{2^L}$.

Cleanup Phases. We must next eliminate all minority opinions, while still relying on the invariant $g = \sum_v v \cdot \text{bias}$ to ensure correctness. Note that it is possible with low probability to have a greater count of minority opinions (with smaller values of `bias`), so only relying on counts of positive and negative biases would give possibilities of error.

We first remove any minority agents with large bias, by using an additional subpopulation of Reserve agents that enable additional split reactions for large values of $|\text{bias}| > \frac{1}{2^t}$. Then after cancel reactions with the bulk of majority agents, the only minority agents left must have $|\text{bias}| < \frac{1}{2^{t+2}}$.

To then remove minority agents with small bias, we allow agents with larger bias to “consume” agents with smaller bias, such as an interaction between agents $+\frac{1}{4}$ and $-\frac{1}{256}$. Here the positive agent can be thought to hold the entire bias $+\frac{1}{4} - \frac{1}{256} = +\frac{63}{256}$, but since this value is not in the allowable states, it can only store that its bias is in the range $+\frac{1}{8} \leq \text{bias} \leq +\frac{1}{4}$. Without knowing its exact `bias`, this agent cannot participate in future averaging interactions. However, we show there are enough available majority agents to eliminate all remaining minority via these *consumption reactions*. Thus with high probability, all minority agents are eliminated.

A final phase checks for the presence of both positive and negative `bias`, and if one has been completely eliminated, it stabilizes to the correct output. In the case where both are present, this is a detectable error, where we can move to a slow, correct backup that uses the original inputs. Due to the low probability of this case, it contributes negligibly to the total expected time.

3.2 Intuitive description of each phase

Our full protocol is broken up into 11 consecutive phases. We describe each phase intuitively before presenting full pseudocode in Section 3.4. Note that some further separation of phases was done to create more straightforward proofs of correctness, so simplicity of the proofs was optimized over simplicity of the full protocol pseudocode. It is likely possible to have simpler logic that still solves majority via the same strategy.

Phase 0: “Population splitting” [7] divides agents into roles used in subsequent phases: Main, Reserve, Clock. In timed phases (those not marked as *Untimed* or *Fixed-resolution clock*, including the current phase), Clock agents count from $\Theta(\log n)$ to 0 to cause the switch to the next phase after $\Theta(\log n)$ time.

“Standard” population splitting uses reactions such as $x, x \rightarrow r_1, r_2$ to divide agents into two roles r_1, r_2 . This takes $\Theta(n)$ time to converge, which can be decreased to $\Theta(\log n)$ time via $r_1, x \rightarrow r_1, r_2$ and $r_2, x \rightarrow r_2, r_1$, while maintaining that $\#r_1$ and $\#r_2$ are both $n/2 \pm \sqrt{n}$ WHP. However, since all agents initially have an opinion, but Clock and Reserve agents do not hold an opinion, agents that adopt role Clock or Reserve must first pass off their opinion to a Main agent.

From each interacting pair of unassigned agents, one will take the Main role and hold the opinions of both agents, interpreting A as $+1$ and B as -1 . This Main agent will then be allowed to take at most one other opinion (in an additional reaction that enables rapid convergence of the population splitting), and holding 3 opinions can end up with a bias in the range $\{-3, -2, -1, 0, +1, +2, +3\}$.

Phase 1: Agents do “integer averaging” [6] of biases in the set $\{-3, \dots, +3\}$ via reactions $i, j \rightarrow \lfloor \frac{i+j}{2} \rfloor, \lceil \frac{i+j}{2} \rceil$. Although taking $\Theta(n)$ time to converge in some cases, this process is known [45] to result in three consecutive values in $O(\log n)$ time. If those three values are detected to be $\{-1, 0, +1\}$ in the next phase, the algorithm continues.

Phase 2: (*Untimed*) Agents propagate the set of opinions (signs of biases) remaining after Phase 1 to detect if only one opinion remains. If so, we have converged on a majority consensus, and the algorithm halts here (see Fig. 6). At this point, this is essentially the exact majority protocol of [41], which takes $O(\log n)$ time with an initial gap $\Omega(n)$, but longer for sublinear gaps (e.g., $\Omega(n)$ time for a gap of 1). Thus, if agents proceed beyond this phase (i.e., if both opinions A and B remain at this point), we will use later that the gap was smaller than $0.025 \cdot \#\text{Main}$. With low probability both opinions remain but some agent has $|\text{bias}| > 1$, in which case we proceed directly to a slow stable backup protocol in Phase 10.

Phase 3: (*Fixed-resolution clock*) The key goal at this phase is to use cancel and split reactions to average the bias across the population to give almost all agents the majority opinion. Biased agents hold a field $\text{exponent} \in \{-L, \dots, -1, 0\}$, describing the magnitude $|\text{bias}| = 2^{\text{exponent}}$, a quantity we call the agent’s *mass*. Cancel reactions eliminate opposite biases $+\frac{1}{2^i}, -\frac{1}{2^i} \rightarrow 0, 0$ with the same exponent ; cancel reactions strictly reduce total mass. Split reactions $\pm\frac{1}{2^i}, 0 \rightarrow \pm\frac{1}{2^{i+1}}, \pm\frac{1}{2^{i+1}}$ give half of the bias to an unbiased agent, decrementing the exponent ; split reactions preserve the total mass. The unbiased \mathcal{O} agents, with $\text{role} = \text{Main}$, $\text{opinion} = \text{bias} = 0$, act as the fuel for split reactions.

We want to obtain tighter synchronization in the exponents than Fig. 1a, approximating the ideal synchronized behavior of the $O(\log^2 n)$ time algorithm of Fig. 1c while using only $O(\log n)$ time. To achieve this, the Clock agents run a “fixed resolution” clock that keeps them roughly synchronized (though not perfectly; see Fig. 2) as they count their “minutes” from 0 up to $L' = kL$, using $O(1)$ time per minute. This is done via “drip” reactions $C_i, C_i \rightarrow C_i, C_{i+1}$ (when minute i gets sufficiently populated, pairs of C_i agents meet with sufficient likelihood to increment the minute) and $C_j, C_i \rightarrow C_j, C_j$ for $i < j$ (new higher minute propagates by epidemic).⁹ If randomized transitions are allowed, by lowering the probability p of the drip reaction, the clock rate can be slowed by a constant factor.

⁹This clock is similar to the power-of-two-choices leaderless phase clock of [5], where the agent with smaller (or equal) minute increments their clock ($C_j, C_i \rightarrow C_j, C_{i+1}$ for $i \leq j$), but increasing the smaller minute by only 1. Similarly to our clock, the maximum minute can increase only with both agents at the same minute. A similar process was analyzed in [15], and in fact was shown to have the key properties needed for our clock to work—an exponentially-decaying back tail and a double-exponentially-decaying front tail—so it seems likely that a power-of-two-choices clock could also work with our protocol.

The randomized variant of our clock with drip probability p is also similar to the “junta-driven” phase clock of [37], but with a linear number $2pn$ of agents in the junta, using $O(1)$ time per minute, rather than the $O(n^\epsilon)$ -size junta of [37], which uses $O(\log n)$ time per minute. There, smaller minutes are brought up by epidemic, and only an agent in the junta seeing another agent at the same minute will increment. The epidemic reaction is exactly the same in both rules. The probability p of a drip reaction can be interpreted as the probability that one of the agents is in the junta. For similar rate of $O(1)$ time per minute phase clock construction see also Dudek and Kosowski work [34].

Although we prove a few lemmas about this generalized clock, and some of our simulation plots in [Section 3.3](#) use $p < 1$, our proofs work even for $p = 1$, i.e., a deterministic transition function, although this requires constant-factor more states (by increasing the number of “minutes per hour”, explained next).

Now the \mathcal{O} agents will use $\Theta(\log n)$ states to store an “hour”, coupled to the C clock agents via $C_{\lfloor i/k \rfloor}, \mathcal{O}_j \rightarrow C_{\lfloor i/k \rfloor}, \mathcal{O}_{\lfloor i/k \rfloor}$ if $\lfloor i/k \rfloor > j$, i.e., every consecutive k Clock minutes corresponds to one Main hour, and clock agents drag \mathcal{O} agents up to the current hour (see [Fig. 3](#)). Our proofs require $k = 45$ minutes per hour when $p = 1$, but smaller values of k work in simulation. For example, the simulation in [Fig. 1d](#) showing intended behavior of this phase used only $k = 3$ minutes per hour with $p = 1$.

This clock synchronizes the exponents because agents with `exponent` = $-i$ can only split down to `exponent` = $-(i + 1)$ with an \mathcal{O} agent that has `hour` $\geq i + 1$. This prevents the biased agents from doing too many splits too quickly. As a result, during hour i , most of the biased agents have `bias` = $\frac{1}{2^i}$, so the cancel reactions $+\frac{1}{2^i}, -\frac{1}{2^i} \rightarrow 0, 0$ happen at a high rate, providing many \mathcal{O} agents as “fuel” for future split reactions. We tune the constants of the clock to ensure hour i lasts long enough to bring most biased agents down to `exponent` = $-i$ via split reactions and then let a good fraction do cancel reactions (see [Fig. 4c](#)).

The key property at the conclusion of this phase is that unless there is a tie, WHP most majority agents end up in three consecutive exponents $-l, -(l + 1), -(l + 2)$, with a negligible mass of any other Main agent (majority agents at lower/higher exponents, minority agents at any exponent, or \mathcal{O} agents).¹⁰ Phases 5-7 use this fact to quickly push the rest of the population to a configuration where *all* minority agents have exponents strictly below $-(l + 2)$; [Phase 8](#) then eliminates these minority agents quickly.

Phase 4: (*Untimed*) The special case of a tie is detected by the fact that, since the total bias remains the initial gap g , if all biased agents have minimal exponent $-L$, g has magnitude less than 1:

$$\begin{aligned} |g| &= \left| \sum_{a.\text{role}=\text{Main}} a.\text{bias} \right| \leq \sum_{a.\text{role}=\text{Main}} |a.\text{bias}| \\ &\leq \sum_{a.\text{role}=\text{Main}} \frac{1}{2^L} < \frac{n}{2^{\lceil \log_2(n) \rceil}} \leq 1. \end{aligned}$$

The initial gap g is integer valued, so $|g| < 1 \implies g = 0$. Thus this condition implies there is a tie with probability 1; the converse that a tie forces all biased agents to exponent $-L$ holds with high probability. If only exponent $-L$ is detected, the algorithm halts here with all agents reporting output \top (see [Fig. 7](#)). Otherwise, the algorithm proceeds to the next phase.

Phase 5, Phase 6: Using the key property of [Phase 3](#), these phases WHP pull all biased agents above exponent $-l$ down to exponent $-l$ or below using the Reserve R agents. The R ’s activate themselves in [Phase 5](#) by sampling the exponent of the first biased agent they meet. This ensures WHP that sufficiently many Reserve agents exist with exponents $-l, -(l + 1), -(l + 2)$ (distributed similarly to the agents with those exponents). Then in [Phase 6](#), they act as fuel for splits, via $R_i, \pm \frac{1}{2^j} \rightarrow \pm \frac{1}{2^{j+1}}, \pm \frac{1}{2^{j+1}}$ when $|i| > |j|$. The reserve agents, unlike the \mathcal{O} agents in [Phase 3](#), do not change their exponent in response to interactions with Clock agents. Thus sufficiently many reserve agents will remain to allow the small number of biased agents above exponent $-l$ to split down to exponent $-l$ or below.¹¹

¹⁰ l is defined such that if all biased agents were at exponent $-l$, the difference in counts between majority and minority agents would be between $0.4 \cdot \#\text{Main}$ and $0.8 \cdot \#\text{Main}$.

¹¹The reason we do this in two separate phases is to ensure that the Reserve agents have close to the same distribution of exponents that the Main agents have at the end of [Phase 3](#). If Reserve agents allowed split reactions in the same phase that they sample the `exponent` of Main agents, then the splits would disrupt the distribution of the Main agents before all Reserve agents have finished sampling. This would possibly give the Reserve agents a significantly different distribution among levels than the Main agents had at the start. While this may possibly work anyway, we find it is more straightforward to prove if the Reserve agents have a close distribution over `exponent` values to that of the Main agents.

Phase 7: This phase allows more general reactions to distribute the dyadic biases, allowing reactions between agents up to two exponents apart, to eliminate the opinion with smaller exponent: $\frac{1}{2^i}, -\frac{1}{2^{i+1}} \rightarrow \frac{1}{2^{i+1}}, 0$ and $\frac{1}{2^i}, -\frac{1}{2^{i+2}} \rightarrow \frac{1}{2^{i+1}}, \frac{1}{2^{i+2}}$ (and the equivalent with positive/negative biases swapped). Since all agents have exponent $-l$ or below, and many more majority agents exist at exponents $-l, -(l+1), -(l+2)$ than the total number of minority agents anywhere, these (together with standard cancel reactions $\frac{1}{2^i}, -\frac{1}{2^i} \rightarrow 0, 0$) rapidly eliminate all minority agents at exponents $-l, -(l+1), -(l+2)$, while maintaining $\Omega(n)$ majority agents at exponents $\geq -(l+2)$ and $< 0.01n$ total minority agents, now all at exponents $\leq -(l+3)$.

Phase 8: This phase eliminates the last minority agents, while ensuring that if any error occurred in previous phases, some majority agents remain, to allow detecting the error by the presence of both opinions.¹²

The biased agents add a Boolean field `full`, initially `False`, and *consumption* reactions that allow an agent at a larger exponent i to consume (set to mass 0 by setting it to be \mathcal{O}) an agent at an arbitrarily smaller exponent $j < i$. Now the remaining agent represents some non-power-of-two mass $m = 2^i - 2^j$, which it lacks sufficient memory to track exactly. Thus setting the flag `full = True` corresponds to the agent having an uncertain mass m in the range $2^{i-1} \leq m < 2^i$. Because of this uncertainty, full agents are not allowed to consume other smaller levels. However, there are more than enough high-exponent majority agents by this phase to consume all remaining lower exponent minority agents.

Crucially, agents that have consumed another agent and set `full = True` may themselves then be consumed by a third agent (with `full = False`) at an even larger exponent. This is needed because a minority agent at exponent $i \leq -(l+3)$ may consume a (rare) majority agent at exponent $j < i$, but the minority agent itself can be consumed by another majority agent with exponent $k > i$.

Phase 9: (*Untimed*) This is identical to **Phase 2**: it detects whether both biased opinions A and B remain. If not (the likely case), the algorithm halts, otherwise we proceed to the next phase.

Phase 10: (*Untimed*) Agents execute a simple, slow stable majority protocol [22], similar to that of [33, 41] but also handling ties. This takes $\Theta(n \log n)$ time, but the probability that an earlier error forces us into this phase is $O(1/n^2)$, so it contributes negligibly to the total expected time.

3.3 Simulation of full algorithm

In this section we show simulation results, where the complete pseudocode of **Section 3.4** was translated into Java code available on GitHub [2]. In these simulations, we stop the protocol once all agents reach **Phase 9**. For the low probability case that agents switch to the **Phase 10**, the simulator prints an error indicating the agents should switch to slow back up, but as expected this was not observed in our simulations. For all our plots, we collect data from simulations with $n \approx 2^{23}$, p (drip probability) = 0.1. The first simulation in **Fig. 3** shows the relationship between `minute` of Clock agents and `hour` of Main agents. Here we used $k = 5$ minutes per hour to show clearly the relationship and the discrete nature of the hours.

All remaining simulations used the even weaker value $k = 2$, to help see enough low probability behavior that the logic enforcing probability-1 correctness in later phases is necessary. We show 3 simulations corresponding to the 3 different types of initial gap. **Figs. 4** and **5** show constant initial gap $g = +2$. This is our “typical” case, where the simulation eventually stabilizes to the correct output in **Phase 9**. **Fig. 6** shows linear initial gap $g \approx \frac{n}{10}$, which stabilizes in **Phase 2** after quickly cancelling all minority agents. **Fig. 7** shows initial tie $g = 0$, which stabilizes in **Phase 4** after all biased agents reach the minimum `exponent` = $-L = -23$.

All three simulations show various snapshots of configurations of the biased agents. These particular snapshots are at special times marked in **Figs. 4d, 6a** and **7a**. In all cases, an animation is available at GitHub [3], showing the full evolution of these distributions over all recorded time steps from the simulation.

¹²A naïve idea to reach a consensus at this phase is to allow cancel reactions $\frac{1}{2^i}, -\frac{1}{2^i} \rightarrow 0, 0$ between arbitrary pairs of exponents with opposite opinions. However, this has a positive probability of erroneously eliminating the majority. This is because the majority, while it necessarily has larger mass than the minority at this point, could have smaller *count*. For example, we could have 16 *A*’s with `exponent` = -2 and 32 *B*’s with `exponent` = -5 , so *A*’s have mass $16 \cdot 2^{-2} = 4$ and *B*’s have smaller mass $32 \cdot 2^{-5} = 1$, but larger count than *A*.

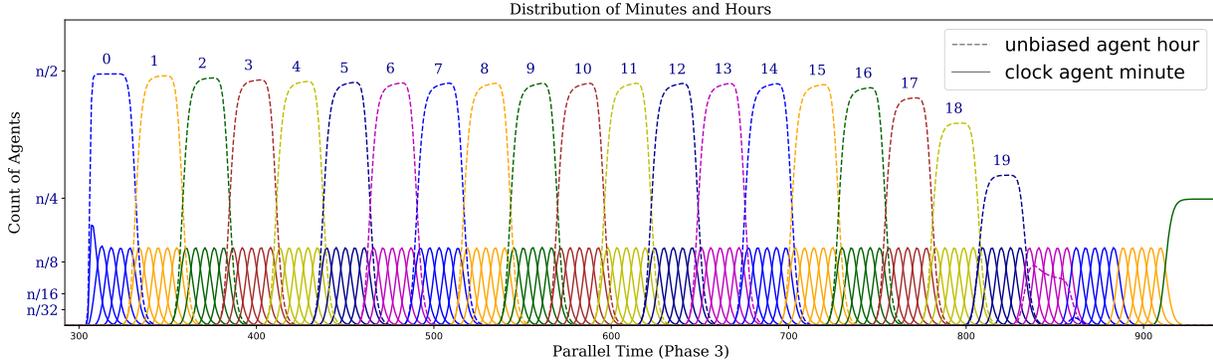


Figure 3: Fixed-resolution phase clock used in Phase 3, with $n \approx 2^{23}$. All Clock agents set `minute` = 0 and count up to $k \cdot L$ with special rules defined in Phase 3. The solid curves show the distributions of the counts at each value of `minute` in the Clock agents. Blocks of $k = 5$ consecutive minutes correspond to a hour in the unbiased \mathcal{O} agents with `role` = Main and `bias` = 0. The dashed curves show the distribution of the counts of each value of `hour`, with the value written on top. We show every k minutes with one color equal to its hour color. Near the end of Phase 3, the \mathcal{O} count falls to 0 as the majority count takes over. This plot used the value $k = 5$ to emphasize the clock behavior. Later plots will all use the weaker constant $k = 2$. In later plots we also omit the `minute` distribution and only show the `hour` from the \mathcal{O} agents.

3.4 Algorithm pseudocode

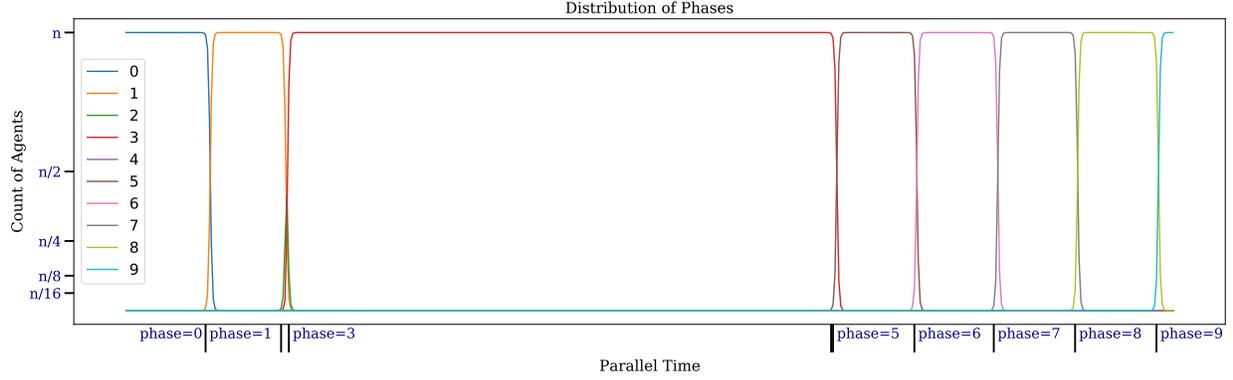
In this section we give a full formal description of the main algorithm.

Every agent starts with a read-only field `input` $\in \{A, B\}$, a field `output` $\in \{A, B, T\}$ corresponding to outputs that the majority is A, B, or a tie. The protocol is broken up into 11 consecutive phases, marked by the additional field `phase` = 0 $\in \{0, \dots, 10\}$. The phase updates via the epidemic reaction $u.\text{phase}, v.\text{phase} \leftarrow \max(u.\text{phase}, v.\text{phase})$. Some fields are only used in particular phases, to ensure the total state space is $\Theta(\log n)$.¹³ Such fields and the initial behavior of an agent upon entering a phase are described in the **Init** section above each phase. Whenever an agent increments their `phase`, they execute **Init** for the new phase (and sequentially for any phases in between if they happen to increment `phase` by more than 1). We refer to the “end of phase i ” to mean the time when the first agent sets `phase` $\leftarrow i + 1$. Note that the agents actually enter each new phase by epidemic, so there is technically no well-defined “beginning of phase i ”. To simplify the analysis, we formally start our arguments for each phase assuming each agent is in the current phase, although technically $\Theta(\log n)$ time will pass between the time the first agent enters phase i and the last agent does.

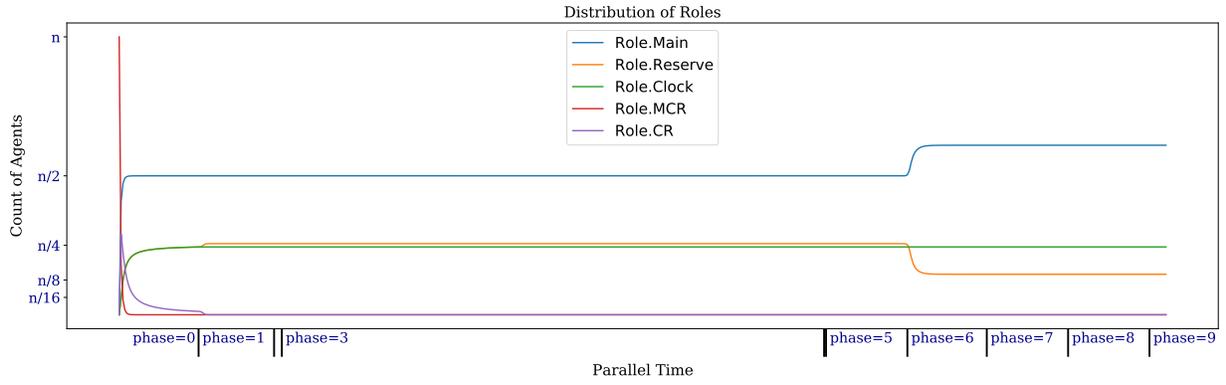
Each timed phase i each requires setting agents to count from $c_i \ln n$ down to 0, where the minimum required value of c_i depends on the phase. These constants can be derived from the technical analysis in Sections 5-7 but for brevity we avoid giving them concrete values in the pseudocode. Our simulations (Section 3.3) that used the same small constant $5 \log_2(n)$ for all counters seem to work, but the proofs require larger constants to ensure the necessary behavior within each phase can complete with high probability $1 - O(1/n^2)$. By increasing these constants c_i (along with changing the phase clock constants p, k ; see Phase 3 and Theorem 6.9), we could also push high probability bound to $1 - O(1/n^c)$ for any desired constant c . For concreteness, use $1 - O(1/n^2)$ for most high probability guarantees, since this is large enough to take appropriate union bounds and ensure the extra time from low probability failures does not contribute meaningfully to the total $O(\log n)$ time bound.

Phase 0 is a timed phase that splits the population into three subpopulations: Main to compute majority, Clock to time the phases and the movement through exponents in Phase 3, and Reserve to aid in cleanup during Phase 6. An agent can only move into role Clock or Reserve by “donating” its opinion to a Main agent, who can collect up to two other opinions in addition to their own, leading to a `bias` of up to ± 3 . After this phase, the populations of the three roles are near the expected one quarter Main, one quarter Clock, and one half Reserve. Lemma 5.2 shows that all initial opinions have been given to assigned Main agents and these

¹³Note that using two fields, both with $O(\log n)$ possible values, requires $O(\log^2 n)$ states, not $O(\log n)$.



(a) The `phase` distribution, giving counts of the agents with each value of `phase`. We show these markers on the horizontal axis in later plots. We have set `max_counter = 5 \log_2(n)`, and removed the counting requirements for `Clock` agents in `Phase 0`. This makes all timed phases `0, 1, 5, 6, 7, 8` take around the same parallel time $2.5 \log_2(n)$. The fixed-resolution clock in `Phase 3` uses $O(\log n)$ time with a larger constant. `Phase 2` and `Phase 4` are untimed, so they end almost immediately and are not labeled above.

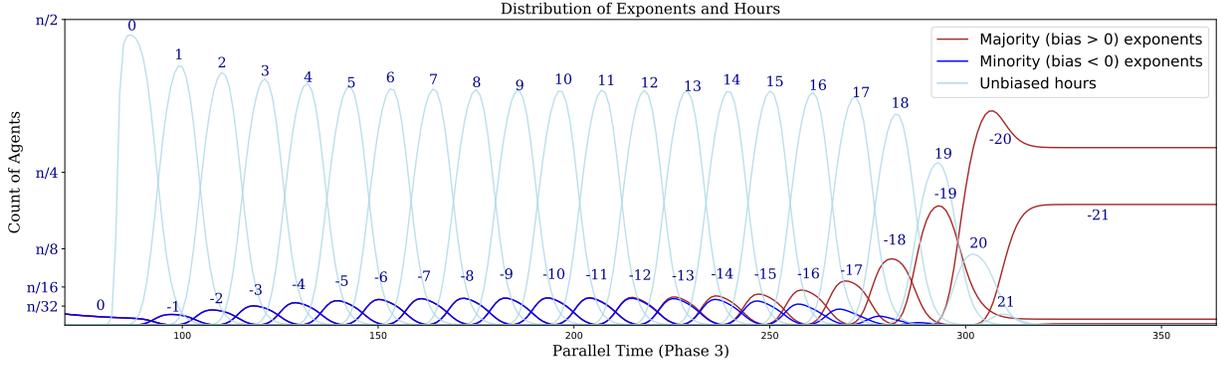


(b) The `role` distribution. All agents start in role `Role_MCR`. By the end of `Phase 0` almost all agents decide on a role. They enter `Phase 1` with $\approx \frac{n}{2}$, $\approx \frac{n}{4}$, and $\approx \frac{n}{4}$ agents in the respective roles `Main`, `Clock`, and `Reserve`. An agent's role remain the same in the following phases except `Phase 6`, where split reactions convert `Reserve` agents into `Main` agents.

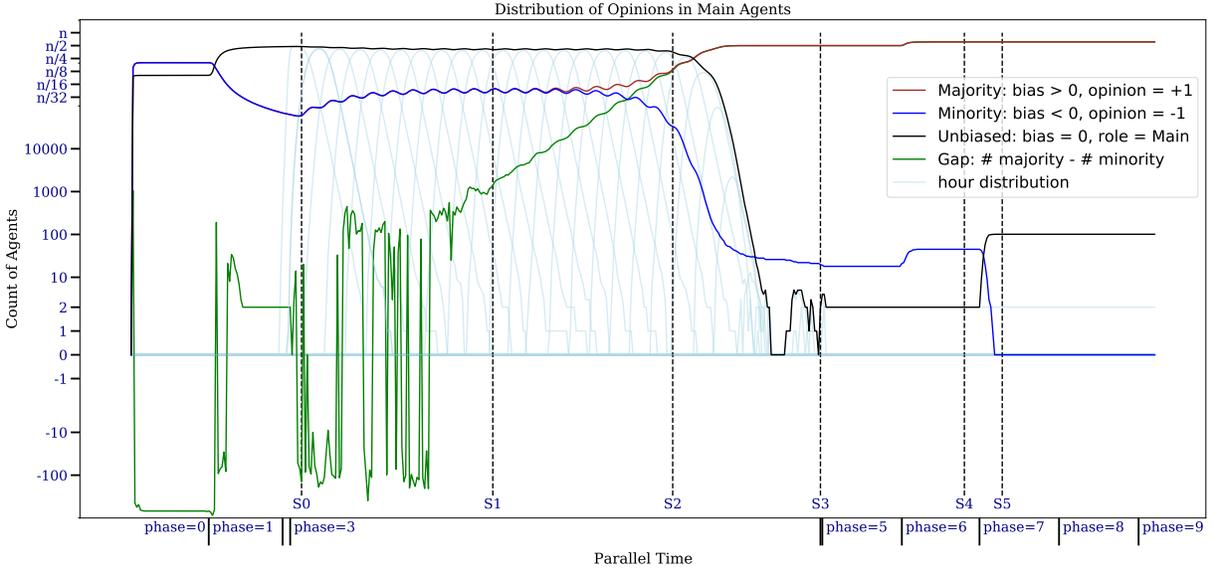
Figure 4: The `phase` and `role` distributions.

Nonuniform Majority_L(u, v). Nonuniform majority algorithm for population sizes n with $L = \lceil \log n \rceil$.
Init: `phase` $\leftarrow 0 \in \{0, 1, \dots, 10\}$ and execute **Init** for `Phase 0`.

- 1: **if** $i.\text{phase} < j.\text{phase}$ where $\{i, j\} = \{u, v\}$ **then**
 - 2: **for** $p = \{i.\text{phase} + 1, \dots, j.\text{phase}\}$ **do**
 - 3: execute **Init** for **Phase** p on agent i
 - 4: $i.\text{phase} \leftarrow j.\text{phase}$
 - 5: execute **Phase** $u.\text{phase}(u, v)$
-

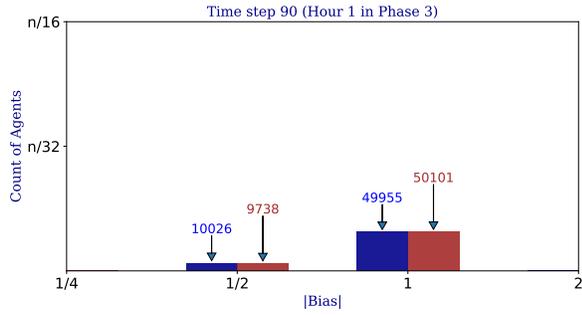


(c) The distribution of **exponent** and **hour** in biased and unbiased Main agents. This plot only shows the time during **Phase 3**, the only time **hour** is used. The values for **exponent** can decrease from 0 to $-L$. As described in **Phase 3**, during hour h , only agents with **exponent** $> -h$ are allowed to split and decrease their **exponent** by one. Thus, the changes of **exponent** are synchronized with the changes in the **hour** values. In this simulation, the **Phase 3** stopped with majority of agents having 3 consecutive values ($-19, -20, -21$) in their **exponent**. The **hour** values are shown behind the next plot in **Fig. 4d**, which shows the totals of these 3 types of Main agents.

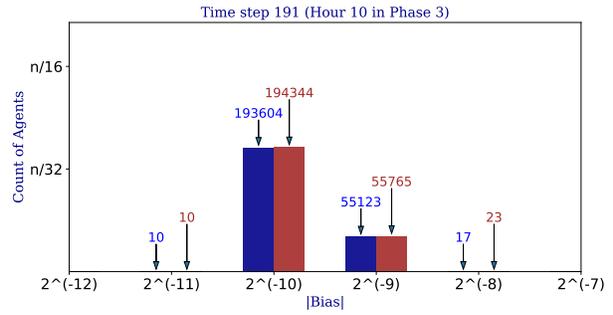


(d) The distribution of **opinion** over all Main agents, with count shown on a log scale. The green line shows the difference in count between majority and minority agents. Special snapshots at times marked S0, S1, S2, S3, S4, S5 are shown in **Fig. 5**. All Main agents are assigned in **Phase 0**, with **bias** $\in \{0, \pm 1, \pm 2, \pm 3\}$, and all initial biases represented held by this Main subpopulation. Here the gap depends randomly on the distribution of $|\mathbf{bias}| \in \{1, 2, 3\}$. Then **Phase 1** brings all **bias** $\in \{-1, 0, +1\}$, so the gap returns to exactly the initial gap of 2, and the biased counts decrease polynomially like $\frac{1}{t}$ from cancel reactions. In the first part of **Phase 3** the gap oscillates randomly about 0, (S0 in **Fig. 5a**). Once we reach a high enough hour / low enough exponent, the doubling trend takes over and the gap undergoes constant exponential growth (S1 in **Fig. 5b**). Finally, this becomes visible as a separation between counts of majority and minority agents (S2 in **Fig. 5c**). **Phase 3** ends with a small but nonzero count of minority agents and the count of unbiased \mathcal{O} agents brought near 0 (S3 in **Fig. 5d**). Then during **Phase 6**, this minority count is amplified slightly by more split reactions bringing the minority exponents down (S4 in **Fig. 5e**). During **Phase 7**, additional cancel reactions bring the minority count to 0 (S5 in **Fig. 5f**). Since minority agents are gone, **Phase 8** has no effect, and the protocol stabilizes to the correct majority output in **Phase 9**.

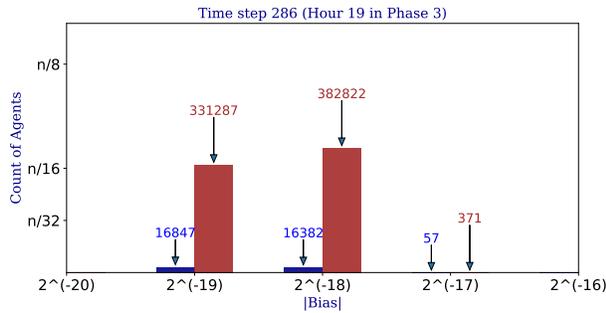
Figure 4: The case of constant initial gap $g = 2$, from simulation with $n = 5122666 \approx 2^{23}$, $p = 0.1$ (drip probability), $k = 2$ (number of minutes per hour). The horizontal axis is in units of parallel time, with the ranges corresponding to each phase marked. All agents converge to the correct majority output in **Phase 9**.



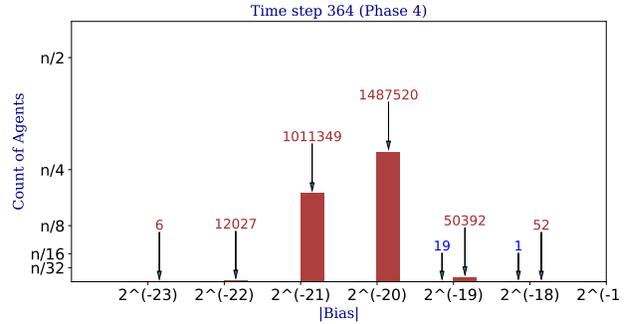
(a) Snapshot S0. At the start of **Phase 3**. The count of minority agents (blue) currently exceeds the count of majority agents because the minorities have done more split reactions. Summing the signed biases, however, gives $+50101 - 49955 + \frac{9738}{2} - \frac{10026}{2} = +146 - 144 = +2$, the invariant initial gap.



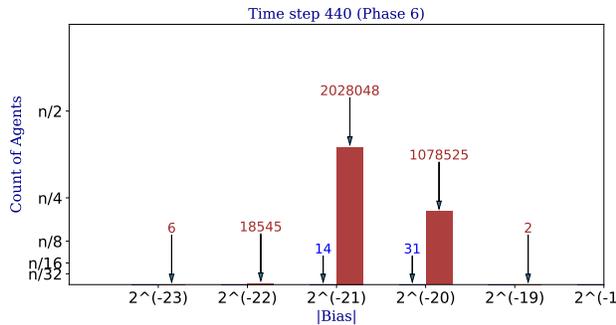
(b) Snapshot S1. A typical distribution at **hour = 10**, when split reactions have brought most agents to **exponent = -10**. Only a few \mathcal{O} agents leaked ahead to **hour = 11** and enabled splits down to **exponent = -11**, and no agents have leaked further ahead than this.



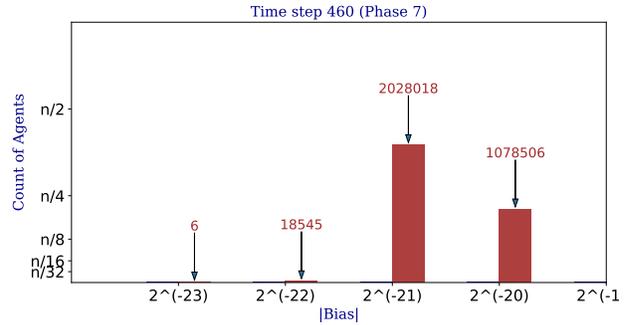
(c) Snapshot S2. We have reached the special exponent $-l = -19$. The count of minority (blue) agents has vastly decreased over the last few hours, and now there will be few more cancel reactions to produce more \mathcal{O} agents.



(d) Snapshot S3. We end **Phase 3** with most Main agents with the majority opinion, and **bias** $\in \{-19, -20, -21\}$ in a range of 3 consecutive values, as shown in **Theorem 6.2**. Only a few minority agents are left.

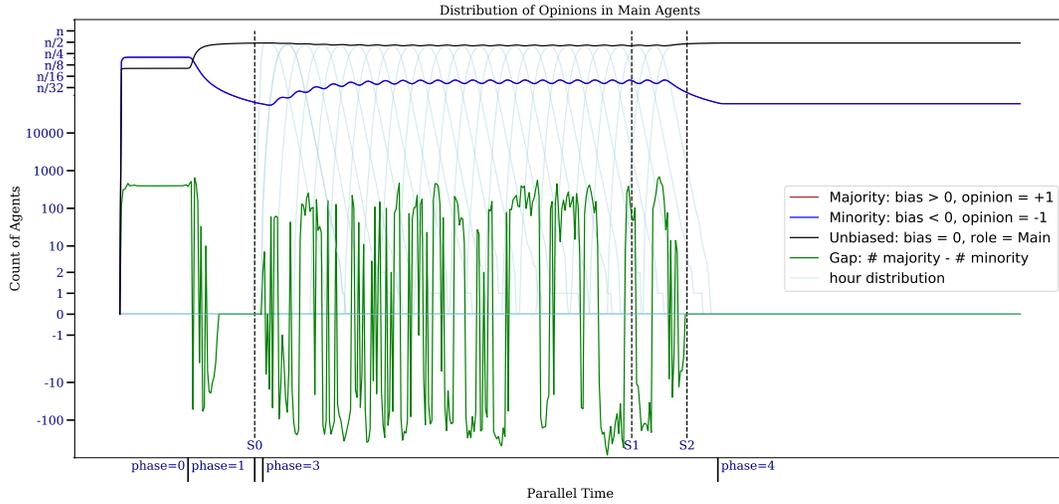


(e) Snapshot S4. After **Phase 6**, where Reserve agents with **sample** $\in \{-19, -20, -21\}$ enabled additional split reactions that brought all minority agents down.

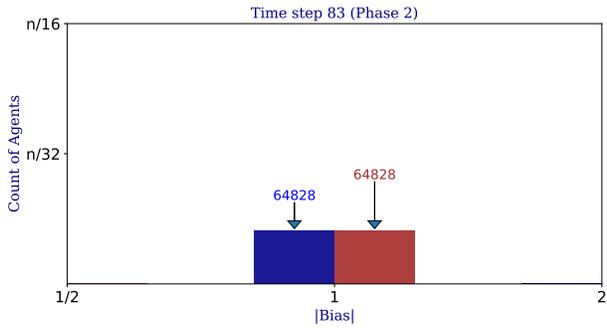


(f) Snapshot S5. After **Phase 7**, additional generalized cancel reactions eliminate all minority agents at the higher **exponent** values. At this point, there are no minority agents left, and we will stabilize to the correct output in **Phase 9**.

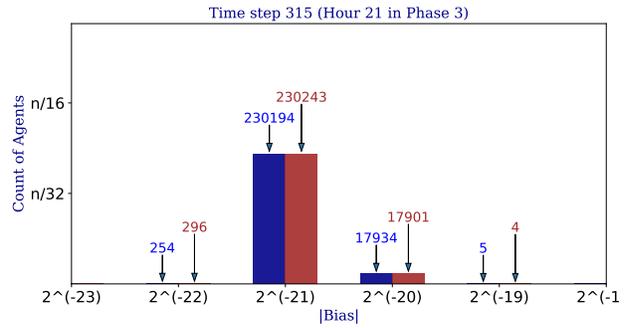
Figure 5: Snapshots S0, S1, S2, S3, S4, S5 from **Fig. 4d**, showing the distribution of **bias** among agents with $|\mathbf{bias}| > 0$. The red and blue bars give the count of A (majority, with $\mathbf{bias} > 0$) and B (minority, with $\mathbf{bias} < 0$) agents respectively. The exact counts are written above the bars, and everywhere not explicitly written the count of is 0. See the link to animations of these plots at the end of this section.



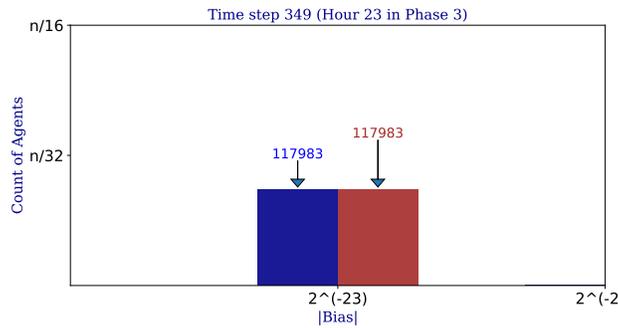
(a) The distribution of opinion in the Main agents in case of a tie. The red line gives the majority count (opinion = +1), the blue line the minority count (opinion = -1), though they overlap everywhere, and the green line their difference. The black line gives the unbiased \mathcal{O} agents. In Phase 3 we see the same qualitative behavior as the first part of Phase 3 with a constant initial gap in Fig. 4d. Now this continues the whole phase, with the gap in counts oscillating about 0 until finally reaching 0 when all biased agents have exponent = -23 at time S2. With no exponent > -23, we stabilize to output T in Phase 4.



(b) Snapshot S0. We enter Phase 3 with an equal number of bias = ±1.



(c) Snapshot S1. We see similar distributions as in Fig. 5b all the way until the end of Phase 3.



(d) Snapshot S2. After reaching synchronous hour 23, split reactions bring all remaining biased agents down to exponent = -23, which is only possible with an initial tie.

Figure 7: The case of a tie, with initial gap $g = 0$, again with $n = 5122666 \approx 2^{23}$, $p = 0.1$, $k = 2$. The simulation converges in Phase 2. Fig. 6a shows the counts of each opinion among the population of Main agents. Three special times S0, S1, S2, and the configurations of biased agents at these snapshots are shown in Figures 7b, 7c, 7d. See the link to animations of these plots at the end of this section.

subpopulations are near their expected fractions, both with high probability.

It is likely that we could use a simpler population splitting scheme. For example, we could simply have each pair of agents with initially opposite opinions change to roles `Clock` and `Reserve`. However, this would mean the number of agents in the role `Main` after `Phase 0` would depend on the initial gap. The current method of population splitting gives us stronger guarantees on the number of each agent in each role, simplifying subsequent analysis.

Phase 0 Initialize Roles. Agent u interacting with agent v .

Init $\text{role} \leftarrow \text{Role}_{\text{MCR}} \in \{\text{Main}, \text{Clock}, \text{Reserve}, \text{Role}_{\text{MCR}}, \text{Role}_{\text{CR}}\}$

$\text{assigned} \leftarrow \text{False} \in \{\text{True}, \text{False}\}$

if $\text{input} = \text{A}$, $\text{bias} \leftarrow +1 \in \{-3, -2, -1, 0, +1, +2, +3\}$

if $\text{input} = \text{B}$, $\text{bias} \leftarrow -1 \in \{-3, -2, -1, 0, +1, +2, +3\}$

we always maintain the invariant $\text{opinion} = \text{sign}(\text{bias}) \in \{-1, 0, +1\}$

if $\text{role} = \text{Clock}$, $\text{counter} \leftarrow c_0 \ln n \in \{0, \dots, c_0 \ln n\}$ only used in the current phase

```

1: if  $u.\text{role} = v.\text{role} = \text{Role}_{\text{MCR}}$  then ▷ Allocate  $\approx \frac{1}{2}$ Main agents
2:    $u.\text{role} \leftarrow \text{Main}; u.\text{bias} \leftarrow u.\text{bias} + v.\text{bias}$ 
3:    $v.\text{bias} \leftarrow 0; v.\text{role} \leftarrow \text{Role}_{\text{CR}}$  ▷  $v$  won't use its bias subsequently
4: if  $i.\text{role} = \text{Role}_{\text{MCR}}, j.\text{role} = \text{Main}, j.\text{assigned} = \text{False}$  where  $\{i, j\} = \{u, v\}$  then
5:    $j.\text{assigned} \leftarrow \text{True}; j.\text{bias} \leftarrow j.\text{bias} + i.\text{bias}$  ▷ Main agents can assign 1 non-Main agent
6:    $i.\text{bias} \leftarrow 0; i.\text{role} \leftarrow \text{Role}_{\text{CR}}$  ▷  $i$  won't use its bias subsequently
7: if  $i.\text{role} = \text{Role}_{\text{MCR}}, j.\text{role} \neq \text{Main}, \text{Role}_{\text{MCR}}, j.\text{assigned} = \text{False}$  where  $\{i, j\} = \{u, v\}$  then
8:    $j.\text{assigned} \leftarrow \text{True}$  ▷ non-Main agents can assign 1 Main agent
9:    $i.\text{role} \leftarrow \text{Main}$ 
10: if  $u.\text{role} = v.\text{role} = \text{Role}_{\text{CR}}$  then ▷ Allocate  $\approx \frac{1}{4}$ Clock agents,  $\approx \frac{1}{4}$ Reserve agents
11:    $u.\text{role} \leftarrow \text{Clock}; u.\text{counter} \leftarrow \Theta(\log n)$ 
12:    $v.\text{role} \leftarrow \text{Reserve}$ 
13: if  $u.\text{role} = v.\text{role} = \text{Clock}$  then ▷ time the phase once we have at least 2 Clock agents
14:   execute Standard Counter Subroutine( $u$ ), execute Standard Counter Subroutine( $v$ )

```

Standard Counter Subroutine(c). Agent c with field counter.

1: $c.\text{counter} \leftarrow c.\text{counter} - 1$

2: if $c.\text{counter} = 0$ then

3: $c.\text{phase} \leftarrow c.\text{phase} + 1$

▷ move to next phase

`Phase 1` is a timed phase that averages the biases in `Main` agents; with high probability at the end of the phase the `bias` fields have three consecutive values, shown in [Lemma 5.3](#). `Phase 1` expects no agent to remain in role `RoleMCR` by this point, signaling an error otherwise.¹⁴

¹⁴`role = RoleMCR` is an error because we need all agents not in role `Main` to have “donated” their bias to a `Main` agent. It is okay for `RoleCR` agents to be undecided about `Clock` versus `Reserve`. All such agents because `Reserve`, which WHP leaves sufficiently many agents in role `Clock` to make the clock interactions sufficiently fast.

Phase 1 Discrete Averaging. Agent u interacting with agent v .

Init if $\text{role} = \text{Role}_{\text{MCR}}$, $\text{phase} \leftarrow 10$ (error, skip to stable backup)

if $\text{role} = \text{Role}_{\text{CR}}$, $\text{role} \leftarrow \text{Reserve}$

if $\text{role} = \text{Clock}$, $\text{counter} \leftarrow c_1 \ln n \in \{0, \dots, c_1 \ln n\}$ only used in the current phase

- 1: **if** $u.\text{role} = v.\text{role} = \text{Main}$ **then**
 - 2: $u.\text{bias} \leftarrow \lfloor \frac{u.\text{bias} + v.\text{bias}}{2} \rfloor$; $v.\text{bias} \leftarrow \lceil \frac{u.\text{bias} + v.\text{bias}}{2} \rceil$
 - 3: **for** $c \in \{u, v\}$ with $c.\text{role} = \text{Clock}$ **do**
 - 4: execute **Standard Counter Subroutine**(c)
-

Phase 2 (an untimed phase) checks to see if the entire minority population was eliminated in **Phase 1** by checking whether both positive and negative biases still exist. It assumes a starting condition where all $\text{bias} \in \{-1, 0, +1\}$ (like the initial condition, but allowing some cancelling to have already happened). So the **Init** checks if any $|\text{bias}| > 1$, which can only happen with low probability (since **Phase 1** WHP reaches the three consecutive values $\{-1, 0, +1\}$), and we consider this an error and simply proceed immediately to **Phase 10**. If not, the minority opinion is gone, and the protocol will stabilize here to the correct output. Otherwise, we proceed to the next phase; note this phase is untimed and proceeds immediately upon detection of conflicting opinions. **Lemma 5.3** shows that starting from a large initial gap, we will stabilize here with high probability.

Phase 2 Output the Consensus. Agent u interacting with agent v .

Init if $|\text{bias}| > 1$, $\text{phase} \leftarrow 10$ (error, skip to stable backup)

$\text{opinions} \leftarrow \{\text{opinion}\} \subseteq \{-1, 0, +1\}$

- 1: $u.\text{opinions}, v.\text{opinions} \leftarrow u.\text{opinions} \cup v.\text{opinions}$ ▷ union opinions
 - 2: **if** $\{-1, +1\} \subseteq \text{opinions}$ **then**
 - 3: $u.\text{phase}, v.\text{phase} \leftarrow u.\text{phase} + 1$ ▷ no consensus, move to next phase
 - 4: **else if** $+1 \in \text{opinions}$ **then**
 - 5: $u.\text{output}, v.\text{output} \leftarrow A$ ▷ current consensus is A
 - 6: **else if** $-1 \in \text{opinions}$ **then**
 - 7: $u.\text{output}, v.\text{output} \leftarrow B$ ▷ current consensus is B
 - 8: **else if** $\text{opinions} = \{0\}$ **then**
 - 9: $u.\text{output}, v.\text{output} \leftarrow T$ ▷ current consensus is T
-

Phase 3 is where the bulk of the work gets done. The biased agents (with non-zero **opinion**) have an additional field **exponent** $\in \{-L, \dots, 0\}$, initially 0, where $L = \lceil \log_2(n) \rceil$, corresponding to holding 2^{exponent} units of mass. The unbiased \mathcal{O} agents have an additional field **hour** $= L \in \{0, \dots, L\}$, and will only participate in split reactions with $-\text{exponent} > \text{hour}$. The field **hour** is set by the **Clock** agents, who have a field **minute** $= 0 \in \{0, \dots, kL\}$, which counts up as **Phase 3** proceeds. Intuitively, there are k minutes in an hour, so $\text{hour} = \lceil \frac{\text{minute}}{k} \rceil$ ranges from 0 to L .¹⁵ **Fig. 3** shows how nonconsecutive minutes in the **Clock** agents may overlap significantly, but non-consecutive hours in the \mathcal{O} agents have negligible overlap. Once a **Clock** agent has **minute** at its maximum value kL , they initialize a new field **counter** $= \Theta(\log n) \in \{0, \dots, \Theta(\log n)\}$ to wait for the end of the phase (waiting for any remaining \mathcal{O} agents to reach **hour** $= L$ and the distribution to settle).

See the overview in **Section 3.2** for an intuitive description of this phase. **Theorem 6.1** gives the main result of **Phase 3** in the case of an initial tie, where all remaining biased agents are at the minimal **exponent** $= -L$. In the other case, **Theorem 6.2** gives the main result that most of the population settle on the majority output with **exponent** $\in \{-l, -(l+1), -(l+2)\}$.

¹⁵The drip reaction $C_i, C_i \rightarrow C_i, C_{i+1}$ implemented by line 5 is asymmetric, but could be made symmetric, i.e., $C_i, C_i \rightarrow C_{i+1}, C_{i+1}$ without affecting the analysis meaningfully.

Phase 3 Synchronized Rational Averaging. Agent u interacting with agent v .

Init if $\text{role} = \text{Main}$ and $\text{opinion} \in \{-1, +1\}$, $\text{exponent} \leftarrow 0 \in \{-L, \dots, -1, 0\}$, and we define $\text{bias} = \text{opinion} \cdot 2^{\text{exponent}}$

if $\text{role} = \text{Main}$ and $\text{opinion} = 0$, $\text{hour} \leftarrow 0 \in \{0, \dots, L\}$

if $\text{role} = \text{Clock}$, $\text{minute} \leftarrow 0 \in \{0, \dots, kL\}$ and $\text{counter} \leftarrow c_3 \ln n \in \{0, \dots, c_3 \ln n\}$

```

1: if  $u.\text{role} = v.\text{role} = \text{Clock}$  then
2:   if  $u.\text{minute} \neq v.\text{minute}$  then
3:      $u.\text{minute}, v.\text{minute} \leftarrow \max(u.\text{minute}, v.\text{minute})$  ▷ clock epidemic reaction
4:   else if  $u.\text{minute} < kL$  then
5:      $u.\text{minute} \leftarrow u.\text{minute} + 1$  with probability  $p$  ▷ clock drip reaction,  $p = 1$  in Theorem 6.9
6:   else ▷ count only when both clocks finished
7:     execute Standard Counter Subroutine( $u$ ), Standard Counter Subroutine( $v$ )
8: if  $m.\text{role} = \text{Main}$ ,  $m.\text{opinion} = 0$  and  $c.\text{role} = \text{Clock}$  where  $\{m, c\} = \{u, v\}$  then
9:    $m.\text{hour} \leftarrow \max(m.\text{hour}, \lfloor \frac{c.\text{minute}}{k} \rfloor)$  ▷ clock update reaction
10: else if  $u.\text{role} = v.\text{role} = \text{Main}$  then
11:   if  $\{u.\text{opinion}, v.\text{opinion}\} = \{-1, +1\}$  and  $u.\text{exponent} = v.\text{exponent} = -h$  then
12:      $u.\text{opinion}, v.\text{opinion} \leftarrow 0$ ;  $u.\text{hour}, v.\text{hour} \leftarrow h$  ▷ cancel reaction
13:   if  $t.\text{opinion} = 0$ ,  $i.\text{opinion} \in \{-1, +1\}$  and  $|t.\text{hour}| > |i.\text{exponent}|$ , where  $\{t, i\} = \{u, v\}$  then
14:      $t.\text{opinion} \leftarrow i.\text{opinion}$  ▷ split reaction
15:      $i.\text{exponent}, t.\text{exponent} \leftarrow i.\text{exponent} - 1$  ▷ sets  $i.\text{bias}, t.\text{bias} \leftarrow i.\text{bias}/2$ 

```

In [Phase 4](#) (an untimed phase), the population checks if all Main agents have reached the minimum $\text{exponent} = -L$, which only happens in the case of a tie. If so, the population will stabilize to the tie output. Otherwise, any Main agent with exponent above L can trigger the move to the next phase.

Phase 4 Output Tie. Agent u interacting with agent v .

Init $\text{output} \leftarrow \top$

```

1: if  $|m.\text{bias}| > 2^{-L}$  where  $m \in \{u, v\}$  then ▷ stable if all  $\text{bias} \in \{-\frac{1}{2^L}, 0, +\frac{1}{2^L}\}$ 
2:    $u.\text{phase}, v.\text{phase} \leftarrow u.\text{phase} + 1$  ▷ end of this phase

```

If there was not a tie detected in [Phase 4](#), then most agents should have the majority opinion, with $\text{exponent} \in \{-l, -(l+1), -(l+2)\}$ in a small range. The next goal is to bring all agents with $\text{exponent} > -l$ down to $\text{exponent} \leq -l$. This will be accomplished by the Reserve agents, whose goal is to let exponents above l do split reactions.

The Reserve agents do this across two consecutive phases. In [Phase 5](#), the Reserve agents become active, and will set $\text{sample} = \perp \in \{\perp, 0, \dots, L\}$ to the exponent of the first biased agent they meet, which is likely in $\{-l, -(l+1), -(l+2)\}$. The Clock agents now only hold a field $\text{counter} = \Theta(\log n) \in \{0, \dots, \Theta(\log n)\}$ to act as a simple timer for how long to wait until moving to the next phase. This allows the Reserve agents to adopt a distribution of exponent values approximately equal to that of the Main agents. This behavior is proven in [Lemma 7.1](#) and [Lemma 7.2](#).

Phase 5 Reserves Sample exponent. Agent u interacting with agent v .

Init if $\text{role} = \text{Reserve}$, $\text{sample} \leftarrow \perp \in \{\perp, -L, \dots, -1, 0\}$

if $\text{role} = \text{Clock}$, $\text{counter} \leftarrow c_5 \ln n \in \{0, \dots, c_5 \ln n\}$

```

1: if  $r.\text{role} = \text{Reserve}$  and  $m.\text{role} = \text{Main}$ ,  $m.\text{opinion} \in \{-1, +1\}$  where  $\{r, m\} = \{u, v\}$  then
2:   if  $r.\text{sample} = \perp$  then ▷ sample exponent of biased agent  $m$ 
3:      $r.\text{sample} \leftarrow m.\text{exponent}$ 
4: for  $c \in \{u, v\}$  with  $c.\text{role} = \text{Clock}$  do
5:   execute Standard Counter Subroutine( $c$ )

```

In **Phase 6**, the Reserve agents can help facilitate more split reactions, with any agent at an exponent above their sampled exponent. Because they have approximately the same distribution across all exponents as Main agents, particular exponents $-l, -(l+1), -(l+2)$, this allows them to bring all agents above $\text{exponent} = -l$ down to $-l$ or below. Again, the Clock agents keep a $\text{counter} \in \{0, \dots, c_6 \ln n\}$. [Lemma 7.2](#) proves that **Phase 6** works as intended, bringing all agents down to $\text{exponent} \leq -l$ with high probability.

Phase 6 Reserve Splits. Agent u interacting with agent v .

Init if $\text{role} = \text{Clock}$, $\text{counter} \leftarrow c_6 \ln n \in \{0, \dots, c_6 \ln n\}$

- 1: **if** $r.\text{role} = \text{Reserve}$ and $m.\text{role} = \text{Main}$, $m.\text{opinion} \in \{-1, +1\}$ where $\{r, m\} = \{u, v\}$ **then**
 - 2: **if** $r.\text{sample} \neq \perp$ and $r.\text{sample} < m.\text{exponent}$ **then**
 - 3: $r.\text{role} \leftarrow \text{Main}$; $r.\text{opinion} \leftarrow m.\text{opinion}$ ▷ split reaction
 - 4: $r.\text{exponent}, m.\text{exponent} \leftarrow m.\text{exponent} - 1$ ▷ sets $r.\text{bias}, m.\text{bias} \leftarrow m.\text{bias}/2$
 - 5: **for** $c \in \{u, v\}$ with $c.\text{role} = \text{Clock}$ **do**
 - 6: execute **Standard Counter Subroutine**(c)
-

Now that all agents are $\text{exponent} \leq -l$, the goal of **Phase 7** is to eliminate any minority agents with exponents $-l, -(l+1), -(l+2)$. This is done by letting the biased agents do generalized cancel reactions that allow their difference in exponents to be up to 2, while still preserving the mass invariant. Again, the Clock agents keep a $\text{counter} \in \{0, \dots, c_7 \ln n\}$. [Lemma 7.5](#) shows that by the end of this phase, any remaining minority agents must have $\text{exponent} < -(l+2)$, with high probability.

Phase 7 High-exponent Minority Elimination. Agent u interacting with agent v

Init if $\text{role} = \text{Clock}$, $\text{counter} \leftarrow c_7 \ln n \in \{0, \dots, c_7 \ln n\}$

- 1: **if** $u.\text{role} = v.\text{role} = \text{Main}$ and $\{u.\text{opinion}, v.\text{opinion}\} = \{-1, +1\}$ **then**
 - 2: **if** $u.\text{exponent} = v.\text{exponent}$ **then**
 - 3: $u.\text{opinion}, v.\text{opinion} \leftarrow 0$ ▷ cancel reaction
 - 4: **else if** $i.\text{exponent} = j.\text{exponent} + 1$ where $\{i, j\} = \{u, v\}$ **then**
 - 5: $i.\text{exponent} \leftarrow i.\text{exponent} - 1$ ▷ gap-1 cancel reaction
 - 6: $j.\text{opinion} \leftarrow 0$ ▷ example bias update: $+\frac{1}{4}, -\frac{1}{8} \rightarrow +\frac{1}{8}, 0$
 - 7: **else if** $i.\text{exponent} = j.\text{exponent} + 2$ where $\{i, j\} = \{u, v\}$ **then**
 - 8: $j.\text{opinion} \leftarrow i.\text{opinion}$ ▷ gap-2 cancel reaction
 - 9: $i.\text{exponent} \leftarrow i.\text{exponent} - 1$ ▷ example bias update: $+\frac{1}{4}, -\frac{1}{16} \rightarrow +\frac{1}{8}, +\frac{1}{16}$
 - 10: $j.\text{exponent} \leftarrow i.\text{exponent} - 2$
 - 11: **for** $c \in \{u, v\}$ with $c.\text{role} = \text{Clock}$ **do**
 - 12: execute **Standard Counter Subroutine**(c)
-

Now that all minority agents occupy exponents below $-(l+2)$, yet a large number of majority agents remain at exponents $-l, -(l+1), -(l+2)$, in **Phase 8**, the algorithm eliminates the last remaining minority opinions at any exponent. It allows opposite-opinion agents of *any* two exponents to react and eliminate the smaller-exponent opinion. The larger exponent $\frac{1}{2^i}$ “absorbs” the smaller $-\frac{1}{2^j}$, setting the smaller to mass 0; the larger now represents mass $\frac{1}{2^i} - \frac{1}{2^j}$, which it lacks the memory to track exactly, so it cannot absorb any further agents (though it can itself be absorbed by $-\frac{1}{2^m}$ for $m > i$). [Lemma 7.6](#) shows that **Phase 8** eliminates any remaining minority agents with high probability. A key property is that it cannot violate correctness since, although the biases held by some agents becomes unknown, the allowed transitions maintain the sign of the bias. Thus the only source of error is failing to eliminate all minority agents, detected in the next phase.

Phase 8 Low-exponent Minority Elimination. Agent u interacting with agent v

Init if $\text{role} = \text{Clock}$, $\text{counter} \leftarrow c_8 \ln n \in \{0, \dots, c_8 \ln n\}$

```

1: if  $u.\text{role} = v.\text{role} = \text{Main}$  and  $\{u.\text{bias}, v.\text{bias}\} = \{-1, +1\}$  then
2:   if  $i.\text{exponent} > j.\text{exponent}$  and  $i.\text{full} = \text{False}$  where  $\{i, j\} = \{u, v\}$  then
3:      $i.\text{full} \leftarrow \text{True}$  ▷ consumption reaction
4:      $j.\text{opinion} \leftarrow 0$ 
5: for  $c \in \{u, v\}$  with  $c.\text{role} = \text{Clock}$  do
6:   execute Standard Counter Subroutine( $c$ )

```

Phase 9 (an untimed phase) acts exactly as **Phase 2**, to check that agents have reached consensus. Note that the initial check for $|\text{bias}| > 1$ is not required, since it is guaranteed to pass if we reach this point: it passed in **Phase 2**, and the population-wide maximum $|\text{bias}|$ could only have decreased in subsequent phases.

Phase 9 Output the Consensus. Exact repeat of **Phase 2**.

In **Phase 10**, the agents give up on the fast algorithm, having determined in **Phase 9** that it failed to reach consensus, or detected an earlier error with role or bias assignment. Instead they rely instead on a slow stable backup protocol. This is a 6-state protocol that stably decides between the three cases of majority A, B, and tie T. They only use the fields $\text{output} = \text{input} \in \{A, B, T\}$, and the initial field $\text{active} = \text{True} \in \{\text{True}, \text{False}\}$. [Lemma 7.7](#) proves this 6-state protocol stably computes majority in $O(n \log n)$ time.

Phase 10 Stable Backup. Agent u interacting with agent v . Similar to 6-state algorithm from [22], slight modification of 4-state algorithm from [33, 41].

Init $\text{output} \leftarrow \text{input}$, $\text{active} \leftarrow \text{True}$

```

1: if  $u.\text{active} = v.\text{active} = \text{True}$  then
2:   if  $\{u.\text{output}, v.\text{output}\} = \{A, B\}$  then
3:      $u.\text{output}, v.\text{output} \leftarrow T$  ▷ cancel reaction
4:   else if  $i.\text{output} \in \{A, B\}$  and  $t.\text{output} = T$  where  $\{i, t\} = \{u, v\}$  then
5:      $t.\text{output} \leftarrow i.\text{output}$ ,  $t.\text{active} \leftarrow \text{False}$  ▷ biased converts unbiased
6: if  $a.\text{active} = \text{True}$  and  $p.\text{active} = \text{False}$  where  $\{a, p\} = \{u, v\}$  then
7:    $p.\text{output} \leftarrow a.\text{output}$  ▷ active converts passive

```

4 Useful time bounds

This section introduces various probability bounds which will be used repeatedly in later analysis.

We will use the following standard multiplicative Chernoff bound:

Theorem 4.1. *Let $X = X_1 + \dots + X_k$ be the sum of independent $\{0, 1\}$ -valued random variables, with $\mu = \mathbb{E}[X]$. Then for any $\delta > 0$, $\mathbb{P}[X \geq (1 + \delta)\mu] \leq \exp\left(-\frac{\delta^2 \mu}{2 + \delta}\right)$.*

We use the standard Azuma inequality for supermartingales:

Theorem 4.2. *Let X_0, X_1, X_2, \dots be a supermartingale such that, for all $i \in \mathbb{N}$, $|X_{i+1} - X_i| \leq c_i$. Then for all $n \in \mathbb{N}$ and $\epsilon > 0$, $\mathbb{P}[(X_n - X_0) - \mathbb{E}[X_n - X_0] \geq \epsilon] \leq \exp\left(-\frac{2\epsilon^2}{\sum_{i=0}^n c_i^2}\right)$.*

In application we will consider potential functions ϕ that decay exponentially, with $\mathbb{E}[\phi_{j+1}] \leq (1 - \epsilon)\phi_j$. In this case, we will take the logarithm $\Phi = \ln(\phi)$, which we will show is a supermartingale upon which we can apply Azuma's Inequality to conclude that ϕ achieves a requisite amount of exponential decay.

We will also use the following two concentration bounds on heterogeneous sums of geometric random variables, due to Janson [39, Theorems 2.1, 3.1]:

Theorem 4.3. *Let X be sum of k independent geometric random variables with success probabilities p_1, \dots, p_k , let $\mu = \mathbb{E}[X] = \sum_{i=1}^k \frac{1}{p_i}$, and let $p^* = \min_{1 \leq i \leq k} p_i$. For all $\lambda \geq 1$, $\mathbb{P}[X \geq \lambda\mu] \leq e^{-p^*\mu(\lambda-1-\ln\lambda)}$. For all $\lambda \leq 1$, $\mathbb{P}[X \leq \lambda\mu] \leq e^{-p^*\mu(\lambda-1-\ln\lambda)}$.*

The expression $\lambda - 1 - \ln \lambda$ is hard to work with if we are not fixing exact constant values for λ . The following Corollary gives asymptotic approximation to the error bound:

Corollary 4.4. *Let X be sum of k independent geometric random variables with success probabilities p_1, \dots, p_k , let $\mu = \mathbb{E}[X] = \sum_{i=1}^k \frac{1}{p_i}$, and let $p^* = \min_{1 \leq i \leq k} p_i$. For any $0 < \epsilon < 1$, we have $(1-\epsilon)\mu \leq X \leq (1+\epsilon)\mu$ with probability $1 - \exp(-\Theta(\epsilon^2 p^* \mu))$.*

Proof. Setting $\lambda = 1 + a$ in [Theorem 4.3](#), where $a = \epsilon$ in the upper bound and $a = -\epsilon$ in the lower bound, by Taylor series approximation of $\ln(1 + a)$, $\lambda - 1 - \ln(\lambda) = a - \ln(1 + a) \geq a - (a - \frac{a^2}{2} + \frac{a^3}{3}) = a^2(\frac{1}{2} - \frac{a}{3}) = \Theta(\epsilon^2)$. The stated inequality then follows from [Theorem 4.3](#). \square

The following lemmas give applications of [Theorem 4.3](#) and [Corollary 4.4](#) to common processes that we will repeatedly analyze. When the fractions we consider are distinct and independent of n , [Corollary 4.4](#) applies to give tight time bounds with very high probability. We also consider cases where we run a process to completion and bring a count to 0. Here we must use [Theorem 4.3](#) and only get a high probability bound with times at most a constant factor above the mean.

First we consider the epidemic process, $i, s \rightarrow i, i$, moving from a constant fraction infected to another constant fraction infected.

Lemma 4.5. *Let $0 < a < b < 1$. Consider the epidemic process starting from a count of $a \cdot n$ infected agents. The expected parallel time t until there is a count $b \cdot n$ of infected agents is*

$$\mathbb{E}[t] = \frac{\ln(b - \frac{1}{n}) - \ln(1 - b + \frac{1}{n}) - \ln(a) + \ln(1 - a)}{2} \sim \frac{\ln(b) - \ln(1 - b) - \ln(a) + \ln(1 - a)}{2}.$$

Let $c = \min(a, 1 - b + \frac{1}{n})$ and $0 < \epsilon < 1$. Then $(1 - \epsilon)\mathbb{E}[t] < t < (1 + \epsilon)\mathbb{E}[t]$ with probability at least $1 - \exp[-\Theta(\epsilon^2 \mathbb{E}[t]nc)]$.

Proof. When there are i infected agents, the probability the next reaction infects another agent and increases this count is $\frac{i(n-i)}{\binom{n}{2}}$. The number of interactions T for the count to increase from $a \cdot n$ to $b \cdot n$ is a sum of geometric random variables, with expected value

$$\begin{aligned} \mathbb{E}[T] &= \sum_{i=a \cdot n}^{b \cdot n - 1} \frac{\binom{n}{2}}{i(n-i)} \sim \frac{1}{2} \sum_{i=a \cdot n}^{b \cdot n - 1} \frac{1}{\frac{i}{n}(1 - \frac{i}{n})} \sim \frac{n}{2} \int_{x=a}^{b - \frac{1}{n}} \frac{dx}{x(1-x)} \\ &= \frac{n}{2} [\ln(x) - \ln(1-x)]_a^{b - \frac{1}{n}} = n \cdot \frac{\ln(b - \frac{1}{n}) - \ln(1 - b + \frac{1}{n}) - \ln(a) + \ln(1 - a)}{2}, \end{aligned}$$

giving the stated value of $\mathbb{E}[t]$ after converting from interactions to parallel time. The minimum probability $p^* = \Theta(\min(a, 1 - b + \frac{1}{n}))$, so using [Corollary 4.4](#) we have Then $(1 - \epsilon)\mathbb{E}[t] < t < (1 + \epsilon)\mathbb{E}[t]$ with probability at least $1 - \exp[-\Theta(\epsilon^2 \mathbb{E}[T]p^*)]$. \square

Note that if $a, (1 - b), \epsilon$ are all constants independent of n , then the bound above is with very high probability. If we wanted to consider the complete epidemic process (which starts with $a = 1/n$ and ends with $(1 - b) = 1/n$), we would have $\mathbb{E}[t] = \Theta(\log n)$ and minimum probability $p^* = \Theta(\frac{1}{n})$. Then the probability bound would become $1 - \exp[-\Theta(\log n)] = 1 - n^{-\Theta(1)}$, which is high probability, but requires carefully considering the constants to get the exact polynomial bound. In our proofs, we only use the very high probability case. For tight bounds on a full epidemic with precise constants, see [24, 42].

Next we consider the cancel reactions $a, b \rightarrow 0, 0$, which are key to the majority protocol.

Lemma 4.6. Consider two disjoint subpopulations A and B of initial sizes $|A| = a \cdot n$ and $|B| = b \cdot n$, where $0 < b < a < 1$. An interaction between an agent in A and an agent in B is a cancel reaction which removes both agents from their subpopulations. Thus after i cancel reactions we have $|A| = a \cdot n - i$ and $|B| = b \cdot n - i$. The expected parallel time t until $d \cdot n$ cancel reactions occur, where $d < b$ is

$$\mathbb{E}[t] = \frac{\ln(b) - \ln(a) - \ln(b - d + \frac{1}{n}) + \ln(a - d + \frac{1}{n})}{2(a - b)}.$$

Let $c = (a - d + \frac{1}{n})(b - d + \frac{1}{n})$ and $0 < \epsilon < 1$. Then $(1 - \epsilon)\mathbb{E}[t] < t < (1 + \epsilon)\mathbb{E}[t]$ with probability at least $1 - \exp[-\Theta(\epsilon^2 \mathbb{E}[t]nc)]$.

The parallel time t until all $b \cdot n$ cancel reactions occur has $\mathbb{E}[t] \sim \frac{\ln n}{2(a-b)}$ and satisfies $t \leq \frac{5 \ln n}{2(a-b)}$ with high probability $1 - O(1/n^2)$.

Again, the first bound is with very high probability if all constants a, b, d are independent of n .

Proof. After i cancel reactions, the probability of the next cancel reaction is $p \sim \frac{2|A||B|}{n^2} = 2(a - \frac{i}{n})(b - \frac{i}{n})$, and the number of interactions until this cancel reaction is a geometric random variable with mean p . The number of interactions T for $d \cdot n$ cancel reactions to occur is a sum of geometrics with mean

$$\begin{aligned} \mathbb{E}[T] &= \sum_{i=0}^{d \cdot n - 1} \frac{1}{2(a - \frac{i}{n})(b - \frac{i}{n})} \sim n \int_{x=0}^{d - \frac{1}{n}} \frac{dx}{2(a-x)(b-x)} = \frac{n}{2(a-b)} \int_{x=0}^{d - \frac{1}{n}} \left(\frac{1}{b-x} - \frac{1}{a-x} \right) dx \\ &= \frac{n}{2(a-b)} \left[-\ln(b-x) + \ln(a-x) \right]_0^{d - \frac{1}{n}} = n \cdot \frac{\ln(b) - \ln(a) - \ln(b - d + \frac{1}{n}) + \ln(a - d + \frac{1}{n})}{2(a-b)}. \end{aligned}$$

Translating to parallel time and using [Corollary 4.4](#) gives the first result, where minimum probability $p^* = \Theta(c)$.

In the second case where $d = b$ and we are waiting for all cancel reactions to occur, then

$$\mathbb{E}[t] \sim \frac{\ln(b) - \ln(a) + \ln(n) + \ln(a-b)}{2(a-b)} = \frac{\ln n}{2(a-b)},$$

and $\mu = \mathbb{E}[T] \sim \frac{n \ln n}{2(a-b)}$. Now the minimum geometric probability is when $i = bn - 1$ and $p^* \sim \frac{2(a-b)}{n}$. Choosing $\lambda = 5$ so that $\lambda - 1 - \ln \lambda > 2$, by [Theorem 4.3](#) we have

$$\mathbb{P} \left[t \geq \frac{5 \ln n}{2(a-b)} \right] \leq \mathbb{P}[T \geq \lambda \mu] \leq e^{-p^* \mu (\lambda - 1 - \ln \lambda)} = \exp \left(-\frac{2(a-b)}{n} \cdot \frac{n \ln n}{2(a-b)} \cdot 2 \right) = n^{-2}.$$

Thus $t \leq \frac{5 \ln n}{2a}$ with high probability.

Note we get the same result for this second case also by assuming a fixed minimal fraction $a - b$ in $|A|$ is always there to cancel the agents from $|B|$, and using the next [Lemma 4.7](#). \square

Now we consider a ‘‘one-sided’’ cancel process, where reactions $a, b \rightarrow a, 0$ change only the b agents into a different state. This process happens for example when Clock agents change the hour of \mathcal{O} agents.

Lemma 4.7. Let $0 < a, b_1, b_2, \epsilon < 1$ be constants with $b_1 > b_2$. Consider a subpopulation A maintaining its size above $a \cdot n$, and B initially of size $b_1 \cdot n$. Any interaction between an agent in A and in B is meaningful, and forces the agent in B to leave its subpopulation.

The expected parallel time t until the subpopulation B reaches size $b_2 \cdot n$ is

$$\mathbb{E}[t] = \frac{\ln(b_1) - \ln(b_2)}{2a},$$

and satisfies $(1 - \epsilon)\mathbb{E}[t] < t < (1 + \epsilon)\mathbb{E}[t]$ with probability at least $1 - \exp[-\Theta(\epsilon^2 \mathbb{E}[t]nab_2)]$.

The parallel time t until the subpopulation B reaches size 0 has $\mathbb{E}[t] \sim \frac{\ln n}{2a}$ and satisfies $t \leq \frac{5 \ln n}{2a}$ with high probability $1 - O(1/n^2)$.

Again, the first bound is with very high probability if constants a, b_2 are independent of n .

Proof. When $|B| = i$, then the probability of a meaningful interaction $p \sim \frac{2ia}{n}$. Then the number of interactions before the next meaningful interaction is a geometric with probability p , and the total number T of interaction is a sum of geometrics with mean

$$\mathbb{E}[T] = \sum_{i=b_2n+1}^{b_1n} \frac{n}{2ia} = \frac{n}{2a} \left(\sum_{i=1}^{b_1n} \frac{1}{i} - \sum_{i=1}^{b_2n+1} \frac{1}{i} \right) \sim \frac{n}{2a} (\ln(b_1n) - \ln(b_2n+1)) \sim n \frac{\ln(b_1) - \ln(b_2 + \frac{1}{n})}{2a}.$$

Translating to parallel time and using [Corollary 4.4](#) gives the first result, where minimum probability $p^* = \Theta(a \cdot b_2)$.

In the case where $b_2 = 0$, we have $\mu = \mathbb{E}[T] \sim \frac{n \ln n}{2a}$ and $\mathbb{E}[t] \sim \frac{\ln(n)}{2a}$. Now the minimum geometric probability $p^* = \frac{2a}{n}$. Choosing $\lambda = 5$ so that $\lambda - 1 - \ln \lambda > 2$, by [Theorem 4.3](#) we have

$$\mathbb{P}\left[t \geq \frac{5 \ln n}{2a}\right] \leq \mathbb{P}[T \geq \lambda \mu] \leq e^{-p^* \mu (\lambda - 1 - \ln \lambda)} = \exp\left(-\frac{2a}{n} \cdot \frac{n \ln n}{2a} \cdot 2\right) = n^{-2}.$$

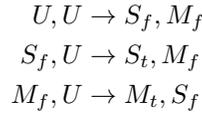
Thus $t \leq \frac{5 \ln n}{2a}$ with high probability. \square

5 Analysis of initial phases

Define \mathcal{M} , \mathcal{C} , and \mathcal{R} to be the sub-populations of agents with roles Main, Clock, Reserve when they first set $\text{phase} = 1$. Let $|\mathcal{M}| = m \cdot n$, $|\mathcal{C}| = c \cdot n$, $|\mathcal{R}| = r \cdot n$, where $m + c + r = 1$.

The population splitting of [Phase 0](#) will set the fractions $m \approx \frac{1}{2}$, $c \approx \frac{1}{4}$, and $r \approx \frac{1}{4}$. The rules also ensure deterministic bounds on these fractions once all Role_{MCR} agents have been assigned. The probability 1 guarantees on subpopulation sizes using this method will be key for a later uniform adaptation of the protocol. We first prove the behavior of the rules used for this initial top level split in [Phase 0](#) between the roles Main and Role_{CR} .

Lemma 5.1. *Consider the reactions*



starting with n U agents. Let $u = \#U$, $s = \#S_f + \#S_t$, and $m = \#M_f + \#M_t$. This converges to $u = 0$ in expected time at most $2.5 \ln n$ and in $12.5 \ln n$ time with high probability $1 - O(1/n^2)$. Once $u = 0$, $\frac{n}{3} \leq s, m \leq \frac{2n}{3}$ with probability 1, and for any $\epsilon > 0$, $\frac{n}{2}(1 - \epsilon) \leq s, m \leq \frac{n}{2}(1 + \epsilon)$ with very high probability.

Proof. Let $s_f = \#S_f$, $s_t = \#S_t$, $m_f = \#M_f$, $m_t = \#M_t$.

First we consider the interactions that happen until $u = 2n/3$. Note that while $u \geq 2n/3$, the probability of the first reaction is $\sim \left(\frac{u}{n}\right)^2 \geq \frac{4}{9}$ and the probability of the other two reactions is $\sim 2\left(\frac{u}{n}\right)\left(\frac{m_f + s_f}{n}\right) \leq 2 \cdot \frac{2}{3} \cdot \frac{1}{3} = \frac{4}{9}$. Thus until $u = 2n/3$, each non-null interaction is the top reaction with at least probability $1/2$. Then by standard Chernoff Bounds, with very high probability, at least a fraction $\frac{1}{2} - \epsilon$ of these reactions are the top reaction, which create a count $s_f + m_f \geq \frac{2}{9}(1 - \epsilon)n > \frac{n}{5}$.

Now notice that this count $s_f + m_f$ can never decrease, so we have $s_f + m_f > \frac{n}{5}$ for all future interactions. Now we can use the rate of the second and third reactions to bound the completion time. The probability of decreasing u is at least $2\left(\frac{u}{n}\right)\left(\frac{1}{5}\right)$, so the number of interactions it takes to decrement u is stochastically dominated by a geometric random variable with probability $p = \frac{2u}{5n}$. Then the number of interactions for u to decrease from $\frac{2n}{3}$ down to 0 is dominated by a sum T of geometric random variables with mean

$$\mathbb{E}[T] = \sum_{u=1}^{2n/3} \frac{5n}{2u} = \frac{5n}{2} \sum_{u=1}^{2n/3} \frac{1}{u} \sim \frac{5n}{2} \ln(2n/3) \sim \frac{5}{2} n \ln n.$$

Now we will apply the upper bound of [Theorem 4.3](#), using $\mu = \frac{5}{2}n \ln n$, $p^* = \frac{2}{5n}$ (when $u = 1$) and $\lambda = 5$, where $\lambda - 1 - \ln(\lambda) > 2$, so we get

$$\mathbb{P}[T \geq \lambda\mu] \leq \exp(-p^*\mu(\lambda - 1 - \ln \lambda)) \leq \exp\left(-\frac{2}{5n} \cdot \frac{5}{2}n \ln n \cdot 2\right) = n^{-2}.$$

Thus with high probability $1 - O(1/n^2)$, this process converges in at most $\frac{\lambda\mu}{n} = 12.5 \ln n$ parallel time.

Observe that the reactions all preserve the following invariant: $s_f + 2s_t = m_f + 2m_t$. The probability-1 count bounds follow by observing that when $u = 0$ (i.e., we have converged) we have $s_f + s_t + m_f + m_t = n$. Maximizing $s = s_f + s_t$ and minimizing $m = m_f + m_t$ is achieved by setting $s_f = 2n/3$, $s_t = 0$, $m_f = 0$, $m_t = n/3$, and symmetrically for maximizing m .

Assume we have $s > m$ at some point in the execution, so $s_f + s_t > m_f + m_t$. Subtracting the invariant equation gives $-s_t > -m_t$, which implies $s_t < m_t$. Together with the first inequality this implies that $s_f > m_f$. Thus the rate of second reaction is higher than the rate of the third reaction, so it is more likely that the next reaction changing the value $s - m$ decreases it.

By symmetry the opposite happens when $m > s$, so we conclude that the absolute value $|m - s|$ is more likely to decrease than increase. Thus we can stochastically dominate $|m - s|$ by a sum of independent coin flips. The high probability count bounds follow by standard Chernoff bounds. \square

Now we can prove bounds on the sizes $|\mathcal{M}| = mn$, $|\mathcal{C}| = cn$, $|\mathcal{R}| = rn$ of Main, Clock, and Reserve agents from the population splitting of [Phase 0](#).

Lemma 5.2. *For any $\epsilon > 0$, with high probability $1 - O(1/n^2)$, by the end of [Phase 0](#), $|\text{Role}_{\text{MCR}}| = 0$, $\frac{n}{2}(1 - \epsilon) \leq |\mathcal{M}| \leq \frac{n}{2}(1 + \epsilon)$ and $|\mathcal{C}|, |\mathcal{R}| \geq \frac{n}{4}(1 - \epsilon)$. If $|\text{Role}_{\text{MCR}}| = 0$ and [Phase 1](#) initializes without error, then with probability 1, $\frac{1}{3}n \leq |\mathcal{M}| \leq \frac{2}{3}n$, $\frac{1}{6}n \leq |\mathcal{R}| \leq \frac{2}{3}n$, and $2 \leq |\mathcal{C}| \leq \frac{1}{3}n$.*

Proof. The top level of splitting of Role_{MCR} into Role_{CR} and Main is equivalent to the reactions of [Lemma 5.1](#), with $U = \text{Role}_{\text{MCR}}$, $M = \text{Main}$, $S = \text{Role}_{\text{CR}}$, and the f and t subscripts representing the Boolean value of the field `assigned`. [Lemma 5.1](#) gives the stated bounds on Main, if there were no further splitting of Role_{CR} .

[Lemma 5.1](#) gives that with high probability, all Role_{MCR} are converted to Role_{CR} and Main in $12.5 \ln n$ time; we begin the analysis at that point, letting s be the number of Role_{CR} agents produced, noting $n/3 \leq s \leq 2n/3$ with probability 1, and for any $\epsilon' \frac{n}{2}(1 - \epsilon') \leq s \leq \frac{n}{2}(1 + \epsilon')$ with high probability.

The splitting of Role_{CR} into Clock and Reserve follows a simpler process that we analyze here. Let $r = |\mathcal{R}|$ and $c = |\mathcal{C}|$ at the end of [Phase 0](#). To see the high probability bounds on r and c , we model the splitting of Role_{CR} by the reaction $U, U \rightarrow R, C$ during [Phase 0](#) and $U \rightarrow R$ at the end of [Phase 0](#), since all un-split Role_{CR} agents become Reserve agents upon leaving [Phase 0](#).

The reaction $U, U \rightarrow R, C$ reduces the count of U from its initial value s ($= n/2 \pm \epsilon'n/2$ WHP) to $\epsilon's$, with the number of interactions between each reaction when $\#U = l$ governed by a geometric random variable with success probability $O(l^2/n^2)$. Applying [Corollary 4.4](#) with $k = s - \epsilon's$, $p_i = O((i + \epsilon's)^2/n^2)$ for $i \in \{1, \dots, k\}$, the reaction $U, U \rightarrow R, C$ takes $O(1)$ time to reduce the count of U from its initial value m ($= n/2 \pm \epsilon'n/2$ WHP) to $\epsilon'm$ with very high probability. This implies that after $O(1)$ time, r and c are both at least

$$\begin{aligned} s/2 - \epsilon's &= s(1/2 - \epsilon') \geq (n/2 - \epsilon'n/2)(1/2 - \epsilon') \\ &= n/4 - \epsilon'n/2 - \epsilon'n/4 + (\epsilon')^2n/2 \\ &> n/4 - \epsilon'n = n/4(1 - 4\epsilon') \end{aligned}$$

with very high probability. Choosing $\epsilon = \epsilon'/4$ gives the high probability bounds on r and c . Thus we require $12.5 \ln n + O(1) \leq 13 \ln n$ time, and for appropriate choice of counter constant c_0 , this happens before the first Clock agent advances to the next phase with high probability.

We now argue the probability-1 bounds. The bound $r \leq 2n/3$ follows from $|\mathcal{M}| \geq n/3$. The bound $r \geq n/6$ follows from $|\mathcal{M}| \leq 2n/3$, so $\text{Role}_{\text{CR}} \geq n/3$ if no $U, U \rightarrow R, C$ splits happen, and the fact that at least half of Role_{CR} get converted to Reserve: exactly half by $U, U \rightarrow R, C$ and the rest by $U \rightarrow R$.

Although the reactions $U, U \rightarrow R, C$ and $U \rightarrow R$ can produce only a single C , there must be at least two Clock agents for [Standard Counter Subroutine](#) to count at all and end [Phase 0](#), so if [Phase 1](#) initializes, $c \geq 2$. The bound $c \leq n/3$ follows from the fact that c is maximized when $|\mathcal{M}| = n/3$ and no $U \rightarrow R$ reactions happen, i.e., all $2n/3$ Role_{CR} agents are converted via $U, U \rightarrow R, C$, leading to $c = n/3$. \square

We now reason about [Phase 1](#), which has different behavior based on the initial gap g .

Lemma 5.3. *If the initial gap $|g| \geq 0.025|\mathcal{M}|$, then we stabilize to the correct output in [Phase 2](#). If $|g| < 0.025|\mathcal{M}|$, then at the end of [Phase 1](#), all agents have $\text{bias} \in \{-1, 0, +1\}$, and the total count of biased agents is at most $0.03|\mathcal{M}|$. Both happen with high probability $1 - O(1/n^2)$.*

Proof. In the case where all agents enter [Phase 1](#), none are still in Role_{MCR} , so every non-Main agent has given their bias to a Main agent via [Line 2](#) or [Line 5](#) of [Phase 0](#). Thus the initial gap $g = \sum_{m.\text{role}=\text{Main}} m.\text{bias}$.

Let $\mu = \lfloor \frac{g}{|\mathcal{M}|} \rfloor$ be the average bias among all Main agents, rounded to the nearest integer. By [\[45\]](#), we will converge to have all $\text{bias} \in \{\mu - 1, \mu, \mu + 1\}$, in $O(\log n)$ time with high probability $1 - O(1/n^2)$. We use Corollary 1 of [\[45\]](#), where the constant $K = O(\sqrt{n})$ and $\delta = \frac{1}{n^2}$. This gives that with probability $1 - \delta$, all $\text{bias} \in \{\mu - 1, \mu, \mu + 1\}$ after a number of interactions

$$t \geq (n - 1)(2 \ln(K + \sqrt{n}) - \ln(\delta) - \ln(2)) \sim n(2 \ln(\sqrt{n}) + \ln(n^2)) = 3n \ln n.$$

Thus after time $3n \ln n$, all $\text{bias} \in \{\mu - 1, \mu, \mu + 1\}$ with high probability $1 - 1/n^2$.

If $|g| > 0.5|\mathcal{M}|$, then $|\mu| \geq 1$, so all remaining biased agents have the majority opinion, and we will stabilize in [Phase 2](#) to the correct majority output.

If $|g| \leq 0.5|\mathcal{M}|$, then $\mu = 0$, so now all $\text{bias} \in \{-1, 0, +1\}$. We will use [Lemma 4.6](#), with the sets of biased agents $A = \{a : a.\text{bias} = +1\}$ and $B = \{b : b.\text{bias} = -1\}$, which have initial sizes $|A| = a \cdot n$ and $|B| = b \cdot n$.

In the first case where $0.025|\mathcal{M}| \leq |g| \leq 0.5|\mathcal{M}|$, we have $a - b \geq 0.025m$ (assuming WLOG that A is the majority). Then by [Lemma 4.6](#), with high probability $1 - O(1/n^2)$, the count of B becomes 0 in at most time $\frac{5 \ln n}{2(a-b)} = \ln n \frac{5}{2 \cdot 0.025m} = \frac{100}{m} \ln n \leq 201 \ln n$. With all minority agents eliminated, we will again stabilize in [Phase 2](#) with the correct output.

In the second case where $|g| < 0.025|\mathcal{M}|$, we can use [Lemma 4.6](#) with constant $d = b - 0.0025m$. Then even with maximal gap $a - b = 0.025m$, with very high probability in constant time we bring the counts down to $b = 0.0025m$ and $a = 0.0275m$. Thus the total count of biased agents is at most $(0.0025m + 0.0275m)n = 0.03|\mathcal{M}|$.

Since all the above arguments take at most $O(\log n)$ time, for appropriate choice of counter constant c_1 , the given behavior happens before the first Clock agent advances to the next phase with high probability. \square

6 Analysis of main averaging [Phase 3](#)

The longest part in the proof is analyzing the behavior of the main averaging [Phase 3](#). The results of this section culminate in the following two theorems, one for the case of an initial tie, and the other for an initial biased distribution.

In the case of an initial tie, we will show that all biased agents have minimal $\text{exponent} = -L$ by the end of the phase.

Theorem 6.1. *If the initial configuration was a tie with gap 0, then by the end of [Phase 3](#), all biased agents have $\text{exponent} = -L$, with high probability $1 - O(1/n^2)$.*

Note that the where all biased agents have $\text{exponent} = -L$ gives a stable configuration in the next [Phase 4](#), with all agents having output = \top . Thus from [Theorem 6.1](#), we conclude that from an initial tie, the protocol will stabilize in [Phase 4](#) with high probability. Conversely, we have already observed that this configuration can only be reached in the case of a tie, since the sum of all biased agents would bound the

magnitude of the initial gap $|g| < 1$. Thus in the other case of a majority initial distribution, the agents will proceed through to [Phase 5](#) with probability 1.

In this majority case, we will show that a large majority of the Main agents have `opinion` set to the majority opinion and `exponent` $\in \{-l, -(l+1), -(l+2)\}$ is in a consecutive range of 3 possible exponents, where the value $-l$ depends on the initial distribution. In addition, we show the upper tail above this exponent $-l$ is very small.

Theorem 6.2. *Assume the initial gap $|g| < 0.025|\mathcal{M}|$. Let the exponent $-l = \lfloor \log_2(\frac{g}{0.4|\mathcal{M}|}) \rfloor$. Let $i = \text{sign}(g)$ be the majority opinion and M be the set of all agents with `role` = Main, `opinion` = i , `exponent` $\in \{-l, -(l+1), -(l+2)\}$. Then at the end of [Phase 3](#), $|M| \geq 0.92|\mathcal{M}|$ with high probability $1 - O(1/n^2)$.*

In addition, the total mass above exponent $-l$ is $\mu_{(>-l)} = \sum_{a.\text{exponent} > -l} |a.\text{bias}| \leq 0.002|\mathcal{M}|2^{-l}$, and the total minority mass is $\beta_- = \sum_{a.\text{opinion} = -i} |a.\text{bias}| \leq 0.004|\mathcal{M}|2^{-l}$.

Note the assumption of small initial gap g that we get from [Phase 1](#) is just for convenience in the proof, to reason uniformly about the base case behavior at hour $h = 0$ for our inductive argument. The rules of [Phase 3](#) would also work as intended with even larger initial gaps, just requiring a variant of the later analysis to acknowledge that the initial `exponent` = 0 is quite close to the final range `exponent` $\in \{-l, -(l+1), -(l+2)\}$.

6.1 Clock synchronization theorems

We will first consider the behavior of the Clock agents in [Phase 3](#). The goal is to show their `minute` fields remain tightly concentrated while moving from 0 up to kL , summarized in [Theorem 6.9](#). See [Fig. 8](#) for simulations of the clock distribution. The back tail behind the peak of the `minute` distribution decays exponentially, since each agent is brought ahead by epidemic at a constant rate and thus their counts each decay exponentially. The front part of the distribution decays much more rapidly. With a fraction f of agents at `minute` = i , the rate of the drip reaction is proportional to pf^2 . This repeated squaring leads to the concentrations at the leading minutes decaying doubly exponentially. The rapid decay is key to showing that very few Clock agents can get very far ahead of the rest.

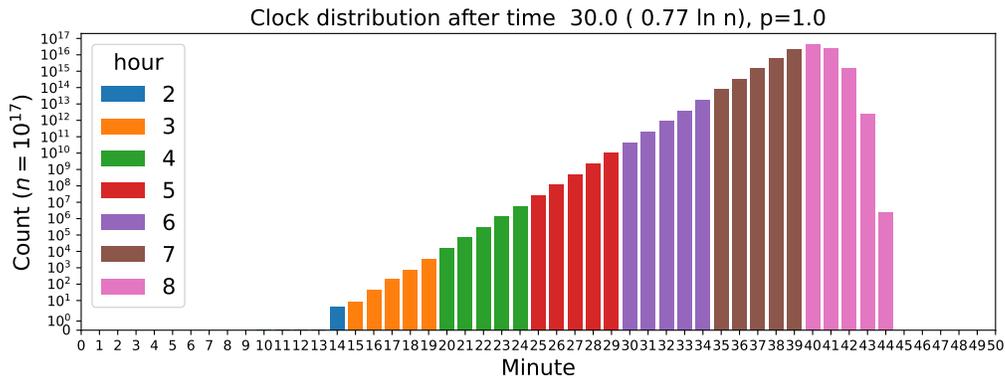
While our algorithm runs for only $O(\log n)$ minutes, the clock behavior we require can be made to continue for $O(n^c)$ minutes for arbitrary c . Our proofs rely on induction on minutes, and all results in this section hold with very high probability. Thus, we can take a union bound over polynomially many minutes and still keep the very high probability guarantees. If the clock runs for a superpolynomial number of minutes, the results of this section no longer hold.

We start by consider an entire population running the clock transitions (lines 1-5 of [Phase 3](#)), so $|\mathcal{C}| = n$. The following lemmas describe the behavior of just this clock protocol. In our actual protocol, the clock agents are a subpopulation $|\mathcal{C}| = c \cdot n$, and these clock transitions only happen when two clock agents meet with probability $\binom{|\mathcal{C}|}{2} / \binom{n}{2} \sim \frac{1}{c^2}$. This more general situation is handled in later theorems applying to our exact protocol.

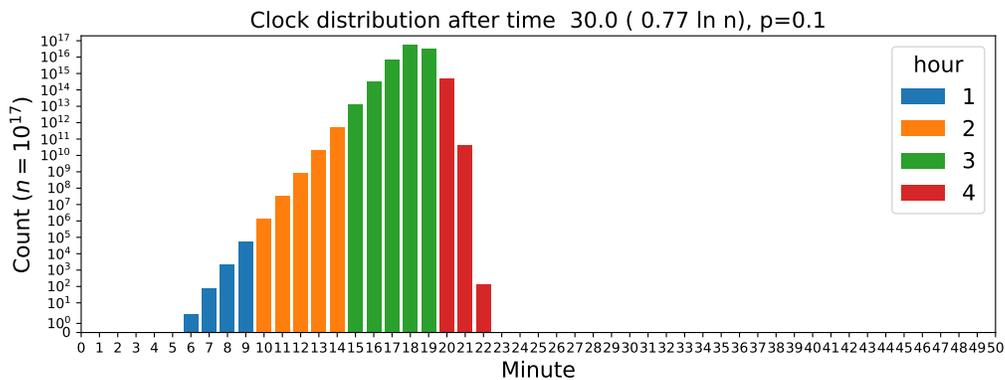
Definitions of values used in subsequent lemma statements. Throughout this section, we reference the following quantities. For each minute i and parallel time t , define $c_{\geq i}(t) = |\{c : c.\text{minute} \geq i\}| / |\mathcal{C}|$ to be the fraction of clock agents at minute i or beyond at time t . Then define $t_{\geq i}^+ = \min\{t : c_{\geq i}(t) > 0\}$, $t_{\geq i}^{0.1} = \min\{t : c_{\geq i}(t) \geq 0.1\}$ and $t_{\geq i}^{0.9} = \min\{t : c_{\geq i}(t) \geq 0.9\}$ to be the first times where this fraction becomes positive, hits 0.1 and hits 0.9.

We first show that the $c_{\geq i+1}$ is significantly smaller than $c_{\geq i}$ while both are still increasing, so the front of the clock distribution decays very rapidly. In our argument, we consider three types of reactions that change the counts at minutes i or above:

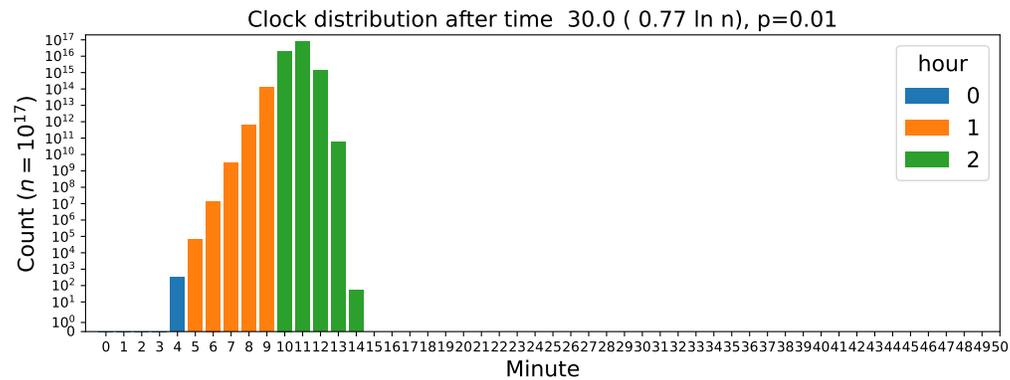
- (i) Drip reactions $i, i \rightarrow i, (i+1)$ ([Line 5](#) in [Phase 3](#))
- (ii) Epidemic reactions $j, k \rightarrow j, j$, for any minutes j, k with $k \leq i < j$ ([Line 3](#) in [Phase 3](#))



(a) Simulating the clock rules with $p = 1$.



(b) Simulating the clock rules with $p = \frac{1}{10}$.



(c) Simulating the clock rules with $p = \frac{1}{100}$.

Figure 8: Simulating the clock rules on a large population of size $n = 10^{17}$, with $k = 5$ minutes per hour and multiple values of p . The `minute` distribution gets tighter for smaller values of p . To make the `hour` distribution tighter, we can also simply make k larger. The full evolution of these distributions can be seen in the example notebook [1], along with the code that ran the large simulation to generate the data.

(iii) Epidemic reactions $j, k \rightarrow j, j$, for any minutes j, k with $k < i \leq j$ (Line 3 in Phase 3)

Note we ignore the drip reactions $(i-1), (i-1) \rightarrow (i-1), i$, which would only help the argument, but we do not have guarantees on the count at minute $i-1$.

We will try to show the relationship pictured in Fig. 8, that $c_{\geq(i+1)}(t) \approx c_{\geq i}(t)^2$. One challenge here is that this relationship will no longer hold with high probability for small values of $c_{\geq i}$. For example, if $c_{\geq i}(t) = n^{-1/2-\epsilon}$, the desired relationship would require $c_{\geq(i+1)}(t) < \frac{1}{n}$, meaning the count above minute i is 0. A drip reaction at minute i could happen with non-negligible probability $\approx n^{-2\epsilon}$.

To handle this difficulty, we define the set of *early drip agents* $D_{\geq(i+1)}(t)$, the set of agents that moved above minute i via a drip reaction at a time $\leq t$ when $c_{\geq i}(t) < n^{-0.45}$, or that were brought above minute i via an epidemic reaction with another early drip agent in $d_{\geq(i+1)}$. Note that the latter group of agents can move to minute $\geq i+1$ after time $t_{\geq(i+1)}^+$. Thus, this set represents the effect of any drip reactions that happen before $c_{\geq i}(t)$ grows large enough for our large deviation bounds to work. We then define $d_{\geq(i+1)}(t) = |D_{\geq(i+1)}(t)|/|\mathcal{C}|$ to be the fraction of early drip agents (including effects that persist beyond this time due to epidemic reactions). The set of agents above minute i that comprise the fraction $c_{\geq(i+1)}$ is then partitioned into the early drip agents that comprise $D_{\geq(i+1)}$, and the rest. By first ignoring these early drip agents $D_{\geq(i+1)}$, we can show that the rest of $c_{\geq(i+1)}$ stays small compared to $c_{\geq i}$.

Lemma 6.3. *With very high probability, if $n^{-0.45} \leq c_{\geq i}(t) \leq 0.1$, then $c_{\geq(i+1)}(t) \leq 0.9pc_{\geq i}(t)^2 + d_{\geq(i+1)}(t)$.*

Proof. The proof will proceed by induction on time t . As a base case, for all t such that $c_{\geq i}(t) < n^{-0.45}$, the statement holds simply by definition of $d_{\geq(i+1)}(t)$, which is equal to $c_{\geq(i+1)}(t)$.

For the inductive step, to show the relationship $c_{\geq(i+1)}(t) < 0.9pc_{\geq i}(t)^2 + d_{\geq(i+1)}(t)$ holds at time t , we will use the inductive hypothesis that $c_{\geq(i+1)}(t-0.1) < 0.9pc_{\geq i}(t-0.1)^2 + d_{\geq(i+1)}(t-0.1)$. Let $x(t) = c_{\geq i}(t)$ and $y(t) = c_{\geq(i+1)}(t) - d_{\geq(i+1)}(t)$, so we need to show $y(t) < 0.9px(t)^2$.

We first lower bound how much $x(t)$ grows by epidemic reactions, in order to show $x(t-0.1) < 0.84x(t)$. Using Lemma 4.5, the expected amount of time for an epidemic to grow from fraction $0.84x$ to x is

$$\frac{1}{2} \left[\ln(x) - \ln(0.84x) + \ln\left(\frac{1-0.84x}{1-x}\right) \right] \leq \frac{1}{2} \left[-\ln(0.84) + \ln\left(\frac{1-0.84 \cdot 0.1}{1-0.1}\right) \right] < 0.096,$$

where we used the fact that $\frac{1-0.84x}{1-x}$ is nondecreasing and $x(t) \leq 0.1$.

The minimum probability p^* is $\Theta(x(t)) = \Omega(n^{-0.45})$, and the expected number of interactions μ is $\Theta(n)$. So the application of Theorem 4.3 will give probability $1 - e^{-\Omega(n^{0.55})}$ for the epidemic to have grown enough within time 0.1. Thus $x(t-0.1) < 0.84x(t)$ with very high probability.

Next we bound how much $y(t)$ grows, by both epidemic reactions and drip reactions. The probability of a drip reaction is at most $px(t)^2$, so the expected number of drip reactions in time 0.1 is $0.1px(t)^2$. A standard Chernoff bound then gives that there are at most $0.11px(t)^2$ drip reactions with very high probability. We assume in the worst case all these drip reactions happen at time $t-0.1$, and then y grows by epidemic starting from $z = y(t-0.1) + 0.11px(t)^2$.

By Lemma 4.5, the expected amount of time for an epidemic to grow from fraction z to $1.23z$ is

$$\frac{1}{2} \left[\ln(1.23z) - \ln(z) + \ln\left(\frac{1-z}{1-1.23z}\right) \right] \geq \frac{1}{2} \ln(1.23) > 0.103.$$

The minimum probability $p^* = \Omega(x(t)^2) = \Omega(n^{-0.9})$, and the expected number of interactions $\mu = \Theta(n)$. So the application of Theorem 4.3 will give probability $1 - e^{-\Omega(n^{0.1})}$. Thus with very high probability

$$\begin{aligned} y(t) &\leq 1.23[y(t-0.1) + 0.11px(t)^2] \\ &\leq 1.23[0.9px(t-0.1)^2 + 0.11px(t)^2] \\ &\leq 1.23[0.9p(0.84x(t))^2 + 0.11px(t)^2] < 0.9px(t)^2. \end{aligned} \quad \square$$

In order to bound $c_{\geq(i+1)}(t)$ as a function of $c_{\geq i}(t)$ only (Theorem 6.5), we need to bound how large the set $D_{\geq i+1}$ of early drip agents can be. The strategy will be to show there is not enough time for the set $D_{\geq i+1}$ to grow very large starting from $t_{\geq i+1}^+$, the first time any agent appears at minute $\geq i+1$ (i.e., the first drip reaction into minute $i+1$), because there are only $O(\log \log n)$ minutes in the front tail (see Fig. 8), so the time between $t_{\geq i+1}^+$ and $t_{\geq i+1}^{0.1}$ is $O(\log \log n)$.

We will first need an upper bound on how long it takes the clock to move from one minute to the next.

Lemma 6.4. $t_{\geq i+1}^{0.1} - t_{\geq i}^{0.1} \leq 2.11 + \frac{1}{2} \ln\left(\frac{1}{p}\right)$ with very high probability.

Proof. First we argue that $c_{\geq i+1}(t_i^{0.1} + \frac{1}{2}) > 0.0045p$. If not, the count at minute i , $c_{\geq i}(t) - c_{\geq i+1}(t) > 0.1 - 0.0045 = 0.0955$ for all $t_i^{0.1} < t < t_i^{0.1} + 0.5$. Then, the probability of a drip reaction is at least $0.0955^2 p > 0.0091p$. By standard Chernoff bounds, we then have that in time $\frac{1}{2}$, there are at least $0.0045p$ drip reactions with very high probability. Thus $c_{\geq i+1}(t_i^{0.1} + \frac{1}{2}) > 0.0045p$ from just those drip reactions alone.

Now we argue that the amount of time it takes for epidemic reactions to bring $c_{\geq(i+1)}$ up to 0.1. By Lemma 4.5, the expected amount of time for an epidemic to grow from fraction 0.0045p to 0.1 is

$$\frac{1}{2} [\ln(0.1) - \ln(0.0045p) + \ln(1 - 0.0045p) - \ln(0.9)] < \frac{\ln(0.1) - \ln(0.0045) - \ln(0.9)}{2} - \frac{\ln p}{2} < 1.603 - \frac{\ln p}{2}.$$

As long as $p = \Theta(1)$, then the minimum probability $p^* = \Theta(p)$ is constant, and by Lemma 4.5, the epidemic takes at most time $1.61 - \frac{\ln p}{2}$ with very high probability.

In total, we then get that $t_{i+1}^{0.1} - t_i^{0.1} \leq 0.5 + 1.61 - \frac{\ln p}{2} = 2.11 + \frac{1}{2} \ln\left(\frac{1}{p}\right)$ with very high probability. \square

Proving that the set $D_{\geq i+1}$ remains small will let us prove the main theorem about the front tail of the clock distribution:

Theorem 6.5. With very high probability, if $n^{-0.4} \leq c_{\geq i}(t) \leq 0.1$, then $c_{\geq(i+1)}(t) < pc_{\geq i}(t)^2$.

Proof. The proof will proceed by induction on the minute i , where the base case is vacuous because $c_{\geq 0}(0) = 1$, so $c_{\geq 0}(t) > 0.1$ for all times t .

The inductive hypothesis will use two claims. The first is that the time $t_{\geq i}^{0.1} - t_{\geq i}^+ = O(\log \log n)$ for minute i . The second is that $d_{\geq i+1}(t_{\geq i}^{0.1}) = O(n^{-0.85})$. Note that using this second claim along with Lemma 6.3 proves the Theorem statement at minute i : when $n^{-0.4} \leq c_{\geq i}(t) \leq 0.1$, by Lemma 6.3 we have

$$c_{\geq(i+1)}(t) < 0.9pc_{\geq i}(t)^2 + d_{\geq i+1}(t) \leq pc_{\geq i}(t)^2,$$

because $c_{\geq i}(t)^2 \geq n^{-0.8}$, so the $d_{\geq i+1}$ term is negligible for sufficiently large n .

We will prove the first claim in two parts. First we argue that $t_{\geq i}^{0.1} - t_{\geq i}^+ = O(\log \log n)$, because the width of the front tail is at most $2 \log \log n$. We will show that at $t_{\geq i}^+$, when the first agent arrives at minute $= i$, we already have $c_{\geq j}(t_{\geq i}^+) \geq 0.1$, where $j = i - 2 \log \log n$ (for $i < 2 \log \log n$, we just have $j = 0$ and there is nothing to show because the width of the front tail can be at most i). First we move $\log \log n$ levels back to $k = i - \log \log n$, to show $c_{\geq k}(t_{\geq i}^+) \geq n^{-0.4}$.

Assume, for the sake of contradiction that $c_{\geq k}(t_{\geq i}^+) < n^{-0.4}$. Between $t_{\geq k}^+$ and $t_{\geq k}^{n^{-0.4}}$, consider the number of drips that happen from levels $k+1$ and above. By the inductive hypothesis, we have $c_{\geq k+1}(t) \leq pc_{\geq k}(t)^2 < n^{-0.8}$ during this whole time. Thus in any interaction of this period the probability of a drip above level $k+1$ is at most $p \cdot (n^{-0.8})^2 \leq n^{-1.6}$. The interval length is $t_{\geq k}^{n^{-0.4}} - t_{\geq k}^+ = O(\log \log n)$ by the inductive hypothesis, so the probability of having at least $\log \log n$ drips during this interval is at most

$$\left(\frac{O(n \log \log n)}{\log \log n} \right) (n^{-1.6})^{\log \log n} \leq \left(\frac{O(n \log \log n)}{n^{1.6}} \right)^{\log \log n} = n^{-\omega(1)}.$$

This implies $c_{\geq i}(t_{\geq k}^{n^{-0.4}}) = 0$ with very high probability.

Thus with very high probability, we already have $c_{\geq k}(t_{\geq i}^+) \geq n^{-0.4}$. Then we can iterate the inductive hypothesis $c_l(t) \leq p(c_{l-1}(t))^2$ for the $\log \log n$ minutes $l = k, k-1, \dots, j$, which implies $c_{\geq j}(t_{\geq i}^+) \geq 0.1$. Now that we have shown the width of the front tail is at most $2 \log \log n$, we use [Lemma 6.4](#), which shows that each minute takes $O(1)$ time, so it takes $O(\log \log n)$ time between $t_{\geq i}^+$ when the first agent gets to **minute** = i and $t_{\geq i}^{0.1}$, when the fraction at **minute** $\geq i$ reaches 0.1.

Now we prove the second claim, arguing that in this $O(\log \log n)$ time, $d_{\geq i+1}$ can grow to at most $O(n^{-0.85})$. By definition of $d_{\geq i+1}$, the drip reactions that increase $d_{\geq i+1}$ happen with probability at most $p(n^{-0.45})^2 = pn^{-0.9}$. By a standard Chernoff bound, the number of drip reactions in $O(\log \log n)$ time is $O(\log \log n \cdot n^{-0.9}) = O(n^{-0.89})$. Then we assume in the worst case this maximal number of drip reactions happen, and then $d_{\geq i+1}$ can grow by epidemic. By [Lemma 4.5](#), the time for an epidemic to grow from fraction $O(n^{-0.89})$ to $\Omega(n^{-0.85})$ is $\Omega(\log n) > O(\log \log n)$ with very high probability. Thus, with very high probability, we still have $d_{\geq i+1} = O(n^{-0.85})$ within $O(\log \log n)$ time. \square

Now the relationship proven in [Theorem 6.5](#) implies that $c_{\geq(i+1)}(t_{\geq i}^{0.1}) \leq 0.01p$. We can now use this bound to get lower bounds on the time required to move from one minute to the next.

Lemma 6.6. *With very high probability, $t_{i+1}^{0.1} - t_i^{0.1} \geq \frac{1}{2} \ln\left(1 + \frac{2}{9p}\right) - 0.01$.*

Proof. We start at time $t_{\geq i}^{0.1}$, where by [Theorem 6.5](#) we have $c_{\geq i+1}(t_{\geq i}^{0.1}) \leq 0.01p$ with very high probability.

Define $x = x(t) = c_{\geq i}(t)$ and $y = y(t) = c_{\geq i+1}(t)$. The number of interactions for $Y = yn$ to increase by 1 is a geometric random variable with mean $\frac{1}{\mathbb{P}[(\text{i})] + \mathbb{P}[(\text{ii})]}$, where the drip reaction **(i)** has probability $\mathbb{P}[(\text{i})] \sim p(x-y)^2 \leq p(1-y)^2$ and the epidemic reaction **(ii)** has probability $\mathbb{P}[(\text{ii})] \sim 2y(1-y)$. Assuming in the worst case that $y(t_i^{0.1}) = 0.01p$, then the number of interactions $T = (t_{i+1}^{0.1} - t_i^{0.1})n$ for Y to increase from $0.1pn$ to $0.1n$ is a sum of independent geometric random variables with mean

$$\begin{aligned} \mathbb{E}[T] &\geq \sum_{Y=0.01pn}^{0.1n-1} \frac{1}{p(1-Y/n)^2 + 2(Y/n)(1-Y/n)} \sim n \int_{0.01p}^{0.1} \frac{dy}{p(1-y)^2 + 2y(1-y)} \\ &= n \int_{0.01p}^{0.1} \frac{dy}{(1-y)(p + (2-p)y)} = n \int_{0.01p}^{0.1} \frac{1/2}{1-y} + \frac{1-p/2}{p + (2-p)y} dy \\ &= n \left[-\frac{1}{2} \ln(1-y) + \frac{1}{2} \ln(p + (2-p)y) \right]_{0.01p}^{0.1} \\ &= \frac{n}{2} [-\ln(0.9) + \ln(1 - 0.01p) + \ln(p + 0.1(2-p)) - \ln(p + 0.01p(2-p))] \\ &= \frac{n}{2} \left[-\ln(0.9) + \ln\left(\frac{1 - 0.01p}{p - 0.01p^2 + 0.02p}\right) + \ln(0.9p + 0.2) \right] \\ &\geq \frac{n}{2} \left[-\ln(0.9) + \ln\left(\frac{0.9p + 0.2}{1.02p}\right) \right] \geq \frac{n}{2} \left[-0.02 + \ln\left(1 + \frac{2.04}{9p}\right) \right] \end{aligned}$$

Note that the probability in the geometric includes the term $p(1-y)^2 \geq 0.81p$ since $y \leq 0.1$, thus the minimum geometric probability $p^* \geq 0.81p$ is bounded by a constant. Also the mean $\mu = \Theta(n)$ so by [Corollary 4.4](#), $\mathbb{P}[T \leq n(-0.01 + \frac{1}{2} \ln(1 + \frac{2}{9p})] \leq \exp(-\Theta(n))$, so the time $t_{i+1}^{0.1} - t_i^{0.1} = T/n \geq \frac{1}{2} \ln(1 + \frac{2}{9p}) - 0.01$ with very high probability. \square

The worst case upper bound for the dripping probability, $p(1-y)^2$, used in the above Lemma, is weakest when $p = 1$. We now give a special case lower bound that is stronger in the deterministic case with $p = 1$.

Lemma 6.7. *With very high probability, $t_{i+1}^{0.1} - t_i^{0.1} \geq 0.45$.*

Proof. We assume in the worst case that $p = 1$. Then by [Theorem 6.5](#), we have $c_{\geq i+1}(t_i^{0.1}) \leq 0.01$. This initial fraction will grow by epidemic to at most $0.01 \cdot e^{2 \cdot 0.45} (1 + \epsilon)$ with very high probability by [Lemma 4.5](#). We must also consider all agents that later make it to minute $i + 1$ by a drip reaction, and how much they grow by epidemic.

By time $t_i^{0.1} + s$, the fraction $c_{\geq i}$ could have increased to at most $0.1 + s$, since at most 1 agent can increase its minute in each interaction. Then the probability of a drip reaction at this time is at most $(0.1 + s)^2$. By standard Chernoff bounds, there will be at most $(1 + \epsilon)(0.1 + s)^2 n^{0.5}$ drip reactions in the next $n^{0.5}$ interactions with very high probability. Then by [Lemma 4.5](#), these agents can grow by epidemic by at most a factor $(1 + \epsilon)e^{2 \cdot (0.45 - s)}$ by time $t_i^{0.1} + 0.45$, with very high probability. Summing over consecutive groups of $n^{0.5}$ interactions at time $i \cdot \frac{n^{0.5}}{n}$ and using the union bound, we get a total bound

$$\begin{aligned} c_{\geq i+1}(t_i^{0.1} + 0.45) &\leq 0.01 \cdot e^{2 \cdot 0.45} (1 + \epsilon) + \sum_{i=0}^{0.45 n^{0.5}} (1 + \epsilon)(0.1 + n^{-0.5} i)^2 n^{0.5} \cdot e^{2 \cdot (0.45 - n^{-0.5} i)} \\ &\sim (1 + \epsilon) \cdot e^{0.9} \left[0.01 + \int_0^{0.45} (1 + s)^2 e^{-2s} ds \right] \leq 0.45. \end{aligned}$$

□

We now summarize all bounds on the length of a clock minute in a single theorem:

Theorem 6.8. *Let $t_{\geq i+1}^{0.1} - t_{\geq i}^{0.1}$ be the time between when a fraction 0.1 of agents have `minute` $\geq i$ and when a fraction 0.1 of agents have `minute` $\geq i + 1$. Then with very high probability,*

$$\max \left(0.45, \frac{1}{2} \ln \left(1 + \frac{2}{9p} \right) - 0.01 \right) \leq t_{\geq i+1}^{0.1} - t_{\geq i}^{0.1} \leq 2.11 + \frac{1}{2} \ln \left(\frac{1}{p} \right)$$

These bounds are shown in [Fig. 9](#), along with sampled minute times from simulation. This suggests the actual time per minute is roughly $0.75 + \frac{1}{2} \ln \left(\frac{1}{p} \right)$.

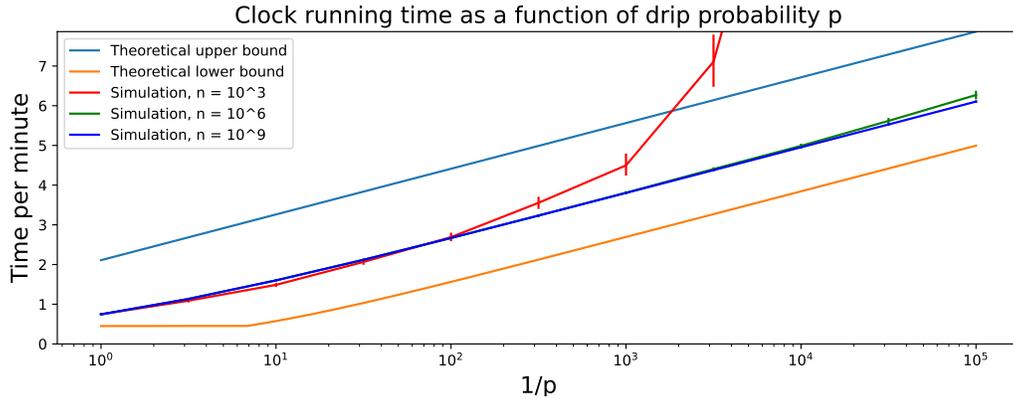


Figure 9: The upper and lower bounds from [Theorem 6.8](#) for the time of one clock minute, along with samples from simulation. For each value of n , 100 minute times were sampled, taking $t_{i+1}^{0.1} - t_i^{0.1}$ for $i = 9, \dots, 18$ over 10 independent trials. All our proofs assume p is constant, and for any fixed value of p , will only hold for sufficiently large n . The case $n = 10^3$ shows that when $p = O(1/n)$, the bounds no longer hold. This is to be expected because the expected number of drips becomes too small for large deviation bounds to still hold.

We can now build from these theorems to get bounds on the values of `hour`. For a clock agent a , define $a.\text{hour} = \lfloor \frac{a.\text{minute}}{k} \rfloor$. Define $\text{start}_h = \min (t : |\{a : a(t).\text{hour} \geq h\}| \geq 0.9|\mathcal{C}|)$ be the first time when the

fraction of clock agents at hour h or beyond reaches 0.9 and $\text{end}_h = \min(t : |\{a : a(t).\text{hour} > h\}| \leq 0.001|\mathcal{C}|)$ be the first time when the fraction of clock agents beyond hour h reaches 0.001. Define the *synchronous hour* h to be the parallel time interval $[\text{start}_h, \text{end}_h]$, i.e. when fewer than 0.1% are in any hour beyond h and at least $(90 - 0.1) = 89.9\%$ are in hour h . Note that if $\text{end}_h < \text{start}_h$ then this interval is empty, but we show this happens with low probability. We choose the threshold $0.001|\mathcal{C}|$ to be a sufficiently small constant for later proofs.

Recall $c = |\mathcal{C}|/n$ is the fraction of clock agents and k is the number of minutes per hour.

Theorem 6.9. *Consider a fraction c of agents running the clock protocol, with $p = 1$. Then for every synchronous hour h , its length $\text{end}_h - \text{start}_h \geq \frac{1}{c^2}[0.45k - 3.1]$, with very high probability. The time between consecutive synchronous hours $\text{start}_{h+1} - \text{start}_h \leq \frac{1}{c^2}[2.11k + 2.2]$ with very high probability.*

Proof. Note that the previous lemmas assumed $|\mathcal{C}| = n$, so the entire population was running the clock algorithm. In reality, we have a fraction $c = |\mathcal{C}|/n$ of clock agents. The reactions we considered only happen between two clock agents, which interact with probability $\sim c^2$. Thus we can simply multiply the bounds from our lemmas by the factor $\frac{1}{c^2}$ to account for the number of regular interactions for the requisite number of clock interactions to happen, which is very close to its mean with very high probability by standard Chernoff bounds.

Because our definition references time $t_{>i}^{0.9}$ when the fraction reaches 0.9, we will first bound the time it takes an epidemic to grow from 0.1 to 0.9. By [Lemma 4.5](#), this takes parallel time t , where

$$\mathbb{E}[t] \sim \frac{\ln(0.9) - \ln(1 - 0.9) - \ln(0.1) + \ln(1 - 0.1)}{2} = \ln(9) < 2.2.$$

Since $\ln(9)(1 + \epsilon) < 2.2$, this completes within parallel time 2.2 with very high probability.

Since $c.\text{hour} = h \iff hk \leq c.\text{minute} < (h + 1)k$, the times $\text{start}_h = t_{hk}^{0.9}$ and $\text{end}_h = t_{hk}^{0.001}$. For the upper bound, by [Lemma 6.4](#) we have $t_{i+1}^{0.1} - t_i^{0.1} \leq 2.11/c^2$ with very high probability. We start at $t_{hk}^{0.9} \geq t_{hk}^{0.1}$, and sum over the k minutes in hour h , then add at most time $\frac{2.2}{c^2}$ between $t_{(h+1)k}^{0.1}$ and $t_{(h+1)k}^{0.1}$. This gives that $\text{start}_{h+1} - \text{start}_h \leq \frac{1}{c^2}[2.11k + 2.2]$ with very high probability.

For the lower bound, by [Theorem 6.5](#), at time $t_{(h+1)k-2}^{0.1}$, we have $c_{>(h+1)k} \leq (0.1^2)^2 = 10^{-4} < 0.001$, so $t_{(h+1)k-2}^{0.1} < \text{end}_h$. Then $\text{start}_h \leq t_{hk}^{0.1} + \frac{2.2}{c^2}$. Using [Lemma 6.7](#), we have $t_{i+1}^{0.1} - t_i^{0.1} \leq 0.45/c^2$ with very high probability, for each $i = hk, \dots, hk + k - 3$. All together, this gives $\text{end}_h - \text{start}_h \geq \frac{1}{c^2}[0.45(k - 2) - 2.2] = \frac{1}{c^2}[0.45k - 3.1]$. \square

We will set $k = 45$ to give us sufficiently long synchronous hours for later proofs. Since every hour takes constant time with very high probability, all L hours will finish within $O(\log n)$ time.

We finally show one more lemma concerning how the Clock agents affect the Main agents. The Clock agents will change the hour of the \mathcal{O} agents via [Line 9](#) of [Phase 3](#). There are a small fraction $0.001|\mathcal{C}|$ of Clock agents that might be running too fast and thus have a larger hour than the synchronized hour. We must now show these agents are only able to affect a small fraction of the Main agents. Intuitively, the Clock agents with $\text{hour} > h$ bring up the hour of both \mathcal{O} agents and other Clock agents. The following lemma will show they don't affect too many Main agents before also bringing in a large number of Clock agents.

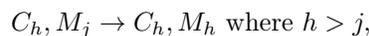
We will redefine $c_{>h} = c_{>h}(t) = |\{c : c.\text{hour} > h\}|/|\mathcal{C}|$ to be the fraction of Clock agents beyond hour h , and similarly $m_{>h} = m_{>h}(t) = |\{m : m.\text{hour} > h\}|/|\mathcal{M}|$ to be the fraction of Main agents beyond hour h .

Lemma 6.10. *For all times $t \leq \text{end}_h$, we have $m_{>h}(t) \leq 1.2c_{>h}(\text{end}_h) = 0.0012$ with very high probability.*

Proof. We have $c_{>h}(\text{end}_h) = 0.001$ by the definition of synchronous hour h . Thus it suffices to show that $m_{>h}(t) \leq 1.2c_{>h}(\text{end}_h)$.

Recall $c = |\mathcal{C}|/n$ and $m = |\mathcal{M}|/n$, so $c \cdot c_{>h}$ and $m \cdot m_{>h}$ are rescaled to be fractions of the whole population.

We will assume in the worst case that every Main agent can participate in the clock update reaction



so the probability of clock update reaction is $2c \cdot c_{>h} \cdot m(1 - m_{>h})$. Among the Clock agents, we will now only consider the epidemic reactions $C_h, C_j \rightarrow C_h, C_h$ between agents in different hours $h > j$, so the probability of the clock epidemic reaction is $2c^2 \cdot c_{>h}(1 - c_{>h})$.

We use the potential $\Phi(t) = m_{>h}(t) - 1.1 \cdot c_{>h}(t)$. Note that initially $\Phi(0) = 0$ since both terms are 0. The desired inequality is $\Phi(\text{end}_h) \leq 0.1c_{>h}(\text{end}_h) = 0.0001$, and we will show this holds with very high probability by Azuma's Inequality because Φ is a supermartingale. The clock update reaction increases Φ by $\frac{1}{|\mathcal{M}|} = \frac{1}{mn}$. The clock epidemic reaction decreases Φ by $\frac{1.1}{|C|} = \frac{1.1}{cn}$. This gives an expected change

$$\begin{aligned} \mathbb{E}[\Phi(t + 1/n) - \Phi(t)] &= \frac{1}{mn} [2c \cdot c_{>h} \cdot m(1 - m_{>h})] - \frac{1.1}{cn} [2c^2 \cdot c_{>h}(1 - c_{>h})] \\ &= \frac{2c \cdot c_{>h}}{n} [(1 - m_{>h}) - 1.1(1 - c_{>h})] \\ &\leq \frac{2c \cdot c_{>i}}{n} [1 - 1.1(0.999)] < 0. \end{aligned}$$

Thus the sequence $(\Phi_j = \Phi(\frac{j}{n}))_{j=0}^{n \cdot \text{end}_h}$ is a supermartingale, with bounded differences $|\Phi_{j+1} - \Phi_j| \leq \max(\frac{1}{mn}, \frac{r}{cn}) = O(\frac{1}{n})$. So we can apply Azuma's Inequality ([Theorem 4.2](#)) to conclude

$$\mathbb{P}[\Phi_{n \cdot \text{end}_h} \geq \delta] \leq \exp\left(-\frac{\delta^2}{2 \sum_{j=0}^{n \cdot \text{end}_h} O(\frac{1}{n})^2}\right) = \exp(-O(n)\delta^2),$$

since by [Theorem 6.9](#), we have time $\text{end}_h = O(1)$ with very high probability. Thus we can choose $\delta = 0.1c_{>h}(\text{end}_h) = 0.0001$ to conclude that $m_{>h}(\text{end}_h) \leq 1.2c_{>h}(\text{end}_h)$ with very high probability $1 - \exp(-\Omega(n)) = 1 - O(n^{-\omega(1)})$.

The lemma statement for times $t \leq \text{end}_h$ simply follows from the monotonicity of the value $m_{>h}$, since agents only decrease their hour field. \square

6.2 Phase 3 exponent dynamics

We now analyze the behavior of the Main agents in [Phase 3](#). We will show their exponent fields stay roughly synchronized with the hour of the Clock agents, decreasing from 0 toward $-L$. We first use the above results on the clock to conclude that during synchronous hour h , the tail of either \mathcal{O} agents with hour $> h$ or biased agents with exponent $< -h$ is sufficiently small.

Lemma 6.11. *Until time end_h , the count $|\{\mathcal{O} : \mathcal{O}.\text{hour} > h\} \cup \{b : b.\text{exponent} < -h\}| \leq 0.0024|\mathcal{M}|$ with very high probability.*

Proof. By [Lemma 6.10](#), the count of \mathcal{O} agents that have been pulled up by a Clock agent to hour $> h$ is at most $0.0012|\mathcal{M}|$ with very high probability. Each of these agents could participate in a split reaction that results in two biased agents with exponent $< -h$, increasing the total count of $|\{\mathcal{O} : \mathcal{O}.\text{hour} > h\} \cup \{b : b.\text{exponent} < -h\}|$ by 1. Thus this total count can be at most twice as large: $\leq 0.0024|\mathcal{M}|$. \square

Our main argument will proceed by induction on synchronous hours. During each synchronous hour h , we will first show that at least a constant fraction $0.77|\mathcal{M}|$ of agents have bias = T with hour = h . Then we will show that split reactions bring most biased agents down to exponent = $-h$. Finally we will show that cancel reactions will reduce the count of biased agents, leaving sufficiently many \mathcal{O} agents for the next step of the induction.

Recall that the initial gap

$$g = (|\{a : a.\text{input} = \text{A}\}| - |\{b : b.\text{input} = \text{B}\}|) = \sum_{m.\text{role}=\text{Main}} m.\text{bias}$$

is both the difference between the counts of A and B agents in the initial configuration and the invariant value of the total bias. We define $\beta_+ = \beta_+(t) = \sum_{a.\text{opinion}=+1} |a.\text{bias}|$ and $\beta_- = \beta_-(t) = \sum_{b.\text{opinion}=-1} |b.\text{bias}|$

to be the total unsigned bias of the positive A agents at time t and negative B agents at time t . Then the net bias $g = \beta_+ - \beta_-$ is the invariant.

We also define $\mu = \mu(t) = \beta_+ + \beta_-$ to be the total unsigned bias, which we interpret as “mass”. Note that split reactions preserve μ , while cancel reactions strictly decrease μ . The initial configuration has mass $\mu(0) = n$, and if we eliminate all of the minority opinion then we will get $\beta_- = 0$ and $\mu = g$. Thus decreasing the mass shows progress toward reaching consensus. Also, if every biased agent decreased their exponent by 1, this would exactly cut the μ in half. We will show an upper bound on μ that cuts in half after each synchronous hour, which implies that on average all biased agents are moving down one exponent.

We also define $\mu_{(>-i)}(t) = \sum_{m.\text{exponent} > -i} |m.\text{bias}|$ as the total mass of all biased agents above exponent $-i$. Note that a bound $\mu_{(>-i)} \leq x2^{-i+1}$ gives a bound on total count $|\{a : a.\text{exponent} > -i\}| \leq x$, since even if all agents above exponent $-i + 1$ split down to exponent $-i + 1$, they would have count at most x . Also note that $\mu(t)$ and $\mu_{(>-i)}(t)$ are nonincreasing in t , since the mass above a given exponent can never increase.

This inductive argument on synchronous hours will stop working once we reach a low enough exponent that the gap has been sufficiently amplified. We define $g_i = g \cdot 2^i$, which we call the relative gap to hour i / exponent $-i$, because if all agents had `exponent = -i`, then a gap $g_i = |\{a : a.\text{opinion} = +1\}| - |\{b : b.\text{opinion} = -1\}|$ would maintain the invariant $g = \sum_v v.\text{bias} = \sum_{a.\text{opinion}=+1} \frac{1}{2^i} - \sum_{b.\text{opinion}=-1} \frac{1}{2^i} = \frac{g_i}{2^i}$. Note that $g_0 = g$, and the relative gap doubles as we increment the hour and decrement the exponent. So if the Main agents have roughly synchronous exponents, there will be some minimal exponent where the relative gap has grown to exceed the number of Main agents $|\mathcal{M}|$, and there are not enough agents available to maintain the invariant using lower exponents.

We now formalize this idea of a minimal exponent. In the case where there is an initial majority, $g_i \neq 0$, and because we assume the majority is A, we have $g_i > 0$. Define the minimal exponent $-l = \lfloor \log_2(\frac{g}{0.4|\mathcal{M}|}) \rfloor$ to be the unique exponent corresponding to relative gap $0.4|\mathcal{M}| \leq g_l < 0.8|\mathcal{M}|$. For larger exponents $-i \geq -l + 5 = -(l - 5)$, we have $g_i \leq g_{l-5} < 0.025|\mathcal{M}|$. Thus for hours $0, 1, \dots, l - 5$ and exponents $0, -1, \dots, -l + 5$, the gap is still very small. The small gap makes the inductive argument stronger, and we will use this small gap to show a high rate of cancelling keeps shrinking the mass μ in half each hour and keeps the count of \mathcal{O} agents large, above a constant fraction $0.77|\mathcal{M}|$.

For the last few hours $l - 4, \dots, l$ and exponents $-l + 4, \dots, -l$, the doubling gap weakens the argument. Thus we have separate bounds for each of these hours. These weaker bounds acknowledge the fact that the majority count is starting to increase while the count of \mathcal{O} agents is starting to decrease.

Theorem 6.12. *With very high probability the following holds. For all synchronous hours $h = 0, \dots, l$, during times $[\text{start}_h + \frac{2}{c}, \text{end}_h]$, the count \mathcal{O}_h of \mathcal{O} agents at hour = h , obeys $|\mathcal{O}_h| \geq \tau_h|\mathcal{M}|$. Then by time $\text{start}_h + \frac{2}{c} + \frac{41}{m}$, the mass $\mu_{(>-h)}$ above exponent $-h$ satisfies $\mu_{(>-h)} \leq 0.001 \cdot 2^{-h+1}$. Then by time $\text{start}_h + \frac{2}{c} + \frac{47}{m} \leq \text{end}_h$, the total mass $\mu(\text{end}_h) \leq \rho_h|\mathcal{M}|2^{-h}$.*

The constant $\rho_h = 0.1$ for $h \leq l - 5$. Then we have $\rho_{l-4} = 0.104$, $\rho_{l-3} = 0.13$, $\rho_{l-2} = 0.212$, $\rho_{l-1} = 0.408$, and $\rho_l = 0.808$.

The constant $\tau_h \geq 0.97 - 2\rho_{(h-1)}$, so $\tau_h \geq 0.77$ for $h \leq l - 4$, and the minimum value $\tau_l \geq 0.15$.

Note that in the case of a tie, l is undefined since we always have gap $g_i = 0$. Here the stronger inductive argument will hold for all hours and exponents. In the tie case, we will technically define $l = L + 5$ so the stronger $h \leq l - 5$ bounds apply to all hours.

We will prove the three sequential claims via three separate lemmas. The first argument of [Lemma 6.13](#), where the clock brings a large fraction of \mathcal{O} agents up to hour h , will need parallel time $\frac{2}{c}$. The second argument of [Lemma 6.15](#), where the \mathcal{O} agents reduce the mass above exponent $-h$ by split reactions, will need parallel time $\frac{41}{m}$. The third argument of [Lemma 6.16](#), where cancel reactions at exponent $-h$ reduce the total mass, will need parallel time $\frac{6}{m}$. Thus the total time we need in a synchronous hour is

$$\frac{2}{c} + \frac{47}{m} \leq \frac{49}{c} \leq \frac{17}{c^2} \leq \frac{0.45 \cdot 45 - 3.1}{c^2},$$

where we use that $c < \frac{1}{3} < m$ by [Lemma 5.2](#). Thus by [Theorem 6.9](#), since we have constant $k = 45$ minutes per hour, each synchronous hour is long enough with very high probability.

The argument proceeds by induction, with each lemma using the previous lemmas and the inductive hypotheses at previous hours. The base case comes from [Lemma 5.3](#), where we have that the initial gap $|g| < 0.025m$, so $l - 5 \geq 0$, and the starting mass is at most $0.03|\mathcal{M}|$. This mass bound gives the base case for [Lemma 6.16](#), whereas since $h = 0$ is the minimum possible hour, the base cases for [Lemma 6.13](#) and [Lemma 6.15](#) are trivial.

Lemma 6.13. *With very high probability the following holds. For each hour $h = 0, \dots, l$, during the whole synchronous hour $[\text{start}_h, \text{end}_h]$, the count of \mathcal{O} agents $|\mathcal{O}| \geq (0.9976 - 2\rho_{(h-1)})|\mathcal{M}|$. Then during times $[\text{start}_h + \frac{2}{c}, \text{end}_h]$, the count of \mathcal{O} agents at hour = h , $|\mathcal{O}_h| \geq (0.97 - 2\rho_{(h-1)})|\mathcal{M}|$.*

Proof. We show the first claim, that during synchronous hour h , the count $|\mathcal{O}| \geq (0.9976 - 2\rho_{(h-1)})|\mathcal{M}|$ by process of elimination. By [Lemma 6.16](#), by time $\text{end}_{h-1} < \text{start}_h$, the total mass $\mu \leq \rho_{(h-1)}|\mathcal{M}|2^{-h+1}$, which implies the total count $\{a : a.\text{exponent} \geq -h\} \leq 2\rho_{(h-1)}|\mathcal{M}|$ in all future configurations, since total mass is nonincreasing. Then by [Lemma 6.11](#), the count $|\{\mathcal{O} : \mathcal{O}.\text{hour} > h\} \cup \{b : b.\text{exponent} < -h\}| \leq 0.0024|\mathcal{M}|$ until time end_h . This leaves $|\mathcal{M}| - 0.0024|\mathcal{M}| - 2\rho_{(h-1)}|\mathcal{M}| = (0.9976 - 2\rho_{(h-1)})|\mathcal{M}|$ agents who must be \mathcal{O} agents with $\text{hour} \geq h$ until time end_h .

By definition of time start_h , there are at least $0.9|\mathcal{C}|$ Clock agents with $\text{hour} \geq h$ that will bring these \mathcal{O} agents up to hour = h . We need to bring all but $0.0276|\mathcal{M}|$ of these agents up to hour = h to achieve the desired bound $|\mathcal{O}_h| \geq (0.97 - 2\rho_{(h-1)})|\mathcal{M}|$. We can apply [Lemma 4.7](#), with $a = 0.9c$, $b_1 = (0.9976 - 2\rho_{(h-1)})m < (0.9976 - 2 \cdot 0.1)m$, and $b_2 = 0.0276m$ to conclude this happens after parallel time t , where $t < (1 + \epsilon)\mathbb{E}[t]$ with very high probability. The expected time

$$\mathbb{E}[t] = \frac{\ln(b_1) - \ln(b_2)}{2a} \leq \frac{\ln(0.7976m) - \ln(0.0276m)}{2 \cdot 0.9c} \leq \frac{1.89}{c}.$$

Thus for small constant $\epsilon > 0$, we have $t < (1 + \epsilon)\frac{1.89}{c} < \frac{2}{c}$ with very high probability. \square

In order to reason about the total mass $\mu_{(>-h)}$ above exponent $-h$, we will define the potential function $\phi_{(>-h)}$, where for a biased agent a with $a.\text{exponent} = -i \geq -h + 1$ at time t , $\phi_{(>-h)}(a, t) = 4^{-i+h-1}$. The global potential

$$\phi_{(>-h)}(t) = \sum_{a.\text{exponent} > -h} \phi(a, t) \geq \sum_{a.\text{exponent} > -h} 4^{-h+1+h-1} = |\{a : a.\text{exponent} > -h\}|.$$

Since $\phi_{(>-h)}$ upper bounds the count above exponent h , we can bound the mass $\mu_{(>-h)}(t) \leq 2^{-h+1}\phi_{(>-h)}(t)$. Also note that unlike the mass, $\phi_{(>-h)}(t)$ strictly decreases via split reactions since $4^{-i} > 4^{-i-1} + 4^{-i-1}$. This will let us show that $\phi_{(>-h)}$ exponentially decays to 0 when there are a constant fraction of \mathcal{O} agents to do these splits.

We first show that by hour h exponents significantly far above $-h$ are empty. Letting $q = \lfloor \frac{\ln n}{3} \rfloor$, we will show the potential $\phi_{(>-h+q)}$ hits 0 by hour h .

Lemma 6.14. *For each hour $h = 0, \dots, l$, by time $\text{start}_h + \frac{2}{c}$, the maximum level among all biased agents is at most $-h + q$ with high probability $1 - O(1/n^{12})$.*

Proof. The statement is vacuous for hours $h < q$, so we must only consider hours $h \geq q$. We use the potential $\phi_{(>-h+q)}$, and start the argument at time $t_{\text{start}} = \text{start}_{(h-q)} + \frac{2}{c} + \frac{41}{m}$ where inductively by [Lemma 6.15](#) $\phi_{(>-h+q)}(t_{\text{start}}) \leq 0.001|\mathcal{M}|$. We must show that by time $t_{\text{end}} = \text{start}_h + \frac{2}{c}$, we have $\phi_{(>-h+q)}(t_{\text{end}}) = 0$.

By [Lemma 6.13](#), the count of \mathcal{O} agents with $\text{hour} \geq h - q$, $|\mathcal{O}_{(\geq h-q)}| \geq 0.77|\mathcal{M}|$ during all synchronous hours after $\text{start}_{(h-q)} + \frac{2}{c}$ and up through synchronous hour $l - 5 \geq h - 5$. Thus the interval we consider consists of at least $q - 5$ synchronous hours. By [Theorem 6.9](#), each synchronous hour takes at least time $[1.48(45 - 2) - 2.2]/c^2 \geq 17/c^2 \geq 51/m$, since $c < \frac{1}{3} < m$ by [Lemma 5.2](#). Thus the total time for this argument is at least $t_{\text{end}} - t_{\text{start}} \geq \frac{51(q-5)}{m}$ time.

For each of the $\frac{51(q-5)}{m}n$ interactions in this interval, we consider the expected change to $\phi_{(>-h+q)}$. For each biased agent a with $a.\text{exponent} = -i > -h + q$, a split reaction will change the potential by $\Delta\phi_a \leq 4^{h-1}(2 \cdot 4^{-i-1} - 4^{-i}) = 4^{h-1}(-\frac{1}{2}4^{-i}) = -\frac{1}{2}\phi_a$. Then in each interaction at parallel time t , the expected change in the potential

$$\begin{aligned} \mathbb{E}[\phi_{(>-h+q)}(t+1/n) - \phi_{(>-h+q)}(t)] &\leq \sum_{a.\text{exponent} > (-h+q)} \mathbb{P}[a \text{ splits}] \cdot \Delta\phi_a \\ &\leq \sum_{a.\text{exponent} > (-h+q)} \frac{2 \cdot 0.77m}{n} \cdot -\frac{1}{2}\phi_a = -\frac{0.77m}{n}\phi_{(>-h+q)}(t). \end{aligned}$$

Then we have $\mathbb{E}[\phi_{(>-h+q)}(t+1/n)|\phi_{(>-h+q)}(t)] \leq (1 - \frac{0.77m}{n})\phi_{(>-h+q)}(t)$.

We now recursively apply this bound to all $\frac{51(q-5)}{m}n$ interactions beginning at time t_{start} :

$$\begin{aligned} \mathbb{E}[\phi_{(>-h+q)}(t_{\text{end}})|\phi_{(>-h+q)}(t_{\text{start}})] &\leq \left(1 - \frac{0.77m}{n}\right)^{\frac{51(q-5)}{m}n} \phi_{(>-h+q)}(t_{\text{start}}) \\ \mathbb{E}[\phi_{(>-h+q)}(t_{\text{end}})] &\leq \exp(-0.77 \cdot 51 \ln n/3) \cdot \exp(51 \cdot 5 \cdot 0.77) \cdot 0.001|\mathcal{M}| \\ &\leq n^{-13} \exp(197) \cdot 0.001mn \\ &= O(n^{1-13}) = O(n^{-12}). \end{aligned}$$

Finally, since $\phi_{(>-h+q)}$ takes nonnegative integer values, we can apply Markov's Inequality to conclude $\mathbb{P}[\phi_{(>-h+q)}(t_{\text{end}}) > 0] = O(1/n^{12})$. \square

Now we can reason about the potential $\phi_{(>-h)}$ during hour h , which will decrease by a large constant factor. The upper bound on $\phi_{(>-h)}$ gives the a bound on the mass of the upper tail $\mu_{(>-h)}$.

Lemma 6.15. *For each hour $h = 0, \dots, l$, by time $\text{start}_h + \frac{2}{c} + \frac{41}{m}$, the potential $\phi_{(>-h)} \leq 0.001|\mathcal{M}|$ with very high probability. This implies the mass above exponent $-h$ is $\mu_{(>-h)} \leq 0.001|\mathcal{M}|2^{-h+1}$.*

Proof. By Lemma 6.13, after time $\text{start}_h + \frac{2}{c}$, we have a count $|\mathcal{O}_h| \geq \tau_h|\mathcal{M}|$, where the weakest bound is at hour l , where $\tau_l = 0.15$.

Inductively, we have $\phi_{(>-h+1)}(\text{start}_h) \leq 0.001|\mathcal{M}|$ by time $\text{start}_{(h-1)} + \frac{2}{c} + \frac{41}{m}$. By Lemma 6.16, by time $\text{end}_{(h-1)}$, the total mass $\mu \leq \rho_{(h-1)}|\mathcal{M}|2^{-h+1} \leq \rho_{(l-1)}|\mathcal{M}|2^{-h+1}$. Thus there are at most $\rho_{(l-1)}|\mathcal{M}| = 0.408|\mathcal{M}|$ biased agents with exponent $-h + 1$, which lets us bound the potential $\phi_{(>-h)}(\text{start}_h + \frac{2}{c}) \leq (0.408 + 4 \cdot 0.001)|\mathcal{M}| = 0.412|\mathcal{M}|$. Thus we must drop the potential by the constant factor 412.

To show that ϕ drops by a constant factor, we will use $\Phi(t) = \ln(\phi_{(>-h)}(t))$, which will be a supermartingale. If agent a at with exponent $-i$ splits, with $\phi_a = 4^{-i+h-1}$, this changes the potential ϕ by $\Delta\phi_a = -\frac{1}{2}\phi_a$. The potential Φ then changes by

$$\Delta\Phi_a = \ln(\phi_{(>-h)} + \Delta\phi_a) - \ln(\phi_{(>-h)}) = \ln\left(1 + \frac{-\frac{1}{2}\phi_a}{\phi_{(>-h)}}\right) \leq -\frac{\phi_a}{2\phi_{(>-h)}}.$$

The expected change in Φ is then

$$\mathbb{E}[\Delta\Phi] \leq \sum_{a.\text{exponent} > -h} \mathbb{P}[a \text{ splits}] \cdot \Delta\Phi_a \leq \sum_{a.\text{exponent} > -h} \frac{2\tau_h m}{n} \cdot -\frac{\phi_a}{2\phi_{(>-h)}} = -\frac{0.15m}{n}.$$

We define the supermartingale $(\Phi_j)_{j=0}^{\frac{41}{m}n}$, where $\Phi_j = \Phi(\text{start}_h + \frac{2}{c} + \frac{j}{n})$. Then the desired inequality is $\phi_{(>h+1)}(\text{start}_h + \frac{2}{c} + \frac{41}{m}) \leq 0.001|\mathcal{M}| \leq \phi_{(>h+1)}(\text{start}_h + \frac{2}{c})/412$, so we need to show $\Phi_{\frac{41}{m}n} - \Phi_0 \leq \ln(\frac{1}{412})$. The expected value

$$\mathbb{E}[\Phi_{\frac{41}{m}n} - \Phi_0] \leq -\frac{0.15m}{n} \cdot \frac{41}{m}n = -6.15 \leq \ln\left(\frac{1}{412}\right) - 0.12.$$

To apply Azuma's Inequality, we will need a bound on the difference $|\Phi_{j+1} - \Phi_j| \leq \max \left| \frac{\phi_a}{2\phi_{(>-h)}} \right|$. During the interval we consider, $\phi_{(>-h)} \geq 0.001|\mathcal{M}|$, since after that the desired inequality will hold. By [Lemma 6.14](#), by time $\text{start}_h + \frac{2}{c}$, the maximum exponent in the population is at most $-h+q$, where $q = \lfloor \frac{\ln n}{3} \rfloor$, so $\phi_a \leq 4^{(-h+q)+h-1} = 4^{q-1}$. Using the fact that $4^q = e^{\ln 4 \ln n / 3} = O(n^{0.47})$, we can bound the largest change in Φ as

$$|\Phi_{j+1} - \Phi_j| \leq \frac{4^{q-1}}{0.002|\mathcal{M}|} = O\left(\frac{4^q}{n}\right) = O(n^{0.47}/n) = O(1/n^{0.53}).$$

Now by Azuma's Inequality ([Theorem 4.2](#)) we have

$$\mathbb{P}\left[\left(\Phi_{\frac{41}{m}n} - \Phi_0\right) - \mathbb{E}\left[\Phi_{\frac{41}{m}n} - \Phi_0\right] \geq 0.12\right] \leq \exp\left(-\frac{0.12^2}{2 \sum_{j=0}^{\frac{41}{m}n} (O(n^{-0.53}))^2}\right) = \exp(-\Theta(n^{0.06})).$$

Thus $\phi_{(>-h)}$ will drop by at least the constant factor 412 with very high probability, giving $\phi_{(>-h)}(\text{start}_h + \frac{2}{c} + \frac{41}{m}) \leq 0.001|\mathcal{M}|$. \square

Lemma 6.16. *For each hour $h = 0, \dots, l$, by time $\text{start}_h + \frac{2}{c} + \frac{47}{m} \leq \text{end}_h$, the total mass $\mu \leq \rho_h |\mathcal{M}| 2^{-h}$ with very high probability. The constant $\rho_h = 0.1$ for all $h \leq l+5$, then the last few constants $\rho_{l-4} = 0.104$, $\rho_{l-3} = 0.13$, $\rho_{l-2} = 0.212$, $\rho_{l-1} = 0.408$, and $\rho_l = 0.808$.*

Proof. We start the argument at time $\text{start}_h + \frac{2}{c} + \frac{41}{m}$. Let $A = \{a : a.\text{opinion} = +1, a.\text{exponent} = -h\}$ and $B = \{b : b.\text{opinion} = -1, b.\text{exponent} = -h\}$. We will show that by time end_h , large number of cancel reactions happen between the agents in A and B to reduce the total mass. Inductively, by end_{h-1} we have total mass $\mu \leq \rho_{(h-1)} |\mathcal{M}| 2^{-h+1}$. We assume in the worst case (since we need the total mass to be small) that we start with the maximum possible amount of mass $\mu = \rho_{(h-1)} |\mathcal{M}| 2^{-h+1}$. We use the upper bound on the gap $g_h \leq \alpha |\mathcal{M}|$, where $\alpha = 0.8 \cdot 2^{h-l}$ and assume in the worst case (since larger gaps reduce the rate of cancelling reactions) the largest possible gap $\beta_+ - \beta_- = g_h 2^{-h} = \alpha |\mathcal{M}| 2^{-h}$. This gives majority mass $\beta_+ = (\rho_{(h-1)} + \frac{\alpha}{2}) |\mathcal{M}| 2^{-h}$ and minority mass $\beta_- = (\rho_{(h-1)} - \frac{\alpha}{2}) |\mathcal{M}| 2^{-h}$.

Since the minority is the limiting reactant in the cancelling reactions, we assume in the worst case the smallest possible minority count at exponent $-h$, where all mass outside of exponent $-h$ is the minority. Then the majority count at exponent $-h$ is $|A| = (\rho_{(h-1)} + \frac{\alpha}{2}) |\mathcal{M}|$. By [Lemma 6.15](#), the mass above exponent $-h$ is $\mu_{(>-h)}(\text{start}_h + t_2) \leq 0.001 |\mathcal{M}| 2^{-h+1}$. By [Lemma 6.11](#), the maximum count below exponent $-h$ is $0.0024 |\mathcal{M}|$, so the mass $\mu_{(<-h)} \leq 0.0024 |\mathcal{M}| 2^{-h-1}$. This leaves a minority count at exponent $-h$ of $|B| = (\rho_{(h-1)} - \frac{\alpha}{2} - 0.002 - 0.0012) |\mathcal{M}|$.

Now we will apply [Lemma 4.6](#) with $a = (\rho_{(h-1)} + \frac{\alpha}{2})m$ and $b = (\rho_{(h-1)} - \frac{\alpha}{2} - 0.0032)m$. First we consider the cases where $h \leq l+5$, where $\rho_h = \rho_{(h-1)} = 0.1$. There the bound on the gap $\alpha = 0.8 \cdot 2^{(l-5)-l} = 0.025$. This gives $a = 0.1125m$ and $b = 0.0843m$. By [Lemma 4.6](#), the time t to cancel a fraction $d = 0.05m$ from both a and b has mean

$$\mathbb{E}[t] \sim \frac{\ln(b) - \ln(a) - \ln(b-d) + \ln(a-d)}{2(a-b)} = \frac{\ln(0.0843) - \ln(0.1125) - \ln(0.0343) + \ln(0.0625)}{2(0.0282m)} < \frac{5.53}{m},$$

so $\mathbb{E}[t](1+\epsilon) < \frac{6}{m}$ and $t < \frac{6}{m}$ with very high probability. Cancelling this fraction d reduces the total mass by $2dn2^{-h} = 0.1 |\mathcal{M}| \cdot 2^{-h} = \rho_{(h-1)} |\mathcal{M}| 2^{-h+1} - \rho_h |\mathcal{M}| 2^{-h}$. Then the total mass is at most $\rho_h |\mathcal{M}| 2^{-h}$ by time $\text{start}_h + \frac{2}{c} + \frac{47}{m}$.

For the remaining levels $h = l-4, l-3, l-2, l-1, l$, we will repeat the above argument and calculation, but now the bound will change so $\rho_h \neq \rho_{(h-1)}$. First our bound on the gap α will double as we move down each level. This causes the worst case fractions a and b to be spread further apart. Then d , the largest fraction which will cancel from both sides within $\frac{6}{m}$ time with high probability, will be smaller. This gives the new mass bound ρ_h , since from the equation $\rho_{(h-1)} |\mathcal{M}| 2^{-h+1} - \rho_h |\mathcal{M}| 2^{-h} = 2dn2^{-h}$, we have $d = (\rho_{(h-1)} - \frac{\rho_h}{2})m$ and $\rho_h = 2(\rho_{(h-1)} - \frac{d}{m})$. This relative total mass bound ρ_h will be growing larger as h decreases. This is

to be expected, since we do in fact see the count of biased agents growing as they reach the final exponents before the count of \mathcal{O} agents runs out (see Fig. 4c and Fig. 4d, where the value $-l = -19$).

The following table shows the constants used in the computations at each of the steps $h = l-5, l-4, \dots, l$. The top row $h = l-5$ corresponds to the exact calculations above, then the following rows are the constants used in the same argument for lower levels.

$h = l-5$	$\alpha = 0.025$	$a = 0.1125m$	$b = 0.0843m$	$d = 0.05m$	$\mathbb{E}[t] < 5.53/m$	$\rho_{(l-5)} = 0.1$
$h = l-4$	$\alpha = 0.05$	$a = 0.125m$	$b = 0.0718m$	$d = 0.048m$	$\mathbb{E}[t] < 5.83/m$	$\rho_{(l-4)} = 0.104$
$h = l-3$	$\alpha = 0.1$	$a = 0.154m$	$b = 0.0508m$	$d = 0.039m$	$\mathbb{E}[t] < 5.66/m$	$\rho_{(l-3)} = 0.13$
$h = l-2$	$\alpha = 0.2$	$a = 0.23m$	$b = 0.0268m$	$d = 0.024m$	$\mathbb{E}[t] < 5.29/m$	$\rho_{(l-2)} = 0.212$
$h = l-1$	$\alpha = 0.4$	$a = 0.412m$	$b = 0.0088m$	$d = 0.008m$	$\mathbb{E}[t] < 2.95/m$	$\rho_{(l-1)} = 0.408$
$h = l$	$\alpha = 0.8$	$a = 0.808m$	$b = 0.0048m$	$d = 0.004m$	$\mathbb{E}[t] < 1.11/m$	$\rho_l = 0.808$

Note that for the last couple hours, the worst case for b is very small, so the amount of cancelling that happens is negligible, and the relative mass bound essentially doubles. We will see later that the minority count does in fact sharply decrease, so it is accurate that the rate of cancellation drops to zero, and the count of biased majority agents is roughly doubling for these last couple rounds. \square

6.3 End of Phase 3

Now that we have proven Theorem 6.12, we will finish analyzing separately the cases of an initial tie and an initial majority opinion.

In the case of a tie, the stronger bounds from Theorem 6.12 hold all through the final hour $h = L$. Thus at hour L we still have a constant fraction $0.77|\mathcal{M}|$ of \mathcal{O} agents. We now show that in the remaining $O(\log n)$ time in Phase 3, while the Clock agents with `hour` = h decrement their counter in Line 6, these \mathcal{O} agents can keep doing split reactions to bring all remaining biased agents down to exponent $-L$. This lets us now prove Theorem 6.1, the main result of the section for the case of a tie.

Theorem 6.1. *If the initial configuration was a tie with gap 0, then by the end of Phase 3, all biased agents have `exponent` = $-L$, with high probability $1 - O(1/n^2)$.*

Proof. We start by using Theorem 6.12, where at hour L , the total mass $\mu(\text{end}_L) \leq 0.1|\mathcal{M}|2^{-L}$. Thus the count of biased agents is at most $0.1|\mathcal{M}|$. We also have at least $0.77|\mathcal{M}|$ \mathcal{O} agents with `hour` = L (this count is actually slightly higher, as we could show almost all of the $0.9|\mathcal{M}|$ \mathcal{O} agents reach `hour` = L quickly by the same argument as Lemma 6.13, but this bound will suffice).

We use the potential $\phi_{(>-L)}$, where by Lemma 6.15, we have $\phi_{(>-L)}(\text{end}_0) \leq 0.001|\mathcal{M}|$. The goal is to now show that $\phi_{(>-L)} = 0$ within the $O(\log n)$ time it takes for any `counter` to hit 0 and trigger the move to the next phase. This proof proceeds identically to the proof of Lemma 6.14, where we had shown that $\mathbb{E}[\phi_{(>-L)}(t+1/n)|\phi_{(>-L)}(t)] \leq (1 - \frac{0.77m}{n})\phi_{(>-L)}(t)$. We again recursively bound the expected value of $\phi_{(>-L)}$ after an additional $16 \ln n$ time.

$$\begin{aligned} \mathbb{E}[\phi_{(>-L)}(\text{end}_L + 16 \ln n)|\phi_{(>-L)}(\text{end}_0)] &\leq \left(1 - \frac{0.77m}{n}\right)^{16n \ln n} \phi_{(>-L)}(\text{end}_0) \\ \mathbb{E}[\phi_{(>-L)}(\text{end}_L + 16 \ln n)] &\leq \exp(-0.77 \cdot 0.25 \cdot 16 \ln n) \cdot 0.001|\mathcal{M}| \\ &\leq n^{-3.08} 0.001mn = O(1/n^2). \end{aligned}$$

Again we conclude by Markov's Inequality that $\mathbb{P}[\phi_{(>-L)}(\text{end}_L + 16 \ln n) > 0] \leq O(1/n^2)$.

Finally we must argue that Phase 3 lasts at least until time $\text{end}_L + 16 \ln n$. We will bound the probability that a Clock agent can decrement its `counter` down to 0 before time $\text{end}_L + 16 \ln n$. Note the if statement on Line 6 will only decrement the counter when both Clock agents have reached the final minute. Even if in the worst case an agent was at the final minute at the beginning of the phase, until time end_{L-1} the count of other Clock agents with `minute` = kL is at most $0.1p|\mathcal{C}| = 0.001|\mathcal{C}|$. By Theorem 6.9, the time between consecutive hours is at most $\frac{3.86k}{c^2} < \frac{290}{c}$ using $k = 45$ and the bound $c > 0.2$ with very high probability from Lemma 5.2. Then the number of interactions until end_{L-1} is at most $\frac{290}{c}Ln$. In each interaction, the

probability of the Clock agent decrementing its counter is at most $2 \cdot 0.001c \cdot \frac{1}{n}$, so the expected number of decrements $\mathbb{E}[X] \leq 290L \cdot 0.002 \leq 0.58(\log_2(n)) \leq 0.84 \ln n$. Then by Chernoff bounds in [Theorem 4.1](#), with $\mu = 0.84 \ln n, \delta = 5$, we have

$$\mathbb{P}[X \geq 5.04 \ln n] \leq \mathbb{P}[X \geq (1 + \delta)\mu] \leq \exp\left(-\frac{\delta^2 \mu}{2 + \delta}\right) = \exp(-3 \ln n) = n^{-3}.$$

Then after time end_{L-1} , we assume that all Clock agents have $\text{minute} = kL$. In this case, the probability a given Clock agent decrements its counter is at most $\frac{2c}{n}$. Thus in $16n \ln n$ interactions, the expected number of decrements $\mathbb{E}[X] \leq 32c \ln n \leq 11 \ln n$, using $c < \frac{1}{3}$ from [Lemma 5.2](#). We again use [Theorem 4.1](#), with $\mu = 11 \ln n, \delta = 0.9$, to conclude

$$\mathbb{P}[X \geq 20.9 \ln n] \leq \mathbb{P}[X \geq (1 + \delta)\mu] \leq \exp\left(-\frac{\delta^2 \mu}{2 + \delta}\right) \leq \exp(-3.07 \ln n) \leq n^{-3}.$$

Then by union bound, Clock agents decrements their counter at most $5.04 \ln n + 20.9 \ln n < 26 \ln n$ times with high probability $1 - O(1/n^2)$. In [Phase 3](#), we initialize $\text{counter} \leftarrow c_3 \ln n$, so setting this counter constant $c_3 = 26$, we conclude that no agent moves to the next phase before time $\text{end}_L + 16 \ln n$ with high probability. This gives enough time to conclude that all biased agents have $\text{exponent} = -L$ by the end of [Phase 3](#). \square

Now we consider the case where there is an initial majority, which we have assumed without loss of generality was **A**. Here the argument of [Theorem 6.12](#) stops working at exponent $-l$, because of the substantial relative gap $0.4|\mathcal{M}| \leq g_l < 0.8|\mathcal{M}|$. Next we will need to show that the count of \mathcal{O} gets brought to almost 0 by split reactions with the majority. In order to show that the count of \mathcal{O} stays small however, we will need some way to bound future cancel reactions. To do this, we will now show that during the last few hours $l-4, \dots, l$, the minority mass β_- drops to a negligible amount. We will then use this tight bound to show the majority consumes most remaining \mathcal{O} agents, and finally that constraints on the majority mass β_+ will force the majority count to remain above 99% at exponents $-l, -(l+1), -(l+2)$ for all reachable configurations in the rest of [Phase 3](#).

We first show inductively that the minority mass becomes negligible during the last few hours $l-5, \dots, l$. We already showed a bound on the total mass in [Lemma 6.16](#), which used an upper bound on the gap. If we want the minority mass to be small, now the worst case is the smallest possible gap. Notice that just using the final mass bound $\mu < \rho_l |\mathcal{M}| 2^{-l} = 0.808 |\mathcal{M}| 2^{-l}$ in the worst case with the smallest possible gap $g_l = 0.4$ at exponent $-l$, would imply $\beta^+ \approx 0.6 |\mathcal{M}| 2^{-l}$ and $\beta^- \approx 0.2 |\mathcal{M}| 2^{-l}$. To get a much tighter upper bound on minority mass, we will make a similar inductive argument to [Lemma 6.16](#), now using the lower bound for each gap g_h , to show enough minority mass cancels at exponent $-h$ during each hour h .

Lemma 6.17. *For each hour $h = l-5, \dots, l$, by time end_h , the minority mass $\beta_- \leq \xi_h |\mathcal{M}| 2^{-h}$ with very high probability. The constants $\xi_{(l-5)} = 0.04375$, $\xi_{(l-4)} = 0.0375$, $\xi_{(l-3)} = 0.0267$, $\xi_{(l-2)} = 0.0145$, $\xi_{(l-1)} = 0.0056$, and $\xi_l = 0.004$.*

Proof. For the base case $h = l-5$, by [Lemma 6.16](#), the total mass $\mu(\text{end}_{(l-5)}) \leq 0.1 |\mathcal{M}| 2^{-l+5}$. The relative gap $0.0125 |\mathcal{M}| \leq g_{(l-5)} < 0.025 |\mathcal{M}|$, so now in the worst case (for minority mass being small), we assume the smallest gap $g_{(l-5)} = 0.0125 |\mathcal{M}|$ and largest total mass $\mu(\text{end}_{(l-5)}) = \beta_+ + \beta_- = 0.1 |\mathcal{M}| 2^{-l+5}$. Since the invariant gap $g = \beta_+ - \beta_- = g_{(l-5)} 2^{-l+5}$, we have $\beta_+ = 0.05625 |\mathcal{M}| 2^{-l+5}$ and $\beta_- = 0.04375 |\mathcal{M}| 2^{-l+5}$, giving an upper bound constant $\xi_{(l-5)} = 0.04375$.

Now we outline the inductive step, for each level $h = l-4, \dots, l$. We will show the constants for the first step $h = l-4$ here, and then list constants for the remaining steps in a table below. Let $A = \{a : a.\text{opinion} = +1, a.\text{exponent} = -h\}$ and $B = \{b : b.\text{opinion} = -1, b.\text{exponent} = -h\}$.

We start the argument at time $\text{start}_h + \frac{2}{c} + \frac{41}{m}$, and with the previous bound $\beta_-(\text{end}_{(h-1)}) = \xi_{(h-1)} |\mathcal{M}| 2^{-h+1} = 0.0875 |\mathcal{M}| 2^{-h}$. We will then assume in the worst case the smallest possible gap $g_h = \alpha |\mathcal{M}| = 0.4 \cdot 2^{h-l} |\mathcal{M}| = 0.025 |\mathcal{M}|$. This gives $\beta_+ = (2\xi_{(h-1)} + \alpha) |\mathcal{M}| 2^{-h}$. We then assume in the worst case that all mass allowed outside exponent $-h$ belongs to the minority (so doesn't reduce by cancelling), giving majority count

$|A| = (2\xi_{(h-1)} + \alpha)|\mathcal{M}|$. We also assume the maximum amount of mass outside exponent $-h$. By Lemma 6.11, before time end_h , the count below exponent $-h$ is at most $0.0024|\mathcal{M}|$, so the mass is at most $0.0024|\mathcal{M}|2^{-h-1}$. By Lemma 6.15, after time $\text{start}_h + \frac{2}{c} + \frac{41}{m}$, the mass above exponent $-h$ is $\mu_{(>-h)} \leq 0.001 \cdot 2^{-h+1}$. This leftover remaining mass at exponent $-h$ gives minority count $|B| = (2\xi_{(h-1)} - 0.0012 - 0.002)|\mathcal{M}|$.

Now we will apply Lemma 4.6 with fractions $a = (2\xi_{(h-1)} + \alpha)m = 0.1125m$ and $b = (2\xi_{(h-1)} - 0.0036)m = 0.0843m$. These match the first fractions a, b used in the proof of Lemma 6.16, and again we use that the time t to cancel a fraction $d = 0.05m$ from both a and b has mean

$$\mathbb{E}[t] \sim \frac{\ln(b) - \ln(a) - \ln(b-d) + \ln(a-d)}{2(a-b)} = \frac{\ln(0.0843) - \ln(0.1125) - \ln(0.0343) + \ln(0.0625)}{2(0.0282m)} < \frac{5.53}{m},$$

so $\mathbb{E}[t](1 + \epsilon) < \frac{6}{m}$ and $t < \frac{6}{m}$ with very high probability. Cancelling this fraction d reduces the minority mass by $dn2^{-h} = 0.05|\mathcal{M}| \cdot 2^{-h}$, giving a new minority mass $\beta_- = (2\xi_{(h-1)} - \frac{d}{m})|\mathcal{M}|2^{-h} = 0.0375|\mathcal{M}|2^{-h}$. Thus for $\xi_h = 0.0375$, the minority mass is at most $\xi_h|\mathcal{M}|2^{-h}$ by time $\text{start}_h + \frac{2}{c} + \frac{47}{m}$.

For the remaining levels $h = l-4, l-3, l-2, l-1, l$, we will repeat the above argument and calculation. The following table shows the constants used in the computations at each of the steps $h = l-4, l-3, \dots, l$. The top row $h = l-4$ corresponds to the exact calculations above, then the following rows are the constants used in the same argument for lower levels.

$h = l-4$	$\alpha = 0.025$	$a = 0.1125m$	$b = 0.0843m$	$d = 0.05m$	$\mathbb{E}[t] < 5.53/m$	$\xi_{(l-4)} = 0.0375$
$h = l-3$	$\alpha = 0.05$	$a = 0.125m$	$b = 0.0718m$	$d = 0.048m$	$\mathbb{E}[t] < 5.83/m$	$\xi_{(l-3)} = 0.0267$
$h = l-2$	$\alpha = 0.1$	$a = 0.154m$	$b = 0.0508m$	$d = 0.039m$	$\mathbb{E}[t] < 5.66/m$	$\xi_{(l-2)} = 0.0145$
$h = l-1$	$\alpha = 0.2$	$a = 0.23m$	$b = 0.0268m$	$d = 0.024m$	$\mathbb{E}[t] < 5.29/m$	$\xi_{(l-1)} = 0.0056$
$h = l$	$\alpha = 0.4$	$a = 0.412m$	$b = 0.0088m$	$d = 0.008m$	$\mathbb{E}[t] < 2.95/m$	$\xi_l = 0.004$

Note that the structure of the proof yields the same sequence of values a, b as the proof of Lemma 6.16, so we have repeated the exact same applications of Lemma 4.6. This time, however, we are using the cancelled fraction d to show the minority mass becomes very small. \square

Now Lemma 6.17 gives that minority mass $\beta_- \leq 0.004|\mathcal{M}|2^{-l}$ by time end_l . The rest of the argument will not try to reason about probabilistic guarantees on what the distribution looks like. Instead, we will make claims purely about reachability, and show that we get to a configuration where invariants force the count of majority agents to remain high in any reachable configuration using the transitions of Phase 3.

Lemma 6.18. *By the time $\text{end}_{(l+2)}$, the count of majority agents with exponent $\in \{-l, -(l+1), -(l+2)\}$ is at least $0.96|\mathcal{M}|$, with very high probability. Then in all reachable configurations where all agents are still in Phase 3, the count of majority agents with exponent $\in \{-l, -(l+1), -(l+2)\}$ is at least $0.92|\mathcal{M}|$.*

Proof. We argue by process of elimination. First we apply Lemma 6.11 to conclude that until time $\text{end}_{(l+2)}$, the count $|\{\mathcal{O} : \mathcal{O}.\text{hour} > l+2\} \cup \{b : b.\text{exponent} < -(l+2)\}| \leq 0.0024|\mathcal{M}|$ with very high probability.

Now Lemma 6.17 gives that minority mass $\beta_- \leq 0.004|\mathcal{M}|2^{-l}$ after time end_l . Then we can bound the maximum minority count at exponent $-(l+2)$ and above by $0.016|\mathcal{M}|$. In addition, these minority agents could eliminate more majority agents by cancel reactions. The count of agents they could cancel with is at most $0.016|\mathcal{M}|$.

This leaves at least $|\mathcal{M}| - 0.0024|\mathcal{M}| - 0.016|\mathcal{M}| - 0.016|\mathcal{M}| = 0.9656|\mathcal{M}|$ agents that are either majority or \mathcal{O} agents at time end_l . We will next show that in hour $l+2$, most of any remaining \mathcal{O} agents set $\text{hour} = l+2$ and then are consumed in split reactions with majority agents with exponent $> -(l+2)$.

The majority mass $\beta_+ \geq g = g_l 2^{-l} \geq 0.4|\mathcal{M}|2^{-l}$. By Lemma 6.15, at most $0.001|\mathcal{M}|2^{-l+1} = 0.002|\mathcal{M}|2^{-l}$ of β_+ can come from exponents above $-l$. Thus we need mass $0.398|\mathcal{M}|2^{-l}$ from majority agents at exponent $-l$ and below. Note that if the count $|\{a : a.\text{opinion} = +1, a.\text{exponent} \in \{-l, -(l+1)\}\}| \leq 0.19|\mathcal{M}|$, then the maximum majority mass could achieve would be $0.81|\mathcal{M}|$ with exponent $= -l-2$ and $0.19|\mathcal{M}|$ with exponent $= -l$, which gives $(\frac{0.81}{4} + 0.19)|\mathcal{M}|2^{-l} < 0.398|\mathcal{M}|2^{-l}$. This implies we must have a count of at least $0.19|\mathcal{M}|$ of majority agents at exponents $-(l+1)$ and $-l$ in all reachable configurations.

Next we show that all but at most $0.0036|\mathcal{M}|$ of these \mathcal{O} agents gets brought up to $\text{hour} = l+2$. We start at time start_{l+2} , when the count $|\mathcal{O}| \leq (0.9656 - 0.398)|\mathcal{M}| = 0.5676|\mathcal{M}|$, and the count of clock agents at

hour = $l + 2$ is at least $0.9|\mathcal{C}|$. Thus we can apply [Lemma 4.7](#), with $a = 0.9c$, $b_1 = 0.5676m$, $b_2 = 0.0056m$ to conclude that at most $0.0056|\mathcal{M}|$ of \mathcal{O} agents are left at the wrong hour after parallel time $t_1 < (1 + \epsilon)\mathbb{E}[t_1]$, where the expected time

$$\mathbb{E}[t_1] = \frac{\ln(b_1) - \ln(b_2)}{2a} = \frac{\ln(0.5676m) - \ln(0.0036m)}{1.8c} \leq \frac{2.82}{c}.$$

Thus with very high probability $t_1 < \frac{3}{c}$.

Next the at least $0.19|\mathcal{M}|$ majority agents at exponents $-(l + 1)$ and $-l$ will eliminate these \mathcal{O} agents by split reactions. We show that at most $0.001|\mathcal{M}|$ of \mathcal{O}_h agents are left. We again apply [Lemma 4.7](#), with $a = 0.19m$, $b_1 = 0.5676m$, $b_2 = 0.01m$ to conclude this takes at most parallel time $t_2 < (1 + \epsilon)\mathbb{E}[t_2]$, where

$$\mathbb{E}[t_2] = \frac{\ln(b_1) - \ln(b_2)}{2a} = \frac{\ln(0.5676m) - \ln(0.001m)}{0.38m} \leq \frac{16.69}{m}.$$

Thus with very high probability $t_2 < \frac{17}{m}$.

Account for these counts $0.0032|\mathcal{M}|$ and $0.001|\mathcal{M}|$ of leftover \mathcal{O} agents, along with the maximal count $0.001|\mathcal{M}|$ of biased agents above exponent $-l$, we are left with $0.9656|\mathcal{M}| - 0.0032|\mathcal{M}| - 0.001|\mathcal{M}| - 0.001|\mathcal{M}| = 0.96|\mathcal{M}|$ agents that must be majority agents with **exponent** $\in \{-l, -(l + 1), -(l + 2)\}$.

Now we argue that no reachable configurations can bring this count below $0.92|\mathcal{M}|$. We have already accounted for the maximum number of minority agents that can do cancel reactions at exponents $-l, -(l + 1), -(l + 2)$. Thus we only have to argue that no amount of split reactions can bring this count down by $0.04|\mathcal{M}|$. Note that reducing the count by $0.04|\mathcal{M}|$ will reduce the mass at these exponents by at least $0.04|\mathcal{M}|2^{-(l+2)}$. It will take at least $0.08|\mathcal{M}|$ agents below exponent $-l - 2$ to account for the same mass, for a total of $0.96|\mathcal{M}| - 0.04|\mathcal{M}| + 0.08|\mathcal{M}| = |\mathcal{M}|$. Thus it is not possible to move more than $0.04|\mathcal{M}|$ majority agents below exponent $-(l + 2)$ by split reactions, because there are not enough Main agents to account for the mass. \square

We can now use the results of [Lemma 6.18](#), [Lemma 6.15](#), and [Lemma 6.17](#), including the high probability requirements from all previous lemmas, to prove [Theorem 6.2](#), the main result for [Phase 3](#) in the non-tie case:

Theorem 6.2. *Assume the initial gap $|g| < 0.025|\mathcal{M}|$. Let the exponent $-l = \lfloor \log_2(\frac{g}{0.4|\mathcal{M}|}) \rfloor$. Let $i = \text{sign}(g)$ be the majority opinion and M be the set of all agents with **role** = Main, **opinion** = i , **exponent** $\in \{-l, -(l + 1), -(l + 2)\}$. Then at the end of [Phase 3](#), $|M| \geq 0.92|\mathcal{M}|$ with high probability $1 - O(1/n^2)$.*

In addition, the total mass above exponent $-l$ is $\mu_{(>-l)} = \sum_{a.\text{exponent} > -l} |a.\text{bias}| \leq 0.002|\mathcal{M}|2^{-l}$, and the total minority mass is $\beta_- = \sum_{a.\text{opinion} = -i} |a.\text{bias}| \leq 0.004|\mathcal{M}|2^{-l}$.

7 Analysis of final phases

7.1 Reserve agent Phases 5 and 6

We now consider the behavior of the Reserve agents in [Phase 5](#) and [Phase 6](#). We first show that with high probability they are also able to set their **sample** field during [Phase 5](#).

Lemma 7.1. *By the end of [Phase 5](#), all Reserve agents have **sample** $\neq \perp$, with high probability $1 - O(1/n^2)$.*

Proof. By [Theorem 6.2](#), we have at least $0.92|\mathcal{M}|$ majority agents with **exponent** $\in \{-l, -(l + 1), -(l + 2)\}$ by the end of [Phase 3](#). Now we can analyze the process of Reserve agents sampling the **exponent** of the first biased agent they encounter in [Phase 5](#) with [Lemma 4.7](#), where subpopulations $A = \{a : a.\text{role} = \text{Main}, a.\text{bias} = \pm 1\}$ with and $B = \mathcal{R}$, comprising fractions $a \geq 0.92m$ and $b_1 = r$. Then by [Lemma 4.7](#), all Reserve agents set their **sample** within time t , where $t \leq \frac{5 \ln n}{2 \cdot 0.92m}$ with high probability $1 - O(1/n^2)$. Thus for appropriate choice of counter constant c_5 , this will happen before the first Clock agent advances to the next phase. \square

We next show that the Reserve agents are able to bring the exponents of all biased agents down to at most $-l$ during [Phase 6](#).

Lemma 7.2. *By the end of [Phase 6](#), all biased agents have `exponent` $\leq -l$, with high probability $1 - O(1/n^2)$.*

Proof. [Theorem 6.2](#), the mass above exponent $-l$ is $\mu_{(>-l)} \leq 0.001|\mathcal{M}|2^{-l+1}$. We must show that all of this mass moves down to at least exponent $-l$ via split reactions with the Reserve agents ([Line 3](#) in [Phase 6](#)). We consider the following two cases based on whether the size of $A_{-l} = \{a : a.\text{opinion} = +1, a.\text{exponent} = -l\}$.

In the first case, $|A| > 0.19|\mathcal{M}|$ and we will show all agents get brought down to `exponent` $\leq -l$. Because of the sampling process in [Phase 5](#), the count of $R_{-l} = \{r : r.\text{role} = \text{Reserve}, r.\text{sample} = -l\}$ has size at least $|R_{-l}| \geq 0.18|\mathcal{R}|$ with very high probability, by standard Chernoff bounds. If all the agents above exponent $-l$ split down to exponent $-l$, they would have count at most $0.002|\mathcal{M}|$, potentially all coming out of the initial count of R_{-l} . Thus for the entirety of [Phase 6](#), we have

$$|R_{-l}| \geq 0.18|\mathcal{R}| - 0.002|\mathcal{M}| = n(0.18r - 0.002m) \geq n(0.18 \cdot 0.24 - 0.002) \geq 0.04n,$$

using the very high probability bound $r > 0.24$ from [Lemma 5.2](#).

In the second case $|A| \leq 0.19|\mathcal{M}|$. Now we cannot make guarantees on the size $|R_{-l}|$, so we will instead reason about $A_{-(l+1)} = \{a : a.\text{opinion} = +1, a.\text{exponent} = -(l+1)\}$ and $R_{-(l+1)} = \{r : r.\text{role} = \text{Reserve}, r.\text{sample} = -(l+1)\}$.

We first show that $|A_{-(l+1)}| > 0.59|\mathcal{M}|$. Recall by definition of l and g_l that the majority mass $\beta_+ \geq g_l 2^{-l} \geq 0.4|\mathcal{M}|2^{-l}$. At most $0.001|\mathcal{M}|2^{-l+1}$ can come from exponents above $-l$ and at most $0.19|\mathcal{M}|2^{-l}$ comes from exponent $-l$. Thus there must be at least mass $(0.398 - 0.19)|\mathcal{M}|2^{-l}$ from $A_{-(l+1)}$ and the exponents below. If $|A_{-(l+1)}| = 0.59|\mathcal{M}|$, then even putting all $0.41|\mathcal{M}|$ remaining Main agents at `exponent` $= -(l+2)$ would only give mass $0.59|\mathcal{M}|2^{-(l+1)} + 0.41|\mathcal{M}|2^{-(l+2)} = 0.3975|\mathcal{M}|2^{-l}$. Thus we require $|A_{-(l+1)}| > 0.59|\mathcal{M}|$. Again by standard Chernoff bounds from the sampling process of [Phase 5](#), this implies the initial size of $|R_{-(l+1)}| \geq 0.58|\mathcal{R}|$.

If all the agents above exponent $-l$ split down to exponent $-(l+1)$, they would have count at most $0.004|\mathcal{M}|$, potentially all coming out of the initial count of R_{-l} . In addition, we can lose count $|A| \leq 0.19|\mathcal{M}|$ from R_{-l} from additional split reactions. This implies that for the entirety of [Phase 6](#), we have count at least

$$\begin{aligned} |R_{-(l+1)}| &\geq 0.58|\mathcal{R}| - 0.004|\mathcal{M}| - 0.19|\mathcal{M}| = n(0.58r - 0.004m - 0.19m) \\ &\geq n(0.58 \cdot 0.24 - (0.004 + 0.19) \cdot 0.51) \geq 0.04n, \end{aligned}$$

using the very high probability bounds $r > 0.24$ and $m < 0.51$ from [Lemma 5.2](#).

Thus in both cases we maintain a count of at least $0.04n$ Reserve agents at level $-l$ or $-(l+1)$. For an agent a with `a.exponent` $> -l$, the probability that in a given interaction agent a does a split reaction ([Line 3](#) in [Phase 6](#)) with an agent in $R_{-l} \cup R_{-(l+1)}$ is at least $\frac{2 \cdot 0.04n \cdot 1}{n^2} = \frac{0.08}{n}$. Now we proceed as in the proofs of [Lemma 6.14](#) and [Theorem 6.1](#), analyzing the potential $\phi_{(>-l)}$ to show it hits 0 in $O(\log n)$ time.

Recall that for each biased agent a with `a.exponent` $= -i > -l$ and local potential $\phi_a = 4^{-i+l-1}$, a split reaction changes the potential by $\Delta\phi_a \leq 4^{l-1}(2 \cdot 4^{-i-1} - 4^{-i}) = 4^{l-1}(-\frac{1}{2}4^{-i}) = -\frac{1}{2}\phi_a$. Then in each interaction at parallel time t , the expected change in the potential

$$\begin{aligned} \mathbb{E}[\phi_{(>-l)}(t+1/n) - \phi_{(>-l)}(t)] &\leq \sum_{a.\text{exponent} > -l} \mathbb{P}[a \text{ splits}] \cdot \Delta\phi_a \\ &\leq \sum_{a.\text{exponent} > -l} \frac{0.08}{n} \cdot -\frac{1}{2}\phi_a = -\frac{0.04}{n}\phi_{(>-l)}(t). \end{aligned}$$

Thus we have $\mathbb{E}[\phi_{(>-l)}(t+1/n) | \phi_{(>-l)}(t)] \leq (1 - \frac{0.04}{n})\phi_{(>-l)}(t)$.

When we begin the argument at time t_{start} at the beginning of [Phase 6](#), we start with $\phi_{(>-l)}(t_{\text{start}}) \leq 0.001|\mathcal{M}|$ from the final iteration of [Lemma 6.15](#). We will wait until time $t_{\text{end}} = t_{\text{start}} + 75 \ln n$, and now recursively bound the potential after these $75n \ln n$ interactions:

$$\begin{aligned}\mathbb{E}[\phi_{(>-l)}(t_{\text{end}})|\phi_{(>-l)}(t_{\text{start}})] &\leq \left(1 - \frac{0.04}{n}\right)^{75n \ln n} \phi_{(>-l)}(t_{\text{start}}) \\ \mathbb{E}[\phi_{(>-h+q)}(t_{\text{end}})] &\leq \exp(-0.04 \cdot 75 \ln n) \cdot 0.001|\mathcal{M}| \\ &\leq n^{-3} \cdot 0.001mn = O(n^{-2})\end{aligned}$$

Finally, since $\phi_{(>-l)}$ takes nonnegative integer values, we can apply Markov's Inequality to conclude $\mathbb{P}[\phi_{(>-l)}(t_{\text{end}}) > 0] = O(1/n^2)$. So after $75 \ln n$ time in [Phase 6](#), all biased agents have **exponent** $\leq -l$, with high probability. For appropriate choice of counter constant c_6 , this happens before any Clock agent advances to the next phase. \square

Recall M is the set of all majority agents with **exponent** $\in \{-l, -(l+1), -(l+2)\}$, where $|M| \geq 0.92|\mathcal{M}|$ at the end of [Phase 3](#) by [Theorem 6.2](#). We finally observe that after [Phase 6](#), M is still large.

Lemma 7.3. *At the end of [Phase 6](#), $|M| \geq 0.87|\mathcal{M}|$ with very high probability.*

Proof. By [Theorem 6.2](#), $|M| \geq 0.92|\mathcal{M}|$ at the end of [Phase 3](#), so the count of all other main agents is at most $0.08|\mathcal{M}|$. By standard Chernoff bounds on the sampling process in [Phase 5](#), the count

$$|\{r : r.\text{role} = \text{Reserve}, r.\text{sample} \notin \{-l, -(l+1), -(l+2)\}\}| \leq 0.09|\mathcal{R}| \leq 0.05|\mathcal{M}|.$$

Split reactions that consume these Reserve agents are the only way to bring agents out of the set M . Thus at the end of [Phase 6](#), the count is still at least

$$|M| \geq 0.92|\mathcal{M}| - 0.09|\mathcal{R}| \geq 0.87|\mathcal{M}|,$$

using $|\mathcal{R}| < \frac{5}{9}|\mathcal{M}|$ with very high probability by [Lemma 5.2](#). \square

7.2 Minority elimination Phases 7 and 8

Next we argue that in [Phase 7](#), the agents in M are able to eliminate all minority agents with **exponent** $\in \{-l, -(l+1), -(l+2)\}$. Note that these minority agents are able to do a cancel reaction with any agent in M , since by design [Phase 7](#) allows reactions between agents with an exponent gap of at most 2. We first argue that $|M|$ must stay large through the entirety of [Phase 7](#):

Lemma 7.4. *At the end of [Phase 7](#), $|M| \geq 0.8|\mathcal{M}|$.*

Proof. We use the bound on minority mass $\beta_- \leq 0.004|\mathcal{M}|2^{-l}$ from [Theorem 6.2](#). This implies the minority count at exponent $-(l+4)$ and above is at most $0.064|\mathcal{M}|$, since a $0.064|\mathcal{M}| \cdot 2^{-(l+4)} \geq \beta_-$. Note that cancelling with these minority agents is the only way for a majority agent in M to change its state in [Phase 7](#). Using the previous bound from [Lemma 7.3](#), the count of M can decrease at most to $|M| = 0.87|\mathcal{M}| - 0.064|\mathcal{M}| \geq 0.8|\mathcal{M}|$. \square

Now we argue that these agents in M eliminate all high exponent minority agents.

Lemma 7.5. *At the end of [Phase 7](#), all minority agents have **exponent** $< -(l+2)$, with high probability $1 - O(1/n^2)$.*

Proof. From [Lemma 7.2](#), all minority agents already have **exponent** $\leq -l$. Let $B_{-l}, B_{-(l+1)}, B_{-(l+2)}$ be the sets of minority agents with **exponent** $= -l, -(l+1), -(l+2)$, respectively. We argue successively that $|B_{-l}| = 0$, then $|B_{-(l+1)}|$, then $|B_{-(l+2)}| = 0$.

The initial bound on minority mass $\beta_- \leq 0.004|\mathcal{M}|2^{-l}$ from [Theorem 6.2](#) implies that initially $|B_{-l}| \leq 0.004|\mathcal{M}|$. Note that an interaction with any agent in M will bring remove an agent from B_{-l} via one of the [Phase 7](#) cancel reactions. Using the bound $|M| \geq 0.8|\mathcal{M}|$ during the entire phase from [Lemma 7.4](#), we

can use [Lemma 4.7](#) to bound the time for $|B_{-l}| = 0$. We have constants $a = 0.8m$ and $b_1 = 0.004m$. Then by [Lemma 7.4](#), with high probability $1 - O(1/n^2)$, the time t_1 to eliminate the count of B_{-l} is at most $t_1 \leq \frac{5 \ln n}{2(0.8m - 0.004m)} \leq 6.41 \ln n$, using the bound $m \geq 0.49$ from [Lemma 5.2](#).

Next we wait to eliminate the count of $B_{-(l+1)}$, by a similar argument. Initially the count is at most $|B_{-(l+2)}| \leq 0.008|\mathcal{M}|$, and this will all cancel in at most time t_2 . By [Lemma 7.4](#), with high probability, $t_2 \leq \frac{5 \ln n}{2(0.8m - 0.008m)} \leq 6.45 \ln n$. The same argument for $B_{-(l+2)}$, initially $|B_{-(l+2)}| \leq 0.016|\mathcal{M}|$, gives $t_3 \leq \frac{5 \ln n}{2(0.8m - 0.016m)} \leq 6.51 \ln n$.

Thus the entire process requires time at most $t_1 + t_2 + t_3 < 20 \ln n$. So for appropriate choice of counter constant c_7 , this will happen before the first Clock agent advances to the next phase. \square

Now we will prove that the remaining $0.8|\mathcal{M}|$ majority agents in M are able to eliminate all remaining minority agents in [Phase 8](#).

Lemma 7.6. *At the end of [Phase 8](#), there are no more minority agents, with high probability $1 - O(1/n^2)$.*

Proof. From [Lemma 7.4](#) and [Lemma 7.5](#), by the end of [Phase 7](#), we have $|M| \geq 0.8|\mathcal{M}|$ and all minority agents have `exponent` $< -(l + 2)$. We assume in the worst case all $0.2|\mathcal{M}|$ other Main agents have the minority opinion. Note by the consumption reaction in [Phase 8](#), every minority agent will be eliminated by an agent in M that still has `full` = `False`. Thus we can apply [Lemma 4.6](#), with $a = 0.8m$ and $b = 0.2m$, to conclude that all remaining minority agents are eliminated in time t , where $t \leq \frac{5 \ln n}{2(0.8m - 0.2m)} \leq 8.5 \ln n$, using $m \geq 0.495$ with very high probability from [Lemma 5.2](#). So for appropriate choice of counter constant c_8 , this will happen before the first Clock agent advances to the next phase. \square

7.3 Fast Stabilization

We next give a time bound for the stable backup:

Lemma 7.7. *The 6-state protocol in [Phase 10](#) stably computes majority in $O(n \log n)$ stabilization time, both in expectation and with high probability.*

Proof. Note that agents in the initial state with `active` = `True` and `output` $\in \{A, B\}$ can only change their state by a cancel reaction in [Line 3](#) with another active agent with the opposite opinion.

First we consider the case where one opinion, without loss of generality `A`, is the majority. We first wait for all active `B` agents to meet an active `A` agent and changed their state to active `T` via [Line 3](#). This is modelled by the ‘‘cancel reaction’’ process described in [Lemma 4.6](#), taking expected $O(n)$ time and $O(n \log n)$ time with high probability. At this point, there are no active `B` agents and at least one active `A` agent remaining. We next wait for this active `A` agent to interact with all remaining active `T` agents, which will then become passive via [Line 5](#). This takes $O(n \log n)$ time in expectation and with high probability by a standard coupon-collector argument, after which point there are only active `A` agents and passive agents. Finally, we wait for these active `A` agents to convert all passive agents to `output` = `A` via [Line 7](#), taking another $O(n \log n)$ time by the same coupon-collector argument.

Next we consider the case where the input distribution is a tie. We first wait for all $n/2$ pairs of active `A` and `B` agents to cancel via [Line 3](#), taking $O(n)$ time in expectation and $O(n \log n)$ time with probability. Consider the last such interaction. At this point, those two agents have become active `T` agents, and there are no active `A` or `B` agents left. Thus after interacting with one of the `T` agents, all passive agents will output `T` via [Line 7](#). This takes $O(n \log n)$ time in expectation and with high probability by the same coupon-collector argument. \square

We are now able to combine all the results from previous sections to prove [Theorem 3.1](#), giving guarantees on the behavior of the protocol:

Theorem 3.1. *There is a nonuniform population protocol [Nonuniform Majority](#), using $O(\log n)$ states, that stably computes majority in $O(\log n)$ stabilization time, both in expectation and with high probability.*

Proof. We first argue that with high probability $1 - O(1/n^2)$, the protocol **Nonuniform Majority** stabilizes to the correct output in $O(\log n)$ time. We consider three cases based on the size of the initial gap $|g|$. For the majority cases where $|g| > 0$, we assume without loss of generality $g > 0$, so **A** is the majority.

If $|g| \geq 0.025|\mathcal{M}|$, then by [Lemma 5.3](#), we stabilize in **Phase 2** to the correct output with high probability $1 - O(1/n^2)$.

If $0 < |g| < 0.025|\mathcal{M}|$, then by [Theorem 6.2](#), with high probability $1 - O(1/n^2)$, we end **Phase 3** with at least 92% of **Main** agents in the set M , with the majority opinion and **exponent** $\in \{-l, -(l+1), -(l+2)\}$. Then after **Phase 5** and **Phase 6**, all biased agents have **exponent** $\leq -l$ with high probability $1 - O(1/n^2)$ by [Lemma 7.2](#). Then after **Phase 7** and **Phase 7**, there are no more minority agents with high probability $1 - O(1/n^2)$ by [Lemma 7.6](#). Thus in **Phase 9**, the majority opinion $+1$ will spread to all agents by epidemic, and we reach a stable correct configuration where all agents have **opinions** $= \{0, +1\}$ and **output** $= \mathbf{A}$.

If $g = 0$, then by [Theorem 6.1](#), with high probability $1 - O(1/n^2)$, we end **Phase 3** with all biased agents at **exponent** $= -L$. Thus in **Phase 4**, we reach a stable correct configuration where all agents have **output** $= \mathbf{T}$, and there are no agents with **exponent** $> -L$ to increment the phase.

Next we justify that **Nonuniform Majority** stably computes majority, since it is always possible to reach a stable correct configuration. By [Lemma 5.2](#), we must produce at least 2 **Clock** agents by the end of **Phase 0**. Thus we must eventually advance through all timed phases **0, 1, 3, 5, 6, 7, 8** using the **counter** field. Also all agents have left the initial role **Role_{MCR}**, otherwise the **Init** of **Phase 1** will trigger all agents to move to **Phase 10** by epidemic. Thus the sum of **bias** across all **Main** agents is the initial gap g , and this key invariant is maintained through all phases.

Using this invariant, if the agents stabilize in **Phase 2**, they have a consensus on their **output**, the sign of their **bias**, and this consensus must be the sign of the sum of the biases, giving the sign of the initial gap. If the agents stabilize in **Phase 4**, all $|\mathbf{bias}| \leq \frac{1}{2^L} \leq \frac{1}{n}$. This implies the gap $|g| \leq |\mathcal{M}| \cdot \frac{1}{n} < 1$, so the gap $g = 0$ and the agents are correctly outputting **T**. Finally, note that in **Phase 8**, when the agents set the field **full** $= \mathbf{True}$, they can no longer actually store in memory the true bias they are holding. For example an agent with **opinion** $= +1$, **exponent** $= -i$, **full** $= \mathbf{True}$, is actually representing the interval $\frac{1}{2^{i+1}} \leq \mathbf{bias} < \frac{1}{2^i}$. Then we still have the guarantee that if we reach stable consensus in **Phase 9**, then this consensus is the sign of the sum of the biases and is thus correct. The final case is that we do not stabilize here, and then move to **Phase 10**, where they agents eventually stabilize to the correct output.

We finally justify that the expected stabilization time is $O(n \log n)$. By [Lemma 7.7](#), the stable backup **Phase 10** will stabilize in expected $O(n \log n)$ time. Note that the high probability guarantees are all at least $1 - O(1/n^2)$, so the time for the stable backup contributes at most $O\left(\frac{n \log n}{n^2}\right) = o(1)$ to the total expected time. Now by [Lemma 5.2](#), with very high probability we have at least $|\mathcal{C}| \geq 0.24n$ **Clock** agents. In this case, the time upper bounds of [Theorem 6.9](#) and upper bounds on the **counter** times using standard Chernoff bound, imply that every timed phase lasts expected $O(\log n)$ time. If we do not stabilize in the untimed phases, then we also pass through by epidemic expected $O(\log n)$ time. Thus we either stabilize or reach the backup **Phase 10** in expected $O(\log n)$ time. There is a final very low probability case that $|\mathbf{Clock}| = o(n)$, but we must at least have $|\mathbf{Clock}| = 2$. Even in this worst case, the time upper bounds of [Theorem 6.9](#) and all counter rounds are at most polynomial in n , whereas the low probability of such a small **Clock** population is smaller than any polynomial. Thus this event adds a negligible contribution, and we conclude the total expected stabilization time is $O(\log n)$. \square

8 Uniform, stable protocols for majority using more states

The algorithm described in [Section 3](#) is *nonuniform*: the set of transitions used for a population of size n depends on the value $\lceil \log n \rceil$. A uniform protocol [\[25, 29, 30\]](#) is a single set of transitions that can be used in any population size. Since it is “uniformly specified”, the transition function is formally defined by a linear-space¹⁶ Turing machine, where the space bound is the maximum space needed to read and write the input

¹⁶That is, the Turing machine is allowed to use a bit more than the space necessary simply to read and write the input and output states, but not significantly more (constant-factor). This allows it to do simple operations, such as integer multiplication,

and output states. The original model [9] used $O(1)$ states and transitions for all n and so was automatically uniform, but many recent $\omega(1)$ state protocols are nonuniform. With the exception of the uniform variant in [17], all $\omega(1)$ state stable majority protocols are nonuniform [4, 5, 8, 14, 16, 21]. The uniform variant in [17] has a tradeoff parameter s that, when set to $O(1)$ to minimize the states, uses $O(\log n \log \log n)$ states and $O(\log^2 n)$ time.

In this section we show that there is a way to make [Nonuniform Majority](#) in [Section 3](#) uniform, retaining the $O(\log n)$ time bound, but the expected number of states increases to $\Theta(\log n \log \log n)$.¹⁷

8.1 Main idea of $O(\log n \log \log n)$ state uniform majority (not handling ties)

Since [Nonuniform Majority](#) uses the hard-coded value $L = \lceil \log n \rceil$, to make the algorithm uniform, we require a way to estimate $\log n$ and store it in a field L (called `logn` below) of each agent. For correctness and speed, it is only required that `logn` be within a constant multiplicative factor of $\log n$.

Gąsieniec and Stachowiak [37] show a uniform $O(\log \log n)$ state, $O(\log n)$ time protocol (both bounds in expectation and with high probability) that computes and stores in each agent a value $\ell \in \mathbb{N}^+$ that, with high probability, is within additive constant $O(1)$ of $\lceil \log \log n \rceil$ (in particular, WHP $\ell \geq \lceil \log \log n \rceil - 3$ [17, Lemma 8]), so $2^\ell = \Theta(\log n)$. (This is the so-called *junta election* protocol used as a subroutine for a subsequent leader election protocol.) Furthermore, agents approach this estimate from below, propagating the maximum estimate by epidemic $\ell', \ell \rightarrow \ell, \ell$ if $\ell' < \ell$. This gives an elegant way to compose the size estimation with a subsequent nonuniform protocol \mathcal{P} that requires the estimate: agents store their estimate `logn` of $\log n$ and use it in \mathcal{P} . Whenever an agent’s estimate `logn` updates—always getting larger—it simply resets \mathcal{P} , i.e., sets the entire state of \mathcal{P} to its initial state. We can then reason as though all agents actually started with their final convergent value of `logn`.¹⁸

To make our protocol uniform, but remove its correctness in the case of a tie, as we explain below, all agents conduct this size estimation, stored in the field `logn`, in parallel with the majority protocol \mathcal{P} of [Section 3](#). Each agent resets \mathcal{P} to its initial state whenever `logn` updates. This gives the stated $O(\log n)$ time bound and $O(\log n \log \log n)$ state bound. Note that in Phase 0, agents count from `counter` = $\Theta(\log n)$ down to 0. It is sufficient to set the constant in the Θ sufficiently large that all agents with high probability receive by epidemic the convergent final value of `logn` significantly before any agent with the same convergent estimate counts down to 0.

Acknowledging that, with small probability, the estimate of $\log n$ could be too low for [Phase 4](#) to be correct, we simply remove [Phase 4](#) and do not attempt to detect ties. So if we permit undefined behavior in the case of a tie (as many existing fast majority protocols do), then this modification of the algorithm otherwise retains stably correct, $O(\log n)$ time behavior, while increasing the state complexity to $O(\log n \log \log n)$.

8.2 How to stably compute ties

With low but positive probability, the estimate of $\log n$ could be too small. For most phases of the algorithm, this would merely amplify the probability of error events (e.g., [Phase 1](#) doesn’t last long enough for agents to converge on biases $\{-1, 0, +1\}$) that later phases are designed to handle. However, the correctness of [Phase 4](#) (which detects ties) requires agents to have split through at least $\log n$ exponents in [Phase 3](#). Since the population-wide bias doubles each time the whole population splits down one exponent, the only way for the whole population to split through $\log n$ exponents is for there to be a tie (i.e., the population-wide bias is 0, so can double unboundedly many times). In this one part of the algorithm, for correctness we require

that require more than constant space, without “cheating” by allowing the internal memory usage of the Turing machine to vastly exceed that required to represent the states.

¹⁷We say “expected” because this protocol has a non-zero probability of using an arbitrarily large number of states. The number of states will be $O(\log n \log \log n)$ in expectation and with high probability.

¹⁸One might hope for a stronger form of composition, in which the size estimation *terminates*, i.e., sets an initially `False` Boolean flag to `True` only if the size estimation has converged, in order to simply prohibit the downstream protocol \mathcal{P} from starting with an incorrect estimate of $\log n$. However, when both states A and B are initially $\Omega(n)$, this turns out to be impossible; $\Omega(n)$ agents will necessarily set the flag to `True` in $O(1)$ time, long before the $O(\log n)$ time required for the size estimation to converge [29, Theorem 4.1].

the estimate to be at least $\log n$ with probability 1. (It can be much greater than $\log n$ without affecting correctness; an overestimate merely slows down the algorithm.)

To correct this error, we will introduce a stable backup size estimate, to be done in [Phase 4](#). Note that there are only a constant number of states with `phase = 4`: `Clock` agents do not store a counter in this phase, and `Main` agents that stay in this phase must have `bias` $\in \{0, \pm \frac{1}{2L}\}$. Thus we can use an additional $\Theta(\log n)$ states for the agents that are currently in `phase = 4` to stably estimate the population size. If they detect that their estimate of L was too small, they simply go to the stable backup [Phase 10](#).

Stable computation of $\lfloor \log n \rfloor$. The stable computation of $\log n$ has all agents start in state L_0 , where the subscript represents the agent’s estimate of $\lfloor \log n \rfloor$. We have the following transitions: for each $i \in \mathbb{N}$, $L_i, L_i \rightarrow L_{i+1}, F_{i+1}$ and, for each $0 \leq j < i$, $F_i, F_j \rightarrow F_i, F_i$. Among the agents in state L_i , half make it to state L_{i+1} , reaching a maximum of L_k at $k = \lfloor \log_2 n \rfloor$.¹⁹ All remaining F agents receive the maximum value k by epidemic. A very similar protocol was analyzed in [20, Lemma 12]. We give a quick analysis below for the sake of self-containment.

Time analysis of stable computation of $\lfloor \log n \rfloor$. The expected time is $\Theta(n \log n)$. For the upper bound, for each i , if j is the count of L_i agents, then the probability of a $L_i, L_i \rightarrow \dots$ reaction is $O(j^2/n^2)$, so for any possible starting count k of L_i agents, it takes expected time at most $O\left(\frac{1}{n} \sum_{j=1}^k \frac{n^2}{j^2}\right) = O(n)$ for the count of L_i agents to get from k to 0 or 1. Assuming in the worst case that no $L_{i+1}, L_{i+1} \rightarrow \dots$ reaction happens until all $L_i, L_i \rightarrow \dots$ reactions complete, then summing over $\lfloor \log n \rfloor$ values of i gives the $O(n \log n)$ time upper bound to converge on one L_k agent, where $k = \lfloor \log n \rfloor$. It takes additional $O(\log n)$ time for the F agents to propagate k by epidemic.

For the lower bound, observe that for each value of i , the *last* reaction $L_i, L_i \rightarrow \dots$ occurs when the count of L_i is either 2 (and will increase at most once more), or 3 (and will never increase again). This takes expected time $\Theta(n)$. Since the count of L_i will increase at most once more, at that time (just before the last $L_i, L_i \rightarrow \dots$ reaction), the count of L_{i-1} is at most 3, otherwise *two* more L_i ’s could be produced, and this would not be the last $L_i, L_i \rightarrow \dots$ reaction. Thus, to produce the L_i needed for this last $L_i, L_i \rightarrow \dots$ reaction, the previous $L_{i-1}, L_{i-1} \rightarrow L_i, F_i$ reaction occurs when the count of L_{i-1} is at most 5, also taking $\Theta(n)$ time. Thus the final reaction at each level takes time $\Theta(n)$ and depends on a reaction at the previous level that also takes time $\Theta(n)$, so summing across $\lfloor \log n \rfloor$ levels gives $\Omega(n \log n)$ completion time.

It is shown in [20, Lemma 12] that the time is $O(n \log^2 n)$ with high probability, i.e., slower than the expected time by a $\log n$ factor. Since the probability is $O(1/n^2)$ that we need to rely on this slow backup, even this larger time bound contributes negligibly to our total expected time.

8.3 Challenges in creating $O(\log n)$ state uniform algorithm

It is worth discussing some ideas for adjusting the uniform protocol described above to attempt to reduce its space complexity to $O(\log n)$ states. The primary challenge is to enable the population size n to be estimated without storing the estimate in any agent that participates in the main algorithm (i.e., the agent has role `Main`, `Clock`, or `Reserve`). If agents in the main algorithm do not store the size, then by [29, Theorem 4.1] they will provably go haywire initially, with agents in every phase, totally unsynchronized, and require the size estimating agents to reset them after having converged on a size estimate that is $\Omega(\log n)$.

The following method would let the `Size` agents reset main algorithm agents, without actually having to store an estimate of $\log n$ in the algorithm agents, but it only works with high probability. The size estimating agents could start a junta-driven clock as in [37], which is reset whenever they update their size estimate. Then, as long as there are $\Omega(n)$ `Size` agents, they could for a phase timed to last $\Theta(\log n)$ time, reset the algorithm agents by *direct* communication (instead of by epidemic). This could put the algorithm agents in a quiescent state where they do not interact with each other, but merely wait for the `Size` agents to

¹⁹More generally, the unique stable configuration encodes the full population size n in binary in the following distributed way: for each position i of a 1 in the binary expansion of n , there is one agent L_i . Thus, these remaining agents lack the space to participate in propagating the value $k = \lfloor \log n \rfloor$ by epidemic, but there are $\Omega(n)$ followers F to complete the epidemic quickly.

exit the resetting phase, indicating that the algorithm agents are able to start interacting again. Since there are $\Omega(n)$ size-estimating agents, each non-size-estimating agent will encounter at least one of them with high probability in $O(\log n)$ time.

The problem is that the reset signal is not guaranteed to reach every algorithm agent. There is some small chance that a Main agent with a bias different from its input does not encounter a Size agent in the resetting phase, so is never reset. The algorithm from that point on could reach an incorrect result when the agent interacts with properly reset agents, since the sum of biases across the population has changed. In our algorithm, by “labeling” each reset with the value $\log n$, we ensure that no matter what states the algorithm agents find themselves in during the initial chaos before size computation converges, every one of them is guaranteed to be reset one last time with the same value of $\log n$. The high-probability resetting described above seems like a strategy that could work to create a high probability uniform protocol using $O(\log n)$ time and states, though we have not thoroughly explored the possibility.

But it seems difficult to achieve probability-1 correctness using the technique of “reset the whole majority algorithm whenever the size estimate updates,” without multiplying the state complexity by the number of possible values of $\log n$. Since we did not need $\lceil \log n \rceil$ exactly, but only a value that is $\Theta(\log n)$, we paid only $\Theta(\log \log n)$ multiplicative overhead for the size estimate, but it’s not straightforward to see how to avoid this using the resetting technique. Of course, one could imagine that the savings could come from reducing the state complexity of the main majority-computing agents in the nonuniform algorithm. However, reducing the nonuniform algorithm’s state complexity to below the $\Omega(\log n)$ lower bound of [5] would provably require the algorithm to be not monotonic or not output dominant. (See Section 9 for a discussion of those concepts.) Another possible approach is to intertwine more carefully the logic of the majority algorithm with the size estimation, to more gracefully handle size estimate updates without needing to reset the entire majority algorithm.

9 Conclusion

There are two major open problems remaining concerning the majority problem for population protocols.

Uniform $O(\log n)$ -time, $O(\log n)$ -state majority protocol Our main $O(\log n)$ state protocol, described in Section 3, is *nonuniform*: all agents have the value $\lceil \log n \rceil$ encoded in their transition function. The uniform version of our protocol described in Section 8 uses $O(\log n \log \log n)$ states. It remains open to find a uniform protocol that uses $O(\log n)$ time and states.

Unconditional $\Omega(\log n)$ state lower bound for stable majority protocols The lower bound of $\Omega(\log n)$ states for (roughly) sublinear time majority protocols shown by Alistair, Aspnes, and Gelashvili [5] applies only to stable protocols satisfying two conditions: *monotonicity* and *output dominance*.

Recall that a *uniform* protocol is one where a single set of transitions works for all population sizes; nonuniform protocols typically violate this by having an estimate of the population size (e.g., the integer $\lceil \log n \rceil$) embedded in the transition function. Monotonicity is a much weaker form of uniformity satisfied by nearly all known nonuniform protocols. While allowing different transitions as the population size grows, monotonicity requires that the transitions used for population size n must also be correct for all smaller population sizes $n' < n$ (i.e., an *overestimate* of the size cannot hurt), and furthermore that the transitions be no slower on populations of size n' than on populations of size n (though the transitions designed for size n may be slower on size n' than the transitions intended for size n'). Typically the nonuniform estimate of $\log n$ is used by the protocol to synchronize phases taking time $\approx T = \Theta(\log n)$, by having each agent individually count from T down to 0 (a so-called “leaderless phase clock”, our [Standard Counter Subroutine](#)). $\Theta(\log n)$ is the time required for an epidemic to reach the whole population and communicate some message to all agents before the phase ends. Most errors in such protocols are the result of some agent not receiving a message before a phase ends, i.e., the clock runs atypically faster than the epidemic. If the size estimate is significantly *larger* than $\log n$, this slows the protocol down, but only increases the probability that all agents receive intended epidemic messages before a phase ends; thus such protocols are monotone.

In our nonuniform protocol [Nonuniform Majority](#), which gives each agent the value $L = \lceil \log n \rceil$, giving a different value of L does not disrupt the stability (only the speed) of the protocol, with one exception: [Phase 3](#) must go for at least $\log n$ exponents, or else (i.e., if L is an underestimate) [Phase 4](#) could incorrectly report a tie. If we consider running [Phase 3](#) on a smaller population size n' than the size n for which it was designed, $L = \lceil \log n \rceil$ or $\lceil \log n' \rceil$ will be an *overestimate*, which does not disrupt correctness. Furthermore, since Clock agents are counting to L in each case, they are just as fast on population size n' as on size n . This implies that [Phase 3](#), thus the whole protocol [Nonuniform Majority](#), is monotone.

Output dominance references the concept of a *stable* configuration \mathbf{c} , in which all agents have a consensus opinion that cannot change in any configuration subsequently reachable from \mathbf{c} . A protocol is *output dominant* if in any stable configuration \mathbf{c} , adding more agents with states already present in \mathbf{c} maintains the property that every reachable stable configuration has the same output (though it may disrupt the stability of \mathbf{c} itself). This condition holds for all known stable majority protocols, including that described in this paper, because they obey the stronger condition that adding states already present in \mathbf{c} does not even disrupt its stability. Such protocols are based on the idea that two agents with the same opinion cannot create the opposite opinion, so stabilization happens exactly when all agents first converge on a consensus output.

To see that our protocol is output dominant, define a configuration to be *silent* if no transition is applicable (i.e., all pairs of agents have a null interaction); clearly a silent configuration is also stable. Although the definition of stable computation allows non-null transitions to continue happening in a stable configuration, many existing stable protocols have the stronger property that they reach a silent configuration with probability 1, including our protocol. It is straightforward to see that any silent configuration has the property required for output dominance, since if no pair of states in a configuration can interact nontrivially, their counts can be increased while maintaining this property. (One must rule out the special case of a state with count 1 that can interact with another copy of itself, which does not occur in our protocol’s stable configurations.)

Monotonicity can be seen as a natural condition that all “reasonable” non-uniform protocols must satisfy, but output dominance arose as an artifact that was required for the lower bound proof strategy of [5]. It remains open to prove an unconditional (i.e., removing the condition of output dominance) lower bound of $\Omega(\log n)$ states for *any* stable monotone majority protocol taking polylogarithmic time, or to show a stable polylogarithmic time monotone majority protocol using $o(\log n)$ states, which necessarily violates output dominance. If the unconditional lower bound holds, then our protocol is simultaneously optimal for both time and states. Otherwise, it may be possible to use $o(\log n)$ states to stably compute majority in polylogarithmic stabilization time with a non-output-dominant protocol. In this case, there may be an algorithm simultaneously optimal for both time and states, or there may be a tradeoff.

$O(\log n)$ time, $o(\log n)$ state non-stable protocol Berenbrink, Elsässer, Friedetzky, Kaaser, Kling, and Radzik [17] showed a non-stable majority protocol (i.e., it has a positive probability of error) using $O(\log \log n)$ states and converging in $O(\log^2 n)$ time. (See [Section 1](#) for more details.) Is there a protocol with $o(\log n)$ states solving majority in $O(\log n)$ time?

We close with questions unrelated to majority.

Fast population protocol for parity The majority problem “ $X_1 > X_2$?” is a special case of a *threshold* predicate, which asks whether a particular weighted sum of inputs $\sum_{i=1}^k w_i X_i > c$ exceeds a constant c . (For majority, $w_1 = 1, w_2 = -1, c = 0$; our protocol extends straightforwardly to other values of w_i , though not other values of c .) The threshold predicates, together with the *mod* predicates, characterize the *semilinear* predicates, which are precisely the predicates stably computable by $O(1)$ state protocols [10] with no time constraints (though $\Theta(n)$ time to stabilize is sufficient for all semilinear predicates [11] and necessary for “most” [13]). A representative mod predicate is *parity*: asking whether an odd or even number of agents exist. (More generally, asking the parity of the number of agents with initial opinion A .) Like majority, parity is solvable by a simple protocol in $O(n)$ time,²⁰ and it is known to require $\Omega(n)$ time for any $O(1)$ state

²⁰Agents start in state L_1 , undergo reactions $L_i, L_j \rightarrow L_{(i+j) \bmod 2}, F_{(i+j) \bmod 2}$; $L_i, F_j \rightarrow L_i, F_i$; $F_i, F_j \rightarrow F_i, F_i$, i.e., F agents adopt the parity of L agents, and two F agents with different parities adopt that of the sender. The first two

protocol to stabilize [13]. Techniques from [4] can be used to show that stabilization requires close to linear time even allowing up to $\frac{1}{2} \log \log n$ states. An interesting open question is to consider allowing $\omega(1)$ states in deciding parity. Can it then be decided in polylogarithmic time? Is there a parity protocol simultaneously optimal for both polylogarithmic time and states, such as the $O(\log n)$ time, $O(\log \log n)$ state protocol for leader election [18]? Or is there a tradeoff?

Fast population protocols for function computation The transition $X, Q \rightarrow Y, Y$, starting with sufficiently many excess agents in state Q , computes the function $f(x) = 2x$, because if we start with x agents in state X , eventually $2x$ agents are in state Y , taking time $O(\log n)$ to stabilize [26]. The similar transition $X, X \rightarrow Y, Q$ computes $f(x) = \lfloor x/2 \rfloor$, but it takes time $\Theta(n)$ to stabilize, as does any $O(1)$ -state protocol computing any linear function with a coefficient not in \mathbb{N} , as well as “most” non-linear functions such as $\min(x_1, x_2)$ (computable by $X_1, X_2 \rightarrow Y, Q$) and $\max(x_1, x_2)$ [13]. Can such functions be computed in sublinear time by using $\omega(1)$ states?

References

- [1] <https://github.com/UC-Davis-molecular-computing/ppsim/blob/main/examples/majority.ipynb>.
- [2] https://github.com/eftekhari-mhs/population-protocols/tree/master/Exact_Majority.
- [3] https://github.com/eftekhari-mhs/population-protocols/tree/master/Exact_Majority/animated_simulations.
- [4] Dan Alistarh, James Aspnes, David Eisenstat, Rati Gelashvili, and Ronald L Rivest. Time-space trade-offs in population protocols. In *SODA 2017: Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2560–2579. SIAM, 2017.
- [5] Dan Alistarh, James Aspnes, and Rati Gelashvili. Space-optimal majority in population protocols. In *SODA 2018: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2221–2239, 2018.
- [6] Dan Alistarh and Rati Gelashvili. Polylogarithmic-time leader election in population protocols. In *42nd International Colloquium on Automata, Languages, and Programming*, volume 9135, pages 479 – 491, 2015.
- [7] Dan Alistarh and Rati Gelashvili. Recent algorithmic advances in population protocols. *ACM SIGACT News*, 49(3):63–73, 2018.
- [8] Dan Alistarh, Rati Gelashvili, and Milan Vojnović. Fast and exact majority in population protocols. In *PODC 2015: Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing*, pages 47–56. ACM, 2015.
- [9] Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, 18(4):235–253, March 2006.
- [10] Dana Angluin, James Aspnes, and David Eisenstat. Stably computable predicates are semilinear. In *PODC 2006: 25th annual ACM Symposium on Principles of Distributed Computing*, pages 292–299, 2006.
- [11] Dana Angluin, James Aspnes, and David Eisenstat. Fast computation by population protocols with a leader. *Distributed Computing*, 21(3):183–199, September 2008.

reactions would take time $O(n \log n)$ for the last remaining L_i to update all F agents directly; the last reaction reduces this time to $O(n)$ [11, Theorem 5].

- [12] Dana Angluin, James Aspnes, and David Eisenstat. A simple population protocol for fast robust approximate majority. *Distributed Computing*, 21(2):87–102, 2008.
- [13] Amanda Belleville, David Doty, and David Soloveichik. Hardness of computing and approximating predicates and functions with leaderless population protocols. In *ICALP 2017: 44th International Colloquium on Automata, Languages, and Programming*, volume 80, pages 141:1–141:14, 2017.
- [14] Stav Ben-Nun, Tsvi Kopelowitz, Matan Kraus, and Ely Porat. An $O(\log^{3/2} n)$ parallel time population protocol for majority with $O(\log n)$ states. In *PODC 2020: Proceedings of the 39th Symposium on Principles of Distributed Computing*, pages 191–199, 2020.
- [15] Petra Berenbrink, Artur Czumaj, Angelika Steger, and Berthold Vöcking. Balanced allocations: The heavily loaded case. *SIAM Journal on Computing*, 35(6):1350–1385, 2006.
- [16] Petra Berenbrink, Robert Elsässer, Tom Friedetzky, Dominik Kaaser, Peter Kling, and Tomasz Radzik. A population protocol for exact majority with $O(\log^{5/3} n)$ stabilization time and $\Theta(\log n)$ states. In *DISC 2018: Proceedings of the 32nd International Symposium on Distributed Computing*, volume 10, pages 1–18, 2018.
- [17] Petra Berenbrink, Robert Elsässer, Tom Friedetzky, Dominik Kaaser, Peter Kling, and Tomasz Radzik. Time-space trade-offs in population protocols for the majority problem. *Distributed Computing*, Aug 2020.
- [18] Petra Berenbrink, George Giakkoupis, and Peter Kling. Optimal time and space leader election in population protocols. In *STOC 2020: Proceedings of the 52nd ACM SIGACT Symposium on Theory of Computing*, page 119–129, 2020.
- [19] Petra Berenbrink, Dominik Kaaser, Peter Kling, and Lena Otterbach. Simple and Efficient Leader Election. In *1st Symposium on Simplicity in Algorithms (SOSA 2018)*, volume 61, pages 9:1–9:11, 2018.
- [20] Petra Berenbrink, Dominik Kaaser, and Tomasz Radzik. On counting the population size. In *PODC 2019: Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, page 43–52, 2019.
- [21] Andreas Bilke, Colin Cooper, Robert Elsässer, and Tomasz Radzik. Brief announcement: Population protocols for leader election and exact majority with $O(\log^2 n)$ states and $O(\log^2 n)$ convergence time. In *PODC 2017: Proceedings of the ACM Symposium on Principles of Distributed Computing*, pages 451–453. ACM, 2017.
- [22] Michael Blondin, Javier Esparza, Stefan Jaax, and Antonín Kučera. Black ninjas in the dark: Formal analysis of population protocols. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '18*, page 1–10, 2018.
- [23] James M Bower and Hamid Bolouri. *Computational modeling of genetic and biochemical networks*. MIT press, 2004.
- [24] Janna Burman, Ho-Lin Chen, Hsueh-Ping Chen, David Doty, Thomas Nowak, Eric Severson, and Chuan Xu. Time-optimal self-stabilizing leader election in population protocols. In *PODC 2021: Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*, 2021.
- [25] Ioannis Chatzigiannakis, Othon Michail, Stavros Nikolaou, Andreas Pavlogiannis, and Paul G. Spirakis. Passively mobile communicating machines that use restricted space. *Theoretical Computer Science*, 412(46):6469–6483, October 2011.
- [26] Ho-Lin Chen, David Doty, and David Soloveichik. Deterministic function computation with chemical reaction networks. *Natural Computing*, 13(4):517–534, 2013.

- [27] Yuan-Jyue Chen, Neil Dalchau, Niranjan Srinivas, Andrew Phillips, Luca Cardelli, David Soloveichik, and Georg Seelig. Programmable chemical controllers made from DNA. *Nature Nanotechnology*, 8(10):755–762, 2013.
- [28] Anne Condon, Monir Hajiaghayi, David Kirkpatrick, and Ján Maňuch. Approximate majority analyses using tri-molecular chemical reaction networks. *Natural Computing*, 19(1):249–270, 2020.
- [29] David Doty and Mahsa Eftekhari. Efficient size estimation and impossibility of termination in uniform dense population protocols. In *PODC 2019: Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pages 34–42, 2019.
- [30] David Doty, Mahsa Eftekhari, Othon Michail, Paul G. Spirakis, and Michail Theofilatos. Brief announcement: Exact size counting in uniform population protocols in nearly logarithmic time. In *DISC 2018: 32nd International Symposium on Distributed Computing*, 2018.
- [31] David Doty and Monir Hajiaghayi. Leaderless deterministic chemical reaction networks. *Natural Computing*, 14(2):213–223, 2015.
- [32] David Doty and David Soloveichik. Stable leader election in population protocols requires linear time. *Distributed Computing*, 31(4):257–271, 2018.
- [33] Moez Draief and Milan Vojnović. Convergence speed of binary interval consensus. *SIAM Journal on control and Optimization*, 50(3):1087–1109, 2012.
- [34] Bartłomiej Dudek and Adrian Kosowski. Universal protocols for information dissemination using emergent signals. In *STOC 2018: Proceedings of the 50th ACM SIGACT Symposium on Theory of Computing*, pages 87–99, 2018.
- [35] Robert Elsässer and Tomasz Radzik. Recent results in population protocols for exact majority and leader election. *Bulletin of EATCS*, 3(126), 2018.
- [36] Leszek Gąsieniec, David Hamilton, Russell Martin, Paul G Spirakis, and Grzegorz Stachowiak. Deterministic population protocols for exact majority and plurality. In *20th International Conference on Principles of Distributed Systems*, 2016.
- [37] Leszek Gąsieniec and Grzegorz Stachowiak. Fast space optimal leader election in population protocols. In *SODA 2018: ACM-SIAM Symposium on Discrete Algorithms*, pages 2653–2667. SIAM, 2018.
- [38] Leszek Gąsieniec, Grzegorz Stachowiak, and Przemysław Uznański. Almost logarithmic-time space optimal leader election in population protocols. In *31st ACM Symposium on Parallelism in Algorithms and Architectures*, pages 93–102, 2019.
- [39] Svante Janson. Tail bounds for sums of geometric and exponential variables. *Statistics and Probability Letters*, 135:1–6, April 2018.
- [40] Adrian Kosowski and Przemysław Uznański. Population protocols are fast. *CoRR*, abs/1802.06872, 2018.
- [41] George B Mertzios, Sotiris E Nikolettseas, Christoforos L Raptopoulos, and Paul G Spirakis. Determining majority in networks with local interactions and very small local memory. In *International Colloquium on Automata, Languages, and Programming*, pages 871–882. Springer, 2014.
- [42] Y. Mocquard, B. Sericola, S. Robert, and E. Anceaume. Analysis of the propagation time of a rumour in large-scale distributed systems. In *2016 IEEE 15th International Symposium on Network Computing and Applications (NCA)*, pages 264–271. IEEE, 2016.

- [43] Yves Mocquard, Emmanuelle Anceaume, James Aspnes, Yann Busnel, and Bruno Sericola. Counting with population protocols. In *14th IEEE International Symposium on Network Computing and Applications*, pages 35–42, 2015.
- [44] Yves Mocquard, Emmanuelle Anceaume, and Bruno Sericola. Optimal proportion computation with population protocols. In *2016 IEEE 15th International Symposium on Network Computing and Applications (NCA)*, pages 216–223, 2016.
- [45] Yves Mocquard, Bruno Sericola, Frédérique Robin, and Emmanuelle Anceaume. Stochastic analysis of average based distributed algorithms. 2020.
- [46] David Soloveichik, Matthew Cook, Erik Winfree, and Jehoshua Bruck. Computation with finite stochastic chemical reaction networks. *Natural Computing*, 7(4):615–633, 2008.
- [47] Niranjana Srinivas, James Parkin, Georg Seelig, Erik Winfree, and David Soloveichik. Enzyme-free nucleic acid dynamical systems. *Science*, 358(6369):eaal2052, 2017.
- [48] Vito Volterra. Variazioni e fluttuazioni del numero d’individui in specie animali conviventi. *Memoria della Reale Accademia Nazionale dei Lincei*, 2:31–113, 1926.