Fast Convergence of k-Opinion Undecided State Dynamics in the Population Protocol Model

Talley Amir
talley.amir@yale.edu
Yale University
New Haven, USA

James Aspnes
james.aspnes@yale.edu
Yale University
New Haven, USA

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

Felix Biermeier
felix.biermeier@uni-hamburg.de
Universität Hamburg
Hamburg, Germany

Christopher Hahn
christopher.hahn-1@uni-hamburg.de
Universität Hamburg
Hamburg, Germany

Dominik Kaaser
dominik.kaaser@tuhh.de
TU Hamburg
Hamburg, Germany

John Lazarsfeld
john.lazarsfeld@yale.edu
Yale University
New Haven, USA

ABSTRACT

We analyze the convergence of the k-opinion Undecided State Dynamics (USD) in the *population protocol* model. For k=2 opinions it is well known that the USD reaches *consensus* with high probability within $O(n \log n)$ interactions. Proving that the process also quickly solves the consensus problem for $k > 2$ opinions has remained open, despite analogous results for larger k in the related parallel *gossip* model. In this paper we prove such convergence: under mild assumptions on k and on the initial number of undecided agents we prove that the USD achieves plurality consensus within $O(kn \log n)$ interactions with high probability, regardless of the initial bias. Moreover, if there is an initial *additive* bias of at least $\Omega(\sqrt{n} \log n)$ we prove that the initial plurality opinion wins with high probability, and if there is a multiplicative bias the convergence time is further improved. Note that this is the first result for $k > 2$ for the USD in the population protocol model. Furthermore, it is the first result for the unsynchronized variant of the USD with $k > 2$ which does not need any initial bias.

CCS CONCEPTS

• **Theory of computation → Distributed algorithms**.

KEYWORDS

population protocols, plurality consensus, randomized algorithms

ACM Reference Format:
Talley Amir, James Aspnes, Petra Berenbrink, Felix Biermeier, Christopher Hahn, Dominik Kaaser, and John Lazarsfeld. 2023. Fast Convergence of k-Opinion Undecided State Dynamics in the Population Protocol Model. In *ACM Symposium on Principles of Distributed Computing (PODC '23), June 19–23, 2023, Orlando, FL, USA.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3583668.3594589

PODC '23, June 19–23, 2023, Orlando, FL, USA

ACM ISBN 979-8-4007-0121-4/23/06…$15.00
https://doi.org/10.1145/3583668.3594589

1 INTRODUCTION

The Undecided State Dynamics (USD) is a simple protocol designed for distributed models of computation where n indistinguishable agents engage in pairwise interactions. The protocol assumes that every agent initially supports one of $k \geq 2$ *opinions*, and the process evolves according to the following transition rules: when an agent x interacts with an agent y and the opinions of x and y differ, x transitions to an *undecided* state. When x interacts with y and x is undecided, x adopts the opinion of y. If y is undecided, or if its opinion is the same as that of x, no updates occur.

Given its suitability as a primitive for other distributed tasks, a substantial amount of recent work has analyzed this process as a protocol for *consensus* under varying settings of two key problem parameters: first, the exact distributed model of pairwise interaction, and second, the number of opinions k. The USD was originally introduced by Angluin et al. [5] for k=2 opinions in the *population protocol* model,[1] where at every discrete time step a single pair of agents is chosen uniformly at random to interact. In this setting, Angluin et al. showed that the USD reaches consensus (a configuration where all agents support the same opinion) in $O(n \log n)$ interactions.[2] Moreover, those authors (and later Condon et al. [20] via a simplified analysis) also showed that the process solves the *approximate majority* problem, meaning the eventual consensus opinion is the one whose initial support was larger, so long as the initial *bias* (the difference between the support of the two opinions) is sufficiently large (specifically, of order $\Omega(\sqrt{n \log n})$).

[1]Independently, Perron et al. [40] analyzed the two opinion USD in the *asynchronous gossip* model of Boyd et al. [18], which can be viewed as the continuous time variant of the population protocol model. For simplicity, our work focused on the latter model, although our results extend easily to the former.

[2]Throughout, all stated results hold with high probability (w.h.p.), meaning with probability $1 - n^{-c}$ for some $c > 0$.

Separately, the USD has also been analyzed in the parallel *gossip* model of communication, where in each synchronous round, *every* agent selects an interaction partner uniformly at random. In this model, Clementi et al. [19] showed convergence results for the case of k=2 opinions that are analogous to those in the population protocol model: the process reaches consensus in $O(\log n)$ synchronous rounds and additionally solves approximate majority when the initial bias is at least $\Omega(\sqrt{n \log n})$. In this model, Becchetti et al. [10] also analyzed the process in the higher-dimensional regime when $k > 2$. Assuming a large enough *multiplicative* bias in the initial supports of opinions, they show that the USD reaches *plurality consensus* in $O(k \log n)$ parallel rounds, meaning that the eventual consensus opinion is the one whose initial support was largest.[3]

Although the population protocol model can be viewed as the asynchronous analog to the synchronous gossip model, the differences in these interaction scheduling modes cause the USD to exhibit significant qualitative differences when run in either setting, even in the case when k=2. This can largely be attributed to the observation that one round of parallel interactions in the gossip model can lead to a constant fraction of agents changing their opinion, whereas at most a single change of opinion can result from each interaction in the population protocol model. These differences, as Clementi et al. [19] remark, have largely prevented any general analysis techniques from transferring between the two models. In particular, it has remained an open problem to analyze the convergence rate of the USD in the population protocol model when $k > 2$.

1.1 Our Contribution

In this work, we close the aforementioned gap and analyze the USD in the population protocol model in the high dimensional $k > 2$ regime.[4] In particular, under mild assumptions, we prove that the USD solves the plurality consensus problem in this model in $O(k \cdot n \log n)$ interactions. Stated informally, we prove the following result:

Theorem 1.1 (informal). *Consider the USD in the population protocol model, and assume a sufficiently small number of initially undecided agents. Then for any* $2 \le k \le O(\sqrt{n}/\log^2 n)$

1. *If the initial support of the plurality opinion is at least* $\Omega(\sqrt{n} \log n)$ *larger (additively) than all other opinions, then the process reaches plurality consensus within* $O(kn \log n)$ *interactions,*
2. *If the initial support of the plurality opinion is a constant multiplicative factor larger than all other opinions, then the process reaches plurality consensus within* $O(kn + n \log n)$ *interactions,*
3. *The process reaches an arbitrary consensus configuration otherwise,*

where each statement holds with high probability.

The exact statement of our main result is given in Theorem 2.1, where the convergence rates have a more precise dependence on the magnitude of the opinion with largest initial support. Roughly speaking, the convergence rate of our result is analogous to that of Becchetti et al. [10] for the gossip model: in that model, plurality consensus is reached within $O(k \cdot \log n)$ *rounds*, while in the population protocol model, we show it takes $O(k \cdot n \log n)$ *interactions*. However, unlike the result of Becchetti et al., our analysis only requires an *additive* bias of $\Omega(\sqrt{n} \log n)$ to reach plurality consensus (rather than a constant *multiplicative* bias); it holds for larger $k = O(\sqrt{n}/\log^2 n)$ (compared to $k = O((n/\log n)^{1/3})$); and we show the process still reaches consensus when starting from a configuration with no initial bias (e.g., when the initial support of each opinion is n/k). On the other hand, when the initial configuration does contain a constant multiplicative bias, our analysis gives a faster convergence rate than in the additive bias regime. Moreover, our convergence rate under a multiplicative bias is faster (when considering its corresponding *parallel time*) than the rate given by Becchetti et al. when the support of the initially largest opinion is close to the average opinion support.[5] In this setting, our results for the population protocol model can be viewed as improvements to the analogous results of Becchetti et al. for the gossip model. If there is a large multiplicative bias (larger than $\log n$) the results by Becchetti et al. give better bounds on the convergence time.

Similar to previous analyses in both models [5, 10, 11, 20] our analysis requires carefully defining a sequence of *phases* throughout which the (qualitative and quantitative) behavior of the process varies. The main challenge is to define appropriate potential functions that allow us to track the progress of the process. In Section 2 we give an overview about the main ideas of our analysis.

1.2 Related Works

The Undecided State Dynamics. The two-opinion USD was introduced independently by Angluin et al. [5] for the population protocol model and by Perron et al. [40] for the closely related (continuous time) asynchronous gossip model. Both works show that the process converges w.h.p. in $O(n \log n)$ steps (respectively, $O(\log n)$ continuous time). Condon et al. [20] give an improved analysis for the two-opinion case in the population model and show the process solves the approximate majority problem assuming an initial additive bias of $\Omega(\sqrt{n \log n})$, which improves over the additive bias of $\omega(\sqrt{n} \log n)$ needed in the analysis of Angluin et al. Similar to our approach, both Angluin et al. and Condon et al. analyze the process in distinct phases that depend on the number of undecided agents and the magnitude of bias in the configuration. In particular, after introducing a suitable structure of phases and sub-phases, the analysis of Condon et al. reduces the convergence of the process to analyzing a sequence of biased, one-dimensional random walks. The boundaries imposed by the phase structure are used to control the magnitudes of the bias, and bounds on the number of interactions needed to complete each phase are derived using standard concentration techniques.

In the parallel gossip model, the convergence of the USD for the $k \ge 2$ opinion case was first studied by Becchetti et al. [10]. Central to their analysis is the introduction of the *monochromatic distance*, which measures the uniformity (i.e., lack of bias) of a configuration. Roughly speaking, this distance is the sum of squares of

[3]Note that when $k > 2$, the initial support of largest may not be a majority, which is why the term *plurality* is used.

[4]Our analysis can also be applied when k=2 and recovers the existing convergence results [5, 20] in this setting.

[5]This is shown explicitly in Section 8.

the support of each opinion, normalized by the square of the most popular opinion. They show convergence within $O(\text{md}(\mathbf{x}) \cdot \log n)$ parallel rounds, where $\text{md}(\mathbf{x})$ is the monochromatic distance of the initial configuration, which is always bounded above by k. This analysis only holds when the initial configuration has a multiplicative bias. In the two-color case, Clementi et al. [19] later present a tight analysis (giving convergence rates that hold for any initial configuration) without using the monochromatic distance, but an analysis for $k > 2$ opinions, starting from any initial configuration in the gossip model still remains open.

In a related strain of research, multiple works [6, 8, 16, 33] have analyzed a *synchronized* variant of the USD where the system alternates between two different phases in a synchronized fashion. In the first phase, all agents perform one step of the USD. In the second phase, all undecided agents adopt an opinion again. The use of so-called *phase clocks* that synchronize the agents allows for a polylogarithmic convergence time regardless of the initial opinion configuration. This improved convergence time comes at the price of making the protocol "less natural": these protocols have a significant state overhead and are typically not *uniform*, meaning that the transition function or state space depend on n.

Other Consensus Dynamics. In the population protocol model, consensus for the case of k=2 opinions is commonly known as the *majority problem.* A large number of works [1, 3, 12, 14, 15, 29] aim to identify the majority opinion even if the initial winning margin is as small as only 1. The best known result [29] solves this *exact* majority problem in $O(n \log n)$ interactions using $O(\log n)$ states, both in expectation. For more details on algorithmic advances in Population Protocols we refer the reader to the surveys by Elsässer and Radzik [30] and Alistarh and Gelashvili [2].

Less is known about exact plurality consensus protocols for $k > 2$ opinions. One line of research focuses on the state space requirements to *always* compute the exact plurality opinion. In [39] the authors show that always correct plurality consensus requires $\Omega(k^2)$ states. The currently best known protocol requires $O(k^6)$ many states [31]. In [7] the authors relax the requirement to *always* return the correct result. They present a protocol for k>2 opinions that may fail with small probability. This negligible error probability allows them to break the lower bound and design a protocol that converges w.h.p. in $O(n\cdot(k \log n + \log^2 n))$ interactions using $O(k+\log n)$ states.

A related family of protocols are the j-Majority processes. The idea is that every agent adopts the majority opinion among a random sample of j other agents (breaking ties randomly). The most simple variant (for j=1) is also known as the so-called Voter process [17, 21, 34, 35, 38]. Here, every agent adopts the opinion of a single, randomly chosen agent. The protocols for j=2 and j=3 have been analyzed under the names of TwoChoices process [22–24] and the 3-Majority dynamics [11, 13, 32]. In the TwoChoices process, lazy tie-breaking towards an agent's original opinion is assumed. Ghaffari and Lengler [32] show for the TwoChoices process with $k = O(\sqrt{n/\log n})$ and for 3-Majority with $k = O(n^{1/3}/\log n)$ that consensus is reached in $O(k \cdot \log n)$ rounds w.h.p. For arbitrary k they show that 3-Majority reaches consensus in $O(n^{2/3} \log^{3/2} n)$ rounds w.h.p. Schoenebeck and Yu [41] analyze the convergence time of a generalization of multi-sample consensus protocols for two opinions on complete graphs and Erdős-Rényi graphs. In the MedianRule process [28] the authors assume that opinions are ordered. In every step every agent then adopts the median of its own opinion and two randomly sampled opinions. This protocol reaches consensus in $O(\log k \log\log n + \log n)$ rounds w.h.p. We remark that in contrast to the MedianRule the USD does not require a total order among the opinions. For further references and additional protocols in similar models we refer the reader to the survey of consensus dynamics by Becchetti et al. [9].

2 BACKGROUND AND OVERVIEW OF RESULTS

In this section, we first introduce some of the preliminaries and notation related to the population protocols model and the USD. We then provide a technical overview of our main result.

Population Protocols. We consider a population protocol for n identical, anonymous agents, where each agent is modeled as a finite state machine with state space Q. Agents interact in pairs drawn uniformly at random. In an interaction (u, v) agent u is called responder and agent v is called initiator. We allow for agents to interact with themselves. The population protocol is defined by its transition function $\delta : Q^2 \to Q^2$.

The undecided state dynamics (USD) is defined as follows. Each agent has either one of k opinions or it is undecided, i.e., $Q = \{1, \dots k, \bot\}$ where $\bot$ stands for undecided. The *undecided state* population protocol is given by the transition function

$$(q, q') \to \begin{cases} (\bot, q') \text{ if } q, q' \neq \bot \wedge q \neq q' \\ (q', q') \text{ if } q = \bot, q' \neq \bot \\ (q, q') \text{ otherwise.} \end{cases}$$

Observe that only the initiator q changes its state.

A configuration $\mathbf{x}(t)$ at time t is a vector $(x_1(t), \dots, x_k(t), u(t))$ of length $k + 1$. For $1 \leq i \leq k$, $x_i(t)$ is the number of agents of Opinion i and $u(t) = n - \sum_{i=1}^{k} x_i(t)$ is the number of undecided agents. In the beginning we assume $x_1(0) \geq x_2(0) \geq \cdots \geq x_k(0)$. For $t > 0$ we define $\max(t)$ as the index of the opinion with the largest support at step t (if there are several opinions with the same maximum support we pick an arbitrary one). Furthermore we introduce the notation $x_{\max}(t) = x_{\max(t)}(t) = \max_{i \in [k]} \{ x_i(t) \}$ for the support of the largest opinion at time t. Note that $x_{\max}(t)$ can refer to the support of different opinions over time.

We call an Opinion i *significant* if $x_i(t) > x_{\max}(t) - \alpha \cdot \sqrt{n} \log n$ for some fixed constant α. An opinion that is not significant is called *insignificant.* A configuration $\mathbf{x}$ has an *additive bias* β if there exists an Opinion m such that for all other opinions $i \neq m$ we have $x_m \geq x_i + \beta$. We say that a configuration $\mathbf{x}$ has a *multiplicative bias* α if there exists an Opinion m such that for all other opinions $i \neq m$ we have $x_m \geq \alpha \cdot x_i$.

In the following we use upper case letters for random variables (for example $\mathbf{X}(t)$ and $U(t)$) and lower case letters ($\mathbf{x}(t)$ and $u(t)$) for fixed configurations or values.

We now state our main theorem. We remark that in our analysis we bound the convergence time in terms of $n/x_1(0)$, where $x_{\max}(0) = x_1(0)$ is the support of the initially largest opinion. Under the assumptions of our theorem, however, we have $x_1(0) > n/(2k)$, which leads to the bounds in terms of k.

Theorem 2.1. *Let $c > 0$ be an arbitrary constant and let $\mathbf{x}(0)$ be an initial configuration with $k \le c \cdot \sqrt{n}/\log^2(n)$ opinions with $u(0) \le (n - x_1(0))/2$ and $x_1(0) \ge x_i(0)$ for all $i \in [k]$. Then w.h.p. all agents agree on Opinion 1 within*

(1) $\mathrm{O}\big(n \log n + n^2/x_1(0)\big) = \mathrm{O}(n \log n + n \cdot k)$ *interactions if $\mathbf{x}(0)$ has a multiplicative bias of at least $1 + \varepsilon$ for an arbitrary constant ε.*
(2) $\mathrm{O}(n^2 \log n/x_1(0)) = \mathrm{O}(k \cdot n \log n)$ *interactions if $\mathbf{x}(0)$ has an additive bias of at least $\Omega\big(\sqrt{n} \log n\big)$.*

Without any bias all agents agree on a significant opinion within $\mathrm{O}\big(n^2 \log n/x_1(0)\big) = \mathrm{O}(k \cdot n \log n)$ interactions w.h.p.

Main Idea of the Analysis. The straightforward approach in the analysis of consensus processes is to track the growth of the support of the plurality Opinion 1 via change of the ratio $x_1(t)/x_i(t)$ over time t. Unfortunately, the change of the support of a single opinion depends on the entire configuration, that is, the support of all other opinions and also the number of undecided agents. Let us fix two opinions i and j with $x_i > x_j$. Then it is possible for the support of Opinion j to grow faster than the support of Opinion i and vice versa, depending on the number of undecided nodes. Hence, to track the progress of the plurality opinion one has to take a close look at the number of undecided nodes. This, in turn, is heavily influenced by the support of all opinions. To cope with this "nonlinearity" we use the potential function $Z_\alpha(t) = n - 2u(t) - \alpha \cdot x_{\max}(t)$, where we use different values of α for different phases. We analyze the drift of $Z_\alpha(t)$ which allows us to show that the number of undecided agents quickly approaches an "unstable equilibrium" u^*. Whenever the process is close to the equilibrium (which changes over time), we can perform a "classical" analysis and show, e.g., that bias between two agents doubles in a certain number of interactions.

Our analysis also handles the case when there is no bias at all. For this we proceed in two steps. First we show that the support difference between two arbitrary but fixed large opinions quickly reaches a value of $\sqrt{n}$ via an anti-concentration bound. From there we bound the probability that the opinions continue to drift apart. In our analysis we rely heavily on existing concentration bounds for the hitting times of one-dimensional random walks with drift, which we can use after establishing the appropriate reductions and potential functions in each phase of the process. The analysis is divided into five parts that correspond to different *phases* of the process. The phases are listed in the following table:

Phase	Section	End Condition	Running Time	Main Lemma
1	Section 3	$u \ge (n - x_{\max})/2$	$\mathrm{O}(n \log n)$	Lemma 3.1
2	Section 4	$\forall i : x_{\max} \ge x_i + \Omega(\sqrt{n} \log n)$	$\mathrm{O}(n^2 \log n/x_{\max})$	Lemma 4.5
3	Section 5	$\forall i : x_{\max} \ge 2x_i$	$\mathrm{O}(n^2 \log n/x_{\max})$	Lemma 5.3
4	Section 6	$x_{\max} \ge 2n/3$	$\mathrm{O}(n^2/x_{\max} + n \log n)$	Lemma 6.4
5	Section 7	$x_{\max} = n$	$\mathrm{O}(n \log n)$	Lemma 7.1

Note that the process does not have to pass through all five phases. For example, the second phase is not needed if there is a large bias in the initial configuration. Our analysis shows that the identity of the majority opinion does not change after the end of the second phase (or not at all if a large enough additive bias is present from the beginning).

3 RISE OF THE UNDECIDED (PHASE 1)

In this section we analyze the running time of Phase 1 which ends as soon as we have a sufficient number of undecided agents (Lemma 3.1). Additionally we show that $x_1(0)$ decreases by at most a constant fraction w.h.p. (Lemma 3.2). Furthermore, an additive and multiplicative bias is preserved as long as $\mathbf{x}(0)$ is an initial configuration with bias. At the end of this section we show an upper bound on the number of undecided agents which holds during the whole running time of the process (Lemma 3.3). This lemma will be used to estimate the running time of the remaining phases.

In the analysis of Lemma 3.1 we use the potential function

$$Z(t) = n - 2u(t) - x_{\max}(t) \ .$$

Observe that Phase 1 ends as soon as $Z(t) \le 0$, since in this case $u(t') \ge n/2 - x_{\max}(t')/2$.

Lemma 3.1. *Let $T_1 = \inf\{t \ge 0 \mid u(t) \ge n/2 - x_{\max}(t)/2\}$. Then $\Pr[T_1 \le \lceil 7n \ln n \rceil] \ge 1 - n^{-3}$.*

Proof. To show the lemma we calculate the expected change in $Z(t)$ for $Z(t) \ge 0$ and apply a drift theorem from [36]. There are three cases. First we consider the case $U(t+1) = u(t) - 1$. In this case a decided agent interacts with an undecided agent, and the latter adopts the opinion of the decided agent. Let $M(t) = \{\, i \in [k] \mid x_i(t) = x_{\max}(t) \,\}$ be the set of all opinions with maximum support at time t. For each Opinion i, an undecided initiator interacts with a responder of Opinion i with probability $x_i(t) \cdot u/n^2$. If $i \in M(t)$, then $Z(t)$ increases by 1. Otherwise $Z(t)$ increases by 2.

Next we consider the case $U(t+1) = u(t) + 1$. In this case a decided initiator interacts with a responder of a different opinion and becomes undecided. For each Opinion i, this happens with probability $x_i(t) \cdot (n - u(t) - x_i(t))/n^2$. If $i \in M(t)$, then $Z(t)$ decreases by 1 or 2. Otherwise $Z(t)$ decreases by 2.

With the remaining probability a step is unproductive and $Z(t)$ does not change. Using these cases, we bound the expected drift of $Z(t)$ as

$$\begin{aligned}
&\mathbb{E}[Z(t) - Z(t+1) \mid \mathbf{X}(t) = \mathbf{x}] \\
&\quad \ge - \sum_{i \in M(t)} \frac{x_i \cdot u}{n^2} - 2 \sum_{i \notin M(t)} \frac{x_i \cdot u}{n^2} + \sum_{i \in M(t)} \frac{x_i(n - u - x_i)}{n^2} \\
&\qquad + 2 \sum_{i \notin M(t)} \frac{x_i(n - u - x_i)}{n^2} \\
&\quad \ge \sum_{i \in [k]} \frac{x_i(n - 2u - x_{\max})}{n^2} + \sum_{i \notin M(t)} \frac{x_i(n - 2u - x_{\max})}{n^2} \\
&\quad \ge \frac{(n-u)(n - 2u - x_{\max})}{n^2} \ge \frac{Z(t)}{2n},
\end{aligned}$$

where we used that $x_i \le x_{\max}$, $Z(t) = n - 2u - x_{\max} \ge 0$, and $u < n/2$ by definition of Phase 1. We now apply the multiplicative drift result of [36] with $r = 3 \ln n$, $s_0 = n - 2u(0) - x_{\max}(0) \le n$, $s_{min} = 1$, $\delta = 1/(2n)$ and get

$$\begin{aligned}
\Pr[T_1 > \lceil 7n \ln n \rceil] &\le \Pr\left[T_1 > \left\lceil \frac{6 \ln n + \ln(n - 2u(0) - x_{\max}(0))}{1/(2n)} \right\rceil\right] \\
&\le e^{-3 \cdot \ln(n)} = n^{-3} \ .
\end{aligned}$$

□

Given the bound on T_1, we proceed to show that both the support of the most popular opinion and the bias of the initial configuration do not decrease too much until time T_1. Recall that initially Opinion 1 has the largest support.

Lemma 3.2. *Let $\alpha, \varepsilon > 0$ be arbitrary constants. Then each of the following statements holds with probability at least $1 - 4n^{-3}$:*

1. *If $x_1(0) - x_i(0) \geq \alpha \cdot \sqrt{n} \log n$, then $X_1(T_1) - X_i(T_1) \geq \alpha/3 \cdot \sqrt{n} \log n$.*
2. *If $x_1(0) \geq (1+\varepsilon) \cdot x_i(0)$, then $X_1(T_1) \geq (1+\varepsilon/(6+5\varepsilon)) \cdot X_i(T_1)$.*
3. *For the largest opinion we have $X_1(T_1) \geq x_1(0)/3$.*

Proof Sketch. For the first statement we show that $\mathbb{E}[(X_1(t) - X_i(t))/(n - U(t))] \geq 0$ and apply a Hoeffding bound. For the second statement we show that

$$\Pr[X_1(t+1) = x_1 + 1 \mid \mathbf{X}(t) = \mathbf{x}] \leq \Pr[X_1(t+1) = x_1 - 1 \mid \mathbf{X}(t) = \mathbf{x}]$$

such that we can bound the development of X_1 by a fair random walk. This enables us to relate the multiplicative bias to the additive bias. The third statement is derived from the first statement by choosing an Opinion i with $x_i(0) = 0$. The full proof can be found in the full version [4]. □

Next we prove the upper bound on the number of undecided agents. The lemma shows that the number of undecided agents stays close to a threshold value $u^* = n \cdot (k-1)/(2k-1) \approx n/2$. Intuitively, this threshold u^* can be regarded as an (unstable) equilibrium for the number of undecided agents: in configurations with more than u^* undecided agents it is more likely that an undecided agent becomes decided than vice versa, whereas in configurations with less than u^* undecided agents it is more likely that a decided agent becomes undecided than vice versa.

Lemma 3.3. *Assume $u(0) \leq (n - x_{\max}(0))/2$. Then*

$$\Pr\left[\forall t \in [n^3] : u(t) \leq \frac{n}{2} - \frac{1}{5c} \cdot \sqrt{n} \log(n)\right] > 1 - n^{-3}.$$

Proof Sketch. We first prove the claim for $u(t) > u^* + 3 \cdot \sqrt{n \log n}$. At the end of the full proof we show how the lemma statement follows out of this. We model the number of undecided agents over time t as a non-lazy random walk $Z(t)$ with state space $\{0, \ldots, n-1\}$. Then we couple $Z(r)$ with a random walk $W(r)$ on the integers with a reflecting barrier at 0 and otherwise fixed transition probabilities. For $W(r)$ we can derive a bound on the probability $\Pr\left[\exists t \in [n^3] : W(t) \geq 3 \cdot \sqrt{n \log n}\right]$. The bound follows since in this case $Z(r) \leq W(r) + Z(0)$. To conclude the proof we show that the lemma statement follows out of our bound stated in terms of u^*. The full proof can be found in the full version [4]. □

4 GENERATION OF AN ADDITIVE BIAS (PHASE 2)

Recall that T_1 is defined as the end of Phase 1. In this section we consider configurations at time T_1 without any additive bias. These configurations will have several significant opinions. We define T_2 as the first time $t \geq T_1$ where $\mathbf{x}(t)$ has only one opinion left which is significant.

Note that $x_{\max}(t) \geq x_{\max}(0)/2 = \Omega(\sqrt{n} \cdot \log^2(n))$ for each interaction t in this phase. This follows from Lemma 3.3 together with the pigeonhole principle. In Lemma 4.5 we show that w.h.p. the running time of this phase is $\mathrm{O}(n^2 \cdot \log n / x_{\max}(T_1))$. To show that result we first need a lower bound (as opposed to the upper bound of Lemma 3.3) on the number of undecided agents. Again, this bound holds until the end of the process.

Lemma 4.1.

$$\Pr\left[\forall t \in [T_1, n^3] : u(t) \geq n/2 - x_{\max}(t)/2 - 8\sqrt{n \cdot \ln n}\right] \geq 1 - n^{-5}.$$

Proof Sketch. Recall that for the proof of Lemma 3.1 we defined $Z(t) = n - 2u(t) - x_{\max}(t)$. We then showed that we have a drift towards zero. We use this for a drift analysis following Theorem 6 in [37]. The full proof can be found in the full version [4]. □

In the following lemma we show that the support of the largest opinion does not shrink by more than a factor of two during Phase 2.

Lemma 4.2. *Let $c > 0$ be an arbitrary constant and define $T = c \cdot n^2 \cdot \log n / x_{\max}(T_1)$. Then*

$$\Pr[\forall t \in [T_1, T_1 + T] : x_{\max}(t) \geq x_{\max}(T_1)/2] \geq 1 - n^{-5}.$$

In Lemma 4.3 we first show that "small opinions" remain small (they only double their support). With small opinion we mean opinions having a support which have support at most $20\sqrt{n \log n}$ and are thus at least a polylogarithmic factor smaller compared to $x_{\max}(t)$. Then in the second part we show that insignificant opinions remain insignificant. Recall that an Opinion i is insignificant if $x_{\max}(t) - x_i(t) = \Omega(\sqrt{n} \log n)$.

Lemma 4.3. *Let $c, c' > 0$ be arbitrary constants and define $T = c \cdot n^2 \cdot \log n / x_{\max}(T_1)$. Assume for Opinion j there exists a time $t_0 \in [T_1, T_1 + T]$ with*

1. *$x_j(t_0) \leq 20\sqrt{n \log n}$. Then*
$$\Pr\left[\forall t \in [t_0, T_1 + T] : x_j(t) \leq 40\sqrt{n \log n}\right] \geq 1 - 2n^{-3}.$$
2. *$x_{\max}(t_0) - x_j(t_0) \geq c' \cdot \sqrt{n} \log n$. Then*
$$\Pr\left[\forall t \in [t_0, T_1 + T] : x_{\max}(t) - x_i(t) \geq c'/2 \cdot \sqrt{n} \log n\right] \leq 1 - 2n^{-3}.$$

Proof Sketch. In the first part we bound the probability for a small Opinion j to grow using Lemma 4.1. This probability is sufficiently small for Opinion j not to double. In the second part, we make a case distinction based on the size of $x_j(t_0)$. If $x_j(t_0)$ is small, then the support of x_j does not double (see Part 1) while x_1 keeps at least half of its support (Lemma 4.2). Otherwise, we use Lemma 4.1 to show that the bias is likely to increase. Then the second part follows from the gambler's ruin problem. The full proof can be found in the full version [4]. □

The following lemma constitutes the foundation of the application of the drift result from [28] which will be used in the proof of Lemma 4.5. In the first part of Lemma 4.4 we consider two important opinions with (almost) the same support. We use an anti-concentration result to show that their support difference quickly reaches $\Omega(\sqrt{n})$. In the second part we again consider two important opinions and give precise bounds on the probability that their support difference increases by a constant factor. Our proof is based on gambler's ruin problem. The proof of this result can be found in the full version [4].

Lemma 4.4. *Fix two opinions i and j and assume there exists $t_0 \geq T_1$ with $x_i(t_0) \geq x_j(t_0) \geq x_{\max}(t_0) - 4\alpha\sqrt{n}\log n$. Let $T = 40 \cdot n^2/x_{\max}(T_1)$ and $\Delta_{ij}(t_0) = x_i(t_0) - x_j(t_0)$. Then*

(1) *If $\Delta_{ij}(t_0) < 4\alpha \cdot \sqrt{n}$ then*

$$\Pr\left[\Delta_{ij}(t_0+T) \geq 4\alpha \cdot \sqrt{n}\right] \geq e^{-\frac{\alpha^2}{16}}.$$

(2) *If $\Delta_{ij}(t_0) \geq 4\alpha \cdot \sqrt{n}$ then*

$$\Pr\left[\Delta_{ij}(t_0+T) \geq \min\{(3/2)\cdot\Delta_{ij}(t_0), 4\alpha\sqrt{n}\log n\}\right] \geq 1-e^{-\frac{\Delta_{ij}(t_0)}{\sqrt{n}}}.$$

Now we are ready to analyze the running time of Phase 2.

Lemma 4.5. *Let*

$$T_2 = \inf\{\, t \geq T_1 \mid \exists i \in [k] : \forall j \neq i : x_i(t) - x_j(t) \geq \alpha\sqrt{n}\log n \,\}.$$

Then

$$\Pr\left[T_2 - T_1 \leq 40 \cdot c \cdot n^2 \cdot \log n / x_{\max}(T_1)\right] \geq 1 - 2n^{-2}.$$

Proof. We define

$$\hat{T} = \inf\left\{\, t \geq T_1 \,\middle|\, u(t) \notin \left[\frac{n - x_{\max}(t)}{2} - 8\sqrt{n\ln n}, \frac{n}{2}\right] \text{ or } \frac{x_{\max}(T_1)}{x_{\max}(t)} < \frac{1}{3} \right\}$$

as a stopping time and $(\hat{\mathbf{X}})_{t\geq T_1}$ as the process with $\hat{\mathbf{X}}(t) = \mathbf{X}(t)$ for all $t \leq \hat{T}$ and $\hat{\mathbf{X}}(t) = \mathbf{X}(\hat{T})$ for $t > \hat{T}$. From Lemma 4.1 it follows that $u(t) \geq (n - x_{\max}(t))/2 - 8 \cdot \sqrt{n \ln n}$ for all $t \in [T_1, n^3]$, w.h.p. From Lemma 3.3 it follows that $u(t) \leq n/2$ for all $t \in [T_1, n^3]$, w.h.p. Finally, Lemma 4.2 gives us that $x_{\max}(t) \geq x_{\max}(T_1)/3$ for all $t \in [T_1, T_1 + cn^2 \log n / x_{\max}(T_1)]$, w.h.p. Thus, $\hat{T} - T_1 = \Omega(n^2 \cdot \log n / x_{\max}(T_1))$ w.h.p. and we can assume that $(\mathbf{X})_{t\geq T_1}$ and $(\hat{\mathbf{X}})_{t\geq T_1}$ are identical for $t \in [T_1, T_1 + \mathrm{O}(n^2 \cdot \log n / x_{\max}(T_1))]$.

Recall that an Opinion i is *significant* at time t if $x_i(t) > x_{\max}(t) - \alpha\sqrt{n}\log n$. In the following we call an Opinion i *important* at time t if $x_i(t) > x_{\max}(t) - 4 \cdot \alpha\sqrt{n}\log n$. In the following we will show that for each pair of important opinions i and j at time T_1 at least one of them becomes unimportant. Furthermore, we show that no unimportant opinion ever becomes significant. From this follows that after $\mathrm{O}(n^2/\hat{x}_{max}(T_1) \cdot \log n)$ only one significant opinion remains.

First we consider a pair of opinions i and j which are important at time T_1 and show that w.h.p. at least one of them becomes unimportant within the next $\tau = 40 \cdot cn^2 \cdot \log n / \hat{x}_{max}(T_1)$ interactions.

We divide the interactions from $[T_1, T_1+\tau]$ into $c_1 \log n$ subphases of length $40 \cdot n^2/\hat{x}(T_1)$ each. For $1 \leq i \leq c \log n$ we define $\ell_1 = 1$ and $\ell_i = 1 + (i-1) \cdot n^2/\hat{x}(T_1)$. Then the ith subphase contains interactions ℓ_i to $(\ell_{i+1} - 1)$. Furthermore, we define t_i as the first interaction in subphase i.

Now we fix an arbitrary subphase i and we consider two cases. If $\hat{x}_i(t_i) - \hat{x}_j(t_i) < 4\alpha\sqrt{n}$ then it follows from Lemma 4.4

$$\Pr\left[\hat{X}_i(t_{i+1}) - \hat{X}_j(t_{i+1}) \geq 4\alpha\sqrt{n}\right] \geq e^{-\frac{\alpha^2}{16}} \tag{1}$$

Otherwise, if $\hat{x}_i(t_i) - \hat{x}_j(t_i) \geq 4\alpha\sqrt{n}$ then

$$\Pr\left[\hat{X}_i(t_{i+1}) - \hat{X}_j(t_{i+1}) \geq \min\{(3/2)\cdot(\hat{x}_i(t_i) - \hat{x}_j(t_i)), 4\alpha\sqrt{n}\log n\}\right] \geq 1 - e^{-(\hat{x}_i(t_i)-\hat{x}_j(t_i))/\sqrt{n}} \tag{2}$$

In either case we call such subphase successful.

In the following we show that in the interval $[T_1, T_1+\tau]$ there is a sufficient amount of consecutive successful subphases such that at least one of the two opinions becomes unimportant. To do so, we define a function $f : [1, c_1 \log n] \to [0, \log\log n]$ which counts the consecutive number of successful subphases.

$$f(i) = \begin{cases} 0 & \text{if } |\hat{x}_i(t_i) - \hat{x}_j(t_i)| < 4\alpha\sqrt{n} \\ j & \text{if } (3/2)^{j-1}\cdot 4\alpha\sqrt{n} \leq |\hat{x}_i(t_i) - \hat{x}_j(t_i)| < (3/2)^j \cdot 4\alpha\sqrt{n} \end{cases}$$

Note that either Opinion i or Opinion j is unimportant at the beginning of subphase i if $f(t_i) = \log\log n$.

We define a random walk W over the state space $[0, \log\log n]$ as follows. W has a reflective state 0 and an absorbing state $\log\log n$. Initially, $W(1) = 0$. For $w \in [0, \log\log n - 1]$ the transition probabilities are defined as follows

$$\Pr[W(t+1) = 1 \mid W(t) = 0] = e^{-\frac{\alpha^2}{16}}$$
$$\Pr[W(t+1) = w+1 \mid W(t) = w] = 1 - e^{-2^w}$$
$$\Pr[W(t+1) = 0 \mid W(t) = w] = e^{-2^w}.$$

To show that either Opinion i or Opinion j becomes unimportant, which is equivalent to our function f taking on the value $\log\log n$, we define coupling between $f(i)$ and $W(i)$ such that $f(i) \geq W(i)$ for all $i \in [1, c_1 \log n]$.

For $i = 1$ the claim holds trivially since we have $W(1) = 0$ and $f(1) \geq 0$. Now assume for $i \geq 1$ that $f(i) \geq W(i)$. Now we consider two cases. In the first case assume $|\hat{x}_i(t_i) - \hat{x}_j(t_i)| < 4\alpha\sqrt{n}$. Then we know $f(i) = 0$ and hence, $W(i) = 0$. It follows from Eq. (1) and $|\hat{x}_i(t_i) - \hat{x}_j(t_i)| \geq 0$

$$\Pr[f(i+1) \geq f(i) + 1 \mid f(i) = 0] \geq e^{-\frac{\alpha^2}{16}} \text{ and}$$
$$\Pr[f(i+1) \geq 0 \mid f(i) = 0] < 1 - e^{-\frac{\alpha^2}{16}}$$

Likewise, from the definition of W it follows

$$\Pr[W(i+1) = W(i) + 1 \mid W(i) = 0] = e^{-\frac{\alpha^2}{16}} \text{ and}$$
$$\Pr[W(i+1) = 0 \mid W(i) = 0] = 1 - e^{-\frac{\alpha^2}{16}}$$

Hence, we can couple the to processes such that the following holds: whenever $W(i)$ is increased by one then $f(i)$ is increased, too. Whenever $f(i)$ is decreased $W(i)$ jumps back to zero.

In the second case we assume

$$4\alpha\sqrt{n} \leq |\hat{x}_i(t_i) - \hat{x}_j(t_i)| < \min\{\, 2(\hat{x}_i(t_i) - \hat{x}_j(t_i)), 4\alpha\sqrt{n}\log n \,\}.$$

Then it follows from Eq. (2) and $|\hat{x}_i(t_i) - \hat{x}_j(t_i)| \geq 0$

$$\Pr[f(i+1) \geq f(i) + 1 \mid f(i) = 0] \geq 1 - e^{-(\hat{x}_i(t_i)-\hat{x}_j(t_i))/\sqrt{n}} \text{ and}$$
$$\Pr[f(i+1) \geq 0 \mid f(i) = 0] < e^{-(\hat{x}_i(t_i)-\hat{x}_j(t_i))/\sqrt{n}}$$

Likewise, from the definition of W it follows

$$\Pr[W(i+1) = W(i) + 1 \mid W(i) = m] = 1 - e^{-2^m} \text{ and}$$
$$\Pr[W(i+1) = 0 \mid W(\ell) = m] = e^{-2^m}$$

Observe that

$$1 - e^{-(\hat{x}_i(t_i)-\hat{x}_j(t_i))/\sqrt{n}} \geq 1 - e^{-2^{f(i)}} \geq 1 - e^{-2^m}.$$

Again, we can couple the to processes such that $f(i) \geq W(i)$.

Finally an application of a known drift result it follows that w.h.p. there exists $i \in [1, c_1 \log n]$ such that $W(i) = \log\log n$. From this follows that there exists a time $t' \leq [T_1, T_1 + \tau]$ such that $\hat{x}_i(t') - \hat{x}_j(t') \geq 4\alpha\sqrt{n}\log n$. This implies, in turn, that at least Opinion j is unimportant. From Statement 2 in Lemma 4.3 it follows that $\hat{x}_{max}(t) - \hat{x}_j(t) \geq 2\alpha\sqrt{n}\log n$ for all $t \in [t', T_1+\tau]$ w.h.p. Hence, the Opinion j does not become significant during the time interval. Finally a union bound over all pairs of initial important opinions at time T_1 yields that all but a single opinion of those important opinions becomes insignificant in the time interval w.h.p.

Now we show that none of the unimportant opinions at time T_1 ever becomes significant during $[T_1, T_1 + \tau]$. First we fix an Opinion j which is unimportant at time T_1. Again from Statement 2 in Lemma 4.3 it follows that $\hat{x}_{max}(t) - \hat{x}_j(t) \geq 2\alpha\sqrt{n}\log n$ for all $t \in [T_1, T_1 + \tau]$ w.h.p. Hence, all unimportant opinions at time T_1 does not become significant during the time interval by a union bound. At last the statement follows because all but a single opinion becomes insignificant and hence, $T_2 - T_1 \leq \tau$. □

5 FROM ADDITIVE TO MULTIPLICATIVE BIAS (PHASE 3)

Recall that T_2 is defined as the end of Phase 2, and $\mathbf{x}(T_2)$ is a configuration with an additive bias of $\Omega(\sqrt{n}\log n)$. In the following we assume w.l.o.g. that $x_1(T_2) \geq x_2(T_2) \ldots \geq x_k(T_2)$.

We start our analysis of Phase 3 with Lemma 5.1 where we show that the support of the largest opinion does not shrink by more than a factor of two. The lemma is the equivalent to Statement 2 of Lemma 4.3 from Phase 2. The full proof can be found in the full version [4].

Lemma 5.1. *Let $c > 0$ be an arbitrary constant and define $T = c \cdot n^2 \cdot \log n / x_1(T_2)$. Then*

$$\Pr[\forall t \in [T_2, T_2 + T] : x_1(t) \geq x_1(T_1)/2] \geq 1 - n^{-5}.$$

We proceed to show that the support difference between Opinion 1 and each other opinion doubles every $O(n^2/x_1(T_2))$ interactions until the ratio between the support of both opinions is sufficiently large. This will be used in Lemma 5.3 to show that after $O(\log n \cdot n^2/x_1(T_2))$ interactions we reach w.h.p. a configuration with a constant factor multiplicative bias.

Lemma 5.2. *Fix an Opinion $i \neq 1$ and assume there exists $t_0 \geq T_2$ with $x_i(t_0) \geq 20\sqrt{n\log n}$ and $x_1(t_0) - x_i(t_0) \geq \alpha\sqrt{n}\log n$. Let $T = 420 \cdot n^2/x_1(T_2)$. Then*

$$\Pr[\exists t \in [t_0, t_0+T] : x_1(t) - x_i(t) \geq \min\{2(x_1(t_0) - x_i(t_0)),\ 3x_i(t)\} \text{ or } x_i(t) < 20\sqrt{n\log n}] \geq 1 - 2n^{-3}.$$

Proof Sketch. The proof follows the analysis of the classical Gambler's ruin problem. That is, starting with $\Delta = x_1(t) - x_i(t)$ we track the evolution of this quantity throughout a sequence of $O(n^2/x_1(t))$ interactions and show that it reaches 2Δ before $\Delta/2$. Here we rely on the bounds on the number of undecided agents (Lemma 3.3 and Lemma 4.1) and on the lower bound on Opinion 1 which holds w.h.p. during time $[T_2, T_2 + 420 \cdot n^2 \cdot \log n/x_1(T_2)]$ (Lemma 4.2). The full proof can be found in the full version [4]. □

Now we are ready to analyze the running time of Phase 3.

Lemma 5.3. *Assume that $\mathbf{x}(T_2)$ is a configuration with $x_1(T_2) - x_i(T_2) \geq \alpha\sqrt{n}\log n$ for all $i \neq 1$. Let*

$$T_3 = \inf\{\, t \geq T_2 \mid \forall i \neq 1 : x_1(t) \geq 2x_i(t) \,\}.$$

Then

$$\Pr\left[T_3 - T_2 \leq 420 \cdot n^2 \cdot \log n / x_1(T_2)\right] \geq 1 - 2n^{-2}.$$

Proof. The main idea of this proof is to repeatedly apply Lemma 5.2 to each Opinion $i \neq 1$ until either the support of Opinion 1 becomes larger than $2n/3$ or the support of Opinion i becomes less than $20 \cdot \sqrt{n\log n}$. In both cases it then follows that the ratio between the support of Opinion 1 and Opinion i is larger than two, and there is a time where there is a multiplicative bias between the first opinion and each other opinion.

Let

$$\hat{T} = \inf\{t \geq T_2 \mid u(t) \notin [(n - x_{\max}(t))/2 - 8 \cdot \sqrt{n\ln n}, n/2] \text{ or } x_1(t) < x_1(T_2)/2.\}$$

be a stopping time. We define $(\hat{\mathbf{X}})_{t\geq T_2}$ as the process with $\hat{\mathbf{X}}(t) = \mathbf{X}(t)$ for all $t \leq \hat{T}$ and $\hat{\mathbf{X}}(t) = \mathbf{X}(\hat{T})$ for $t > \hat{T}$. From Lemma 4.1 it follows that $u(t) \geq (n - x_{\max}(t))/2 - 8 \cdot \sqrt{n\ln n}$ for all $t \in [T_2, n^3]$, w.h.p. From Lemma 3.3 it follows that $u(t) \leq n/2$ for all $t \in [T_2, n^3]$, w.h.p. Finally, Lemma 5.1 gives us that $x_{\max}(t) \geq x_{\max}(T_2)/2$ for all $t \in [T_2, T_2 + cn^2\log n/x_{\max}(T_2)]$, w.h.p. Thus, $\hat{T} - T_2 = \Omega(n^2 \cdot \log n/x_1(T_2))$ w.h.p. and we can assume that $(\mathbf{X})_{t\geq T_2}$ and $(\hat{\mathbf{X}})_{t\geq T_2}$ are identical for $t \in [T_2, T_2 + O(n^2 \cdot \log n/x_{\max}(T_2))]$.

Let $\tau = 420 \cdot n^2 \cdot \log n/x_{\max}(T_2)$ and fix an Opinion $i \neq 1$ at time T_2 with $x_i(T_2) \geq 20\sqrt{n\log n}$. We divide the interactions from $[T_2, T_2 + \tau]$ into $\log n$ subphases of length $420 \cdot n^2/\hat{x}_1(T_2)$ each. For $1 \leq j \leq \log n$ we define $\ell_1 = 1$ and $\ell_j = 1 + (j-1) \cdot 420 \cdot n^2/\hat{x}_1(T_2)$. Then the jth subphase contains interactions ℓ_j to $(\ell_{j+1} - 1)$. Furthermore, we define t_j is the first interaction in subphase j. Now fix an arbitrary subphase j. It follows from Lemma 5.2 that there exists a time $t' \in [t_j, t_{j+1}]$ such that w.h.p. either $\hat{x}_1(t) - \hat{x}_i(t) \geq \min\{2 \cdot (\hat{x}_1(t_j) - \hat{x}_i(t_j)),\ 3 \cdot \hat{x}_i(t)\}$ or $\hat{x}_i(t) < 20\sqrt{n\log n}$.

We apply Lemma 5.2 to each subphase. From the union bound over all subphases and all opinions it follows that after at most $\log n$ subphases w.h.p. there exists for each Opinion i a time $t'_i \in [T_2, T_2 + \tau]$ with either (a) $\hat{x}_1(t'_i) - \hat{x}_i(t'_i) \geq 2n/3$ or (b) $\hat{x}_i(t'_i) < 20\sqrt{n\log n}$ or (c) $\hat{x}_1(t'_i) \geq 4 \cdot \hat{x}_i(t'_i)$. In the following we consider three cases.

Case (a). There exists an Opinion $i \neq 1$ such that $\hat{x}_1(t'_i) - \hat{x}_i(t'_i) \geq 2n/3$. Hence, we have at t'_i a constant multiplicative bias between Opinion 1 and all other opinions $i \neq 1$. From this the statement follows immediately with $T_3 = t'_i$.

Case (b). For Opinion i there exists a t'_i such that $\hat{x}_i(t'_i) < 20\sqrt{n\log n}$. From Lemma 4.3(1) it follows that $\hat{x}_i(t) \leq 40\sqrt{n\log n}$ for all $t \in [t'_i, T_2 + \tau]$ w.h.p. Additionally we know $\hat{x}_1(t) \geq \hat{x}_1(T_2)/2 \geq c'\sqrt{n}\log^2 n$ for all $t \in [t'_i, T_2 + \tau]$. Hence, $\hat{x}_1(t)/\hat{x}_i(t) \gg 2$ for all $t \in [t'_i, T_2 + \tau]$ and, from the viewpoint of Opinion i we have that T_3 can take on an arbitrary value in $[t'_i, T_2 + \tau]$.

Case (c). For Opinion i there exists a t'_i such that $\hat{x}_1(t'_i) \geq 4 \cdot \hat{x}_i(t'_i)$. From the claim below it follows that w.h.p. $\hat{x}_1(t) \geq 2\hat{x}_i(t)$ for all $t \in [t'_i, T_2 + \tau]$ and from the viewpoint of Opinion i we have that T_3 can take on an arbitrary value in $[t'_i, T_2 + \tau]$.

Now Lemma 5.3 follows either immediately from Case (a). Or we can apply Case (b) or Case (c) for each Opinion $i \neq 1$ and then we can choose $T_3 = T_2 + \tau$. It remains to show the following claim. The proof can be found in the full version [4].

Claim 5.3.1. *Let j be an arbitrary subphase and let $t_0 \in [t_j, t_{j+1}]$. Fix an Opinion i and assume $\hat{x}_i(t_0) \in [20 \cdot \sqrt{n \log n}, \hat{x}_1(t_0)/4]$. Then $\hat{x}_1(t) \geq 2 \cdot \hat{x}_i(t)$ for all $t \in [t_0, T_2 + \tau]$.* □

6 FROM MULTIPLICATIVE BIAS TO ABSOLUTE MAJORITY (PHASE 4)

Recall that T_3 is the end of Phase 3 and $\mathbf{X}(T_3)$ is a configuration with multiplicative bias. In this version of the paper we assume that the bias is at least two, the proof of the case of a $(1+\varepsilon)$-bias for any constant ε is deferred to the full version of this paper [4]. It follows from a slightly more involved calculation. In the following we assume w.l.o.g. that $x_1(T_3) > x_2(T_3) \ldots \geq x_k(T_3)$. The main result for this phase is Lemma 6.4, where we show that the multiplicative bias is grown into a unique majority opinion with support at least $2n/3$ within $O(n \log n + n^2/x_1(T_3))$ interactions, w.h.p. To do so we first need an improved bound on the number of undecided agents which we reach at time $T_3 + O(n \log n)$. Additionally we have to show that in the meantime both x_1 and the multiplicative bias decrease only by a small constant fraction (Lemma 6.1 and Lemma 6.2). The proofs of both lemmas are similar to the proofs of Lemma 4.2 and Claim 5.3.1, respectively, and can be found in the full version [4].

Lemma 6.1. *Let $c > 0$ be an arbitrary constant and define $T = c \cdot n^2 \log n / x_1(T_3)$. Then*

$$\Pr[\forall t \in [T_3, T_3 + T] \colon x_1(t) \geq x_1(T_3)/2] \geq 1 - n^{-5}.$$

Lemma 6.2. *Assume that $\mathbf{x}(T_3)$ is a configuration with $x_1(T_3) \geq 2 \cdot x_i(T_3)$ for all $i \neq 1$. Then for all $i \neq i$ it holds*

$$\Pr\left[\forall t \in [T_3, 111 \cdot n^2/x_1(T_3)] \colon x_1(t) \geq 7/4 \cdot x_i(t)\right] \geq 1 - 2n^{-3}.$$

Next we improve the lower bound on the number of undecided agents from Lemma 3.3. Recall that T_4 is the end of Phase 4, defined as $T_4 = \inf\{\, t \geq T_3 \mid x_1(t) \geq 2n/3 \,\}$.

Lemma 6.3. *Let $T_u = \inf\{\, t \geq T_3 \mid u(t) \geq n/2 - 7/8 \cdot x_1(t) \,\}$. Then*

$$\Pr[\min(T_4, T_u) - T_3 \leq \lceil 7n \ln n \rceil] \geq 1 - 4n^{-3}.$$

Proof Sketch. The proof is similar to the proof of Lemma 3.1. The main difference is that we use a modified potential function $Z(t) = n - 2u(t) - 7/8 \cdot x_1(t)$ instead of $Z(t) = n - 2u(t) - x_1(t)$. The expression for the expected drift of this modified potential function becomes slightly more complicated, and to bound it we require the multiplicative bias from Lemma 6.2. The full proof can be found in the full version [4]. □

Now we are ready to analyze the running time of Phase 4.

Lemma 6.4. *Assume that $\mathbf{x}(T_3)$ is a configuration with $x_1(T_3) \geq 2 \cdot x_i(T_3)$ for all $i \neq 1$. Then there exists a constant c such that*

$$\Pr\left[T_4 - T_3 \leq 7n \ln n + 444 \cdot n^2/x_1(T_3)\right] \geq 1 - 2n^{-2}.$$

Proof. To show the statement we require the following two auxiliary results. First we establish in Claim 6.4.1 that the improved bound on the undecided agents from Lemma 6.3 holds throughout the remainder of the phase. As before, we define $T_u = \inf\{\, t \geq T_3 \mid u(t) \geq n/2 - 7/8 \cdot x_1(t)/2 \,\}$ and recall that T_4 denotes the end of the phase. The proof follows along the lines of the proof of Lemma 4.1 with the new $Z(t)$, and can be found in in the full version [4].

Claim 6.4.1.

$$\Pr\left[\forall t \in [T_u, \min\{n^3, T_4\}] \colon u(t) \geq \frac{n}{2} - \frac{7}{16} x_1(t) - 8\sqrt{n \ln n}\right] \geq 1 - 4n^{-3}.$$

Next, in Claim 6.4.2 we bound the number of interactions until the support of Opinion 1 has doubled. Similarly to Lemma 5.2, the proof uses the classical gambler's ruin problem to show that in a sequence of $c \cdot n^2/x_1(t)$ interactions the support of Opinion 1 doubles w.h.p. before it halves. The full proof can be found in the full version [4].

Claim 6.4.2. *Let $\mathbf{x}(t)$ be a configuration with $u(t) \geq n/2 - 7/16 \cdot x_1(t) - 8 \cdot \sqrt{n \ln n}$ and $x_1(t) < 2n/3$. We define $t' = c \cdot n^2/x_1(t)$ for a suitable chosen constant c. Then*

$$\Pr\left[\exists t' \in [t, t + t'] \colon x_1(t') \geq 2x_1(t) \;\text{ or }\; x_1(t) \geq 2n/3\right] \geq 1 - n^{-3}.$$

With these two auxiliary claims we are now ready to show the lemma. We start with a brief overview of the proof. The proof is similar to the proof of Lemma 5.3 but we only have to consider the analog to Case (a). We repeatedly apply Claim 6.4.2 to Opinion 1. Then the support of the largest opinion, $x_1(t)$ doubles every $O(n^2/x_1(t))$ interactions until its support becomes larger than $2n/3$. After doubling at most $\log n$ times, we reach a configuration where $x_1(t) \geq 2n/3$. This will be our time T_4.

To show that there exists a t with $x_1(t) \geq 2n/3$ we define

$$\hat{T} = \inf\{t \geq T_3 + t_0 | u(t) \notin [(n - 7/16 \cdot x_1(t)) - 8 \cdot \sqrt{n \ln n}, n/2] \text{ or } x_1(t) < x_1(T_3)/2\}$$

as a stopping time. Here t_0 is defined as

$$t_0 = \inf\{\, t \colon u(t_0) \geq (n - 7/16 \cdot x_1(t_0)) \,\}.$$

From Lemma 6.3 it follows w.h.p. that $t_0 \leq T_3 + 7n \ln n$.

Let $(\hat{\mathbf{X}}(t))_{t \geq T_3 + t_0}$ denote the process with $\hat{\mathbf{X}}(t) = \mathbf{X}(t)$ for all $t \leq \hat{T}$ and $\hat{\mathbf{X}}(t) = \mathbf{X}(\hat{T})$ for $t > \hat{T}$. From Claim 6.4.1 it follows that $u(t) \geq (n - x_1(T_3 + t_0))/2 - 8 \cdot \sqrt{n \ln n}$ for all $t \in [T_3 + t_0, n^3]$, w.h.p. From Lemma 3.3 it follows that $u(t) \leq n/2$ for all $t \in [T_1, n^3]$, w.h.p. Finally, Lemma 4.2 gives us that $x_1(t) \geq x_1(T_3)/2$ for all $t \in [T_3, T_3 + cn^2 \log n / x_1(T_3)]$, w.h.p. Thus, $\hat{T} - (T_3 + t_0) = \Omega(n^2/x_1(T_3))$ w.h.p. and we can assume that $(\mathbf{X})_{t \geq T_3 + t_0}$ and $(\hat{\mathbf{X}})_{t \geq T_3 + t_0}$ are identical for $t \in [T_3 + t_0, T_3 + t_0 + O(n^2 \cdot \log n / x_1(T_3))]$.

To track the progress of Opinion 1 we divide the interactions from $[T_3 + t_0, T_3 + t_0 + c \cdot n^2/x_1(T_3)]$ into subphases of varying length. Let $T_{(0)} = T_3 + t_0$ and define for $1 \leq \ell \leq \log n$

$$T_{(\ell)} = \inf\{\, t \geq T_{(0)} \mid \hat{x}_1(t) \geq 2^{\ell} \cdot \hat{x}_1(T_{(0)}) \text{ or } \hat{x}_1(t) \geq 2n/3 \,\}.$$

We call the interactions in the interval $\left[T_{(\ell-1)}, T_{(\ell)}\right)$ *subphase ℓ*. Note that by definition of $T_{(\ell)}$, the support of x_1 doubles in every subphase (or $x_1 \geq 2/3n$ and Phase 4 ends). In more detail, for a fixed but arbitrary subphase ℓ it follows from Claim 6.4.2 that the length of subphase ℓ is at most $c \cdot n^2/\hat{x}(T_{(\ell-1)}) \le c \cdot n^2/(2^{\ell-1} \cdot \hat{x}_1(T_{(0)})$, w.h.p. Hence, it follows that there exists a time $t' \in [T_{(\ell-1)}, T_{(\ell-1)} + c \cdot n^2/\hat{x}(T_{(\ell-1)})]$ such that $\hat{x}_1(t') \ge 2^{\ell} \cdot \hat{x}_1(T_{(0)})$ or $\hat{x}_1(t') \ge 2/3 \cdot n$, w.h.p. From the union bound over all subphases we get that after at most $\log n$ subphases there exists w.h.p. a time $t' \in [T_{(0)}, T_{(\log n)}]$ such that $\hat{x}_1(t') \ge 2n/3$. This holds since otherwise $\hat{x}_1(t') \ge 2^{\log n} \cdot \hat{x}_1(T_{(0)}) \ge n \cdot \sqrt{n} \log^2 n > n.$, a contradiction.

Summing up the length of all subphases for $c = 111$ gives us

$$\sum_{i=1}^{\log n} \frac{c \cdot n^2}{2^{i-1} \cdot \hat{x}_1(T_{(0)})} = \frac{c \cdot n^2}{\hat{x}_1(T_{(0)})} \cdot \sum_{i=1}^{\log n} \frac{1}{2^{i-1}} \le 2 \cdot c \cdot \frac{n^2}{\hat{x}_1(T_{(0)})}$$

and hence, $T_4 - T_3 \le 7n \ln n + 4 \cdot c \cdot n^2/\hat{x}_1(T_3)$ as claimed. □

7 FROM ABSOLUTE MAJORITY TO CONSENSUS (PHASE 5)

Recall that T_4 is the end of Phase 4 and $\mathbf{X}(T_4)$ is a configuration where the support of the largest opinion, $x_{\max}(T_4)$ is at least $2n/3$. The fifth phase ends when all agents agree on the Opinion $\max(T_4)$. In the following we assume w.l.o.g. that $x_1(T_4) \ge x_2(T_4) \ge \ldots \ge x_k(T_4)$. In Lemma 7.1 we show that the running time of this phase is $O(n \log n)$. This result is shown via a coupling where we couple the USD on configuration $\mathbf{X}(T_4)$ with the USD on a configuration $\tilde{\mathbf{X}}$ with two opinions only. We show that the time until all agents agree starting from configuration $\mathbf{X}(T_4)$ is majorized by the time starting from configuration $\tilde{\mathbf{X}}$ (see Lemma 7.2).

Lemma 7.1. *Assume that* $\mathbf{x}(T_4)$ *is a configuration with* $x_1(T_4) \ge 2n/3$. *Let* $T_5 = \inf\{t \ge T_4 \mid x_1(t) = n\}$.

$$\Pr[T_5 - T_4 \le c \cdot n \log n] \ge 1 - n^{-3}.$$

W.l.o.g. we assume that $T_4 = 0$. To show this lemma we couple our process $(\mathbf{X}(t))_{t\in\mathbb{N}}$ with k opinions with a process $(\tilde{\mathbf{X}}(t))_{t\in\mathbb{N}}$ with 2-opinions. $\tilde{\mathbf{x}}$ is defined as follows. $\tilde{x}_1(0) = x_1(0), \tilde{x}_2(0) = \sum_{i=2}^{k} x_i(0)$ and $\tilde{u}(0) = u(0)$. We will show in Lemma 7.2 that there exists a coupling such that for all $t \ge 0$ we have $x_1(t) \ge \tilde{x}_1(t)$. From this follows that for all $t \ge 1$ we have $\Pr[x_1(t) = n] \ge \Pr[\tilde{x}_1 = n]$.

Since $x_1(0) \ge 2n/3$, we have $\tilde{x}_1(0) - \tilde{x}_2(0) \ge n/3$. For $k = 2$ it follows from [5] that $(\tilde{\mathbf{X}}(t))_{t\in\mathbb{N}}$ converges w.h.p. to $\tilde{x}_1(t) = n$ in $O(n \log n)$ interaction.

For technical reasons, we show the stronger invariant $x_1(t) \ge \tilde{x}_1(t)$ and $x_1(t) + u(t) \ge \tilde{x}_1(t) + \tilde{u}(t)$ using a step-by-step coupling.

Lemma 7.2. *Consider the two processes* $(\mathbf{X}(t))_{t\in\mathbb{N}}$ *with* k *opinions and* $(\tilde{\mathbf{X}}(t))_{t\in\mathbb{N}}$ *with 2-opinions where* $\tilde{x}_1(0) = x_1(0), \tilde{x}(2) = \sum_{i=2}^{k} x_i(0)$ *and* $\tilde{u}(0) = u(0)$. *Then there exists a coupling between the two processes such that* $\forall t \ge 0$:

$$x_1(t) \ge \tilde{x}_1(t) \text{ and } x_1(t) + u(t) \ge \tilde{x}_1(t) + \tilde{u}(t) \tag{3}$$

Proof. We prove the lemma by induction over t. Obviously, the claim holds for $t = 0$. Fix a time step t where (3) holds. We show that (3) holds for time $t + 1$. In the following, we omit t if clear from the context. We represent the n agents of the configurations $\mathbf{x}(t)$ and $\tilde{\mathbf{x}(t)}$ by n-dimensional vectors $\mathbf{v}(t)$ and $\tilde{\mathbf{v}}(t)$ which are sorted as follows.

$$\tilde{v}_i(t) = \begin{cases} 1 & \text{if } i \in [\tilde{x}_1] \\ \perp & \text{if } i - \tilde{x}_1 \in [u] \\ 2 & \text{if } i - \tilde{x}_1 - u \in \left[\sum_{j=2}^{k} x_j\right] \\ \perp & \text{if } i - \tilde{x}_1 - \tilde{x}_2 \in [u+1, \tilde{u}] \\ 2 & \text{otherwise} \end{cases}$$

$$v_i(t) = \begin{cases} 1 & \text{if } i \in [\tilde{x}_1] \\ \perp & \text{if } i - \tilde{x}_1 \in [u] \\ 2, \ldots, k & \text{if } i - \tilde{x}_1 - u \in \left[\sum_{j=2}^{k} x_j\right] \\ 1 & \text{if } i - \tilde{x}_1 - u - \tilde{x}_2 \in [x_1 - \tilde{x}_1] \\ \perp & \text{otherwise} \end{cases}$$

The definition results in the following two cases.

Case 1: $\tilde{u}(t) \ge u(t)$. In this case the vectors are sorted as follows.

$$\begin{array}{r c|c|c c c c||c|c} \tilde{\mathbf{v}} = & 1\ldots1 & \perp\ldots\perp & 2\ldots2 & 2\ldots2 & \ldots & 2\ldots2 & \perp\ldots\perp & 2\ldots2 \\ \hline \mathbf{v} = & 1\ldots1 & \perp\ldots\perp & 2\ldots2 & 3\ldots3 & \ldots & k\ldots k & 1\ldots1 & 1\ldots1 \\ & |\leftarrow & & & \text{a} & & \rightarrow| & & \end{array}$$

Case 2: $\tilde{u}(t) < u(t)$. In this case the vectors are sorted as follows.

$$\begin{array}{r c|c|c c c c||c|c} \tilde{\mathbf{v}} = & 1\ldots1 & \perp\ldots\perp & 2\ldots2 & 2\ldots2 & \ldots & 2\ldots2 & 2\ldots2 & 2\ldots2 \\ \hline \mathbf{v} = & 1\ldots1 & \perp\ldots\perp & 2\ldots2 & 3\ldots3 & \ldots & k\ldots k & 1\ldots1 & \perp\ldots\perp \\ & |\leftarrow & & & \text{a} & & \rightarrow| & & \end{array}$$

We will use the identity coupling, both processes choose the same pair $(i, j) \in [n]^2$ uniformly at random. Hence, the next interaction is $(v_i(t), v_j(t))$ in the k-opinion process and $(\tilde{v}_i(t), \tilde{v}_j(t))$ in the two-opinion process. Let $a = \tilde{x}_1(t) + \min(u(t), \tilde{u}(t)) + \sum_{j=2}^{k} x_j(t)$. We split the analysis into four cases and exemplify the proof on the first case $i, j \le a$.

Case 1: $i, j \le a$. If $\tilde{v}_i(t) = v_i(t)$ and $\tilde{v}_j(t) = v_j(t)$ both processes perform the same transition and the inductive step is trivially fulfilled. If $\tilde{v}_i(t) \ne v_i(t)$ and $\tilde{v}_j(t) \ne v_j(t)$ we have that $\tilde{v}_i(t), \tilde{v}_j(t) = 2$ and $v_i(t), v_j(t) > 2$. Then $v_i(t+1) = \perp$ and $v_i(t+1) = 2$ resulting in $u(t+1) \ge u(t)$. We have two cases.

If $\tilde{v}_i(t) = v_i$ but $\tilde{v}_j(t) \ne v_j(t)$ the following transitions are possible.

$v_i(t)$	$\tilde{v}_i(t)$	$v_j(t)$	$\tilde{v}_j(t)$	$v_i(t+1)$	$\tilde{v}_i(t+1)$
1	1	>2	2	$\perp$	$\perp$
$\perp$	$\perp$	>2	2	>2	2
2	2	>2	2	$\perp$	2

If $\tilde{v}_i(t) \ne v_i(t)$ and $\tilde{v}_j(t) = v_j(t)$ the following transitions are possible.

$v_i(t)$	$\tilde{v}_i(t)$	$v_j(t)$	$\tilde{v}_j(t)$	$v_i(t+1)$	$\tilde{v}_i(t+1)$
>2	2	1	1	$\perp$	$\perp$
>2	2	$\perp$	$\perp$	>2	2
>2	2	2	2	$\perp$	2

The inductive step now follows in both cases since in all cases $u(t) \ge \tilde{u}(t)$. The remaining cases follow analogously and can be found in the full version [4]. □

8 COMPARISON OF CONVERGENCE RATES WITH BECCHETTI ET AL. [10]

We show that given an initial configuration with a multiplicative bias, our convergence rate from Theorem 2.1 improves over the analogous rate from Becchetti et al. [10] whenever the initial support of the largest opinion x_1 is close to the average opinion size, that is, $x_1 \le n/k \cdot \log n$.

In the regime of an initial multiplicative bias, the analysis of Becchetti et al. of the USD in the gossip model shows the process achieves plurality consensus in $\mathrm{O}(\mathrm{md}(\mathbf{x}(0)) \cdot \log n)$ rounds, where (assuming x_1 has largest initial support)

$$\mathrm{md}(\mathbf{x}(0)) = \sum_{i \in [k]} \left(\frac{x_i(0)}{x_1(0)}\right)^2 .$$

On the other hand, recall our result from Theorem 2.1, which shows convergence towards plurality consensus in the population protocol model in $\mathrm{O}(n \log n + n^2/x_1(0))$ interactions, which is equivalent to $\mathrm{O}(\log n + n/x_1(0))$ *parallel* time.

Considering the range of k for which their result holds, our convergence rate improves the one of Becchetti et al. when $x_1 \le n/k \cdot \log n$. To see this, consider an initial configuration $\mathbf{x}$ and assume that w.l.o.g. $x_1 \ge x_i$ for all $2 \le i \le n$. We calculate

$$\mathrm{md}(\mathbf{x}) \log n = \sum_{i=1}^{k} x_i^2/x_1^2 \log n \ge \frac{k \cdot (n/k)^2}{x_1^2} \log n = \frac{n^2}{k \cdot x_1^2} \log n$$
$$= \frac{n \cdot \log n}{k \cdot x_1} \cdot \frac{n}{x_1}.$$

Hence, $\mathrm{md}(\mathbf{x}) \log n$ gives the better running time when

$$x_1 > \frac{n \cdot \log n}{k}.$$

9 CONCLUSIONS AND OPEN PROBLEMS

We show fast convergence of the USD in the population model, where the exact convergence rates depend on the magnitude of support of the initial plurality opinion and the type of bias (if any) in the initial configuration. Although our result can be viewed as an improvement over the existing, analogous convergence rates for the process in the gossip model [10], our analysis does not readily transfer to that model. Thus it remains open to prove convergence of the $k > 2$ opinion USD with *no initial bias* in the gossip model, and moreover to understand whether there exists a unified framework for analyzing the process in both models simultaneously. Additionally, analyzing the process for $k = \omega(\sqrt{n}/\log^2 n)$ opinions is left for future work.

Separately, we leave as future work analyzing the k-state USD in the presence of adversarial nodes or communication noise. Recent results of d'Amore et al. [26, 27] and of Cruciani et al. [25], which analyze the 2-state USD process (as well as other majority dynamics for $k > 2$) under such settings, suggest that the k opinion USD process is also robust to these noise models. Quantifying the effect of such noise on the convergence rate of the k opinion USD process (for $k > 2$) is thus an interesting open question.

ACKNOWLEDGMENTS

This work is supported by DFG Research Group ADYN under grant DFG 411362735.

REFERENCES

[1] Dan Alistarh, James Aspnes, and Rati Gelashvili. 2018. Space-Optimal Majority in Population Protocols. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*. SIAM, 2221–2239. https://doi.org/10.1137/1.9781611975031.144

[2] Dan Alistarh and Rati Gelashvili. 2018. Recent Algorithmic Advances in Population Protocols. *SIGACT News* 49, 3 (2018), 63–73. https://doi.org/10.1145/3289137.3289150

[3] Dan Alistarh, Rati Gelashvili, and Milan Vojnovic. 2015. Fast and Exact Majority in Population Protocols. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC 2015*. ACM, 47–56. https://doi.org/10.1145/2767386.2767429

[4] Talley Amir, James Aspnes, Petra Berenbrink, Felix Biermeier, Christopher Hahn, Dominik Kaaser, and John Lazarsfeld. 2023. Fast Convergence of k-Opinion Undecided State Dynamics in the Population Protocol Model. *CoRR* abs/2302.12508 (2023). https://doi.org/10.48550/arXiv.2302.12508 arXiv:2302.12508

[5] Dana Angluin, James Aspnes, and David Eisenstat. 2008. A simple population protocol for fast robust approximate majority. *Distributed Computing* 21, 2 (2008), 87–102. https://doi.org/10.1007/s00446-008-0059-z

[6] Gregor Bankhamer, Petra Berenbrink, Felix Biermeier, Robert Elsässer, Hamed Hosseinpour, Dominik Kaaser, and Peter Kling. 2022. Fast Consensus via the Unconstrained Undecided State Dynamics. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022*. SIAM, 3417–3429. https://doi.org/10.1137/1.9781611977073.135

[7] Gregor Bankhamer, Petra Berenbrink, Felix Biermeier, Robert Elsässer, Hamed Hosseinpour, Dominik Kaaser, and Peter Kling. 2022. Population Protocols for Exact Plurality Consensus: How a small chance of failure helps to eliminate insignificant opinions. In *PODC '22: ACM Symposium on Principles of Distributed Computing*. ACM, 224–234. https://doi.org/10.1145/3519270.3538447

[8] Gregor Bankhamer, Robert Elsässer, Dominik Kaaser, and Matjaz Krnc. 2020. Positive Aging Admits Fast Asynchronous Plurality Consensus. In *PODC '20: ACM Symposium on Principles of Distributed Computing*. ACM, 385–394. https://doi.org/10.1145/3382734.3406506

[9] Luca Becchetti, Andrea E. F. Clementi, and Emanuele Natale. 2020. Consensus Dynamics: An Overview. *SIGACT News* 51, 1 (2020), 58–104. https://doi.org/10.1145/3388392.3388403

[10] Luca Becchetti, Andrea E. F. Clementi, Emanuele Natale, Francesco Pasquale, and Riccardo Silvestri. 2015. Plurality Consensus in the Gossip Model. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015*. SIAM, 371–390. https://doi.org/10.1137/1.9781611973730.27

[11] Luca Becchetti, Andrea E. F. Clementi, Emanuele Natale, Francesco Pasquale, Riccardo Silvestri, and Luca Trevisan. 2017. Simple dynamics for plurality consensus. *Distributed Computing* 30, 4 (2017), 293–306. https://doi.org/10.1007/s00446-016-0289-4

[12] Stav Ben-Nun, Tsvi Kopelowitz, Matan Kraus, and Ely Porat. 2020. An $O(\log^{3/2} n)$ Parallel Time Population Protocol for Majority with $O(\log n)$ States. In *PODC '20: ACM Symposium on Principles of Distributed Computing*. ACM, 191–199. https://doi.org/10.1145/3382734.3405747

[13] Petra Berenbrink, Andrea E. F. Clementi, Robert Elsässer, Peter Kling, Frederik Mallmann-Trenn, and Emanuele Natale. 2017. Ignore or Comply?: On Breaking Symmetry in Consensus. In *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC 2017*. ACM, 335–344. https://doi.org/10.1145/3087801.3087817

[14] Petra Berenbrink, Robert Elsässer, Tom Friedetzky, Dominik Kaaser, Peter Kling, and Tomasz Radzik. 2018. A Population Protocol for Exact Majority with O(log5/3 n) Stabilization Time and Theta(log n) States. In *32nd International Symposium on Distributed Computing, DISC 2018 (LIPIcs, Vol. 121)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 10:1–10:18. https://doi.org/10.4230/LIPIcs.DISC.2018.10

[15] Petra Berenbrink, Robert Elsässer, Tom Friedetzky, Dominik Kaaser, Peter Kling, and Tomasz Radzik. 2021. Time-space trade-offs in population protocols for the majority problem. *Distributed Computing* 34, 2 (2021), 91–111. https://doi.org/10.1007/s00446-020-00385-0

[16] Petra Berenbrink, Tom Friedetzky, George Giakkoupis, and Peter Kling. 2016. Efficient Plurality Consensus, Or: the Benefits of Cleaning up from Time to Time. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy (LIPIcs, Vol. 55)*, Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 136:1–136:14. https://doi.org/10.4230/LIPIcs.ICALP.2016.136

[17] Petra Berenbrink, George Giakkoupis, Anne-Marie Kermarrec, and Frederik Mallmann-Trenn. 2016. Bounds on the Voter Model in Dynamic Networks.

In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 146:1–146:15. https://doi.org/10.4230/LIPIcs.ICALP.2016.146

[18] Stephen P. Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. 2006. Randomized gossip algorithms. *IEEE Transactions on Information Theory* 52, 6 (2006), 2508–2530. https://doi.org/10.1109/TIT.2006.874516

[19] Andrea E. F. Clementi, Mohsen Ghaffari, Luciano Gualà, Emanuele Natale, Francesco Pasquale, and Giacomo Scornavacca. 2018. A Tight Analysis of the Parallel Undecided-State Dynamics with Two Colors. In *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2018 (LIPIcs, Vol. 117)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 28:1–28:15. https://doi.org/10.4230/LIPIcs.MFCS.2018.28

[20] Anne Condon, Monir Hajiaghayi, David G. Kirkpatrick, and Ján Manuch. 2017. Simplifying Analyses of Chemical Reaction Networks for Approximate Majority. In *DNA Computing and Molecular Programming - 23rd International Conference, DNA 23 (Lecture Notes in Computer Science, Vol. 10467)*. Springer, 188–209. https://doi.org/10.1007/978-3-319-66799-7_13

[21] Colin Cooper, Robert Elsässer, Hirotaka Ono, and Tomasz Radzik. 2012. Coalescing random walks and voting on graphs. In *ACM Symposium on Principles of Distributed Computing, PODC '12*. ACM, 47–56. https://doi.org/10.1145/2332432.2332440

[22] Colin Cooper, Robert Elsässer, and Tomasz Radzik. 2014. The Power of Two Choices in Distributed Voting. In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014*. Springer, 435–446. https://doi.org/10.1007/978-3-662-43951-7_37

[23] Colin Cooper, Robert Elsässer, Tomasz Radzik, Nicolas Rivera, and Takeharu Shiraga. 2015. Fast Consensus for Voting on General Expander Graphs. In *Distributed Computing - 29th International Symposium, DISC 2015*. Springer, 248–262. https://doi.org/10.1007/978-3-662-48653-5_17

[24] Colin Cooper, Tomasz Radzik, Nicolas Rivera, and Takeharu Shiraga. 2017. Fast Plurality Consensus in Regular Expanders. In *31st International Symposium on Distributed Computing, DISC 2017*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 13:1–13:16. https://doi.org/10.4230/LIPIcs.DISC.2017.13

[25] Emilio Cruciani, Hlafo Alfie Mimun, Matteo Quattropani, and Sara Rizzo. 2021. Phase Transitions of the k-Majority Dynamics in a Biased Communication Model. In *ICDCN '21: International Conference on Distributed Computing and Networking, Virtual Event, Nara, Japan, January 5-8, 2021*. ACM, 146–155. https://doi.org/10.1145/3427796.3427811

[26] Francesco D'Amore, Andrea E. F. Clementi, and Emanuele Natale. 2020. Phase Transition of a Non-linear Opinion Dynamics with Noisy Interactions - (Extended Abstract). In *Structural Information and Communication Complexity - 27th International Colloquium, SIROCCO 2020, Paderborn, Germany, June 29 - July 1, 2020, Proceedings (Lecture Notes in Computer Science, Vol. 12156)*, Andrea Werneck Richa and Christian Scheideler (Eds.). Springer, 255–272. https://doi.org/10.1007/978-3-030-54921-3_15

[27] Francesco D'Amore and Isabella Ziccardi. 2022. Phase Transition of the 3-Majority Dynamics with Uniform Communication Noise. In *Structural Information and Communication Complexity - 29th International Colloquium, SIROCCO 2022, Paderborn, Germany, June 27-29, 2022, Proceedings (Lecture Notes in Computer Science, Vol. 13298)*, Merav Parter (Ed.). Springer, 98–115. https://doi.org/10.1007/978-3-031-09993-9_6

[28] Benjamin Doerr, Leslie Ann Goldberg, Lorenz Minder, Thomas Sauerwald, and Christian Scheideler. 2011. Stabilizing consensus with the power of two choices. In *SPAA 2011: Proceedings of the 23rd Annual ACM Symposium on Parallelism in Algorithms and Architectures*. ACM, 149–158. https://doi.org/10.1145/1989493.1989516

[29] David Doty, Mahsa Eftekhari, Leszek Gasieniec, Eric E. Severson, Przemyslaw Uznanski, and Grzegorz Stachowiak. 2021. A time and space optimal stable population protocol solving exact majority. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021*. IEEE, 1044–1055. https://doi.org/10.1109/FOCS52979.2021.00104

[30] Robert Elsässer and Tomasz Radzik. 2018. Recent Results in Population Protocols for Exact Majority and Leader Election. *Bulletin of the EATCS* 126 (2018), 34 pages. http://bulletin.eatcs.org/index.php/beatcs/article/view/549/546

[31] Leszek Gasieniec, David D. Hamilton, Russell Martin, Paul G. Spirakis, and Grzegorz Stachowiak. 2016. Deterministic Population Protocols for Exact Majority and Plurality. In *20th International Conference on Principles of Distributed Systems, OPODIS 2016 (LIPIcs, Vol. 70)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 14:1–14:14. https://doi.org/10.4230/LIPIcs.OPODIS.2016.14

[32] Mohsen Ghaffari and Johannes Lengler. 2018. Nearly-Tight Analysis for 2-Choice and 3-Majority Consensus Dynamics. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing, PODC 2018*. ACM, 305–313. https://doi.org/10.1145/3212734.3212738

[33] Mohsen Ghaffari and Merav Parter. 2016. A Polylogarithmic Gossip Algorithm for Plurality Consensus. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing, PODC 2016*. ACM, 117–126. https://doi.org/10.1145/2933057.2933097

[34] Yehuda Hassin and David Peleg. 2001. Distributed Probabilistic Polling and Applications to Proportionate Agreement. *Information and Computation* 171, 2 (2001), 248–268. https://doi.org/10.1006/inco.2001.3088

[35] Varun Kanade, Frederik Mallmann-Trenn, and Thomas Sauerwald. 2019. On coalescence time in graphs: When is coalescing as fast as meeting?: Extended Abstract. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019*. SIAM, 956–965. https://doi.org/10.1137/1.9781611975482.59

[36] Johannes Lengler. 2020. Drift Analysis. In *Theory of Evolutionary Computation - Recent Developments in Discrete Optimization*, Benjamin Doerr and Frank Neumann (Eds.). Springer, 89–131. https://doi.org/10.1007/978-3-030-29414-4_2

[37] Johannes Lengler and Angelika Steger. 2018. Drift analysis and evolutionary algorithms revisited. *Combinatorics, Probability and Computing* 27, 4 (2018), 643–666.

[38] Toshio Nakata, Hiroshi Imahayashi, and Masafumi Yamashita. 2000. A probabilistic local majority polling game on weighted directed graphs with an application to the distributed agreement problem. *Networks* 35, 4 (2000), 266–273. https://doi.org/10.1002/1097-0037(200007)35:4<266::AID-NET5>3.0.CO;2-4

[39] Emanuele Natale and Iliad Ramezani. 2019. On the Necessary Memory to Compute the Plurality in Multi-agent Systems. In *Algorithms and Complexity - 11th International Conference, CIAC 2019 (Lecture Notes in Computer Science, Vol. 11485)*. Springer, 323–338. https://doi.org/10.1007/978-3-030-17402-6_27

[40] Etienne Perron, Dinkar Vasudevan, and Milan Vojnovic. 2009. Using Three States for Binary Consensus on Complete Graphs. In *INFOCOM 2009. 28th IEEE International Conference on Computer Communications*. IEEE, 2527–2535. https://doi.org/10.1109/INFCOM.2009.5062181

[41] Grant Schoenebeck and Fang-Yi Yu. 2018. Consensus of Interacting Particle Systems on Erdös-Rényi Graphs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*. SIAM, 1945–1964. https://doi.org/10.1137/1.9781611975031.127