



An Almost Tight Lower Bound for Plurality Consensus with Undecided State Dynamics in the Population Protocol Model

Antoine El-Hayek
Institute of Science and Technology
Austria
Klosterneuburg, Austria
antoine.el-hayek@ist.ac.at

Robert Elsässer
University of Salzburg
Salzburg, Austria
elsa@cs.plus.ac.at

Stefan Schmid
TU Berlin
Berlin, Germany
stefan.schmid@tu-berlin.de

Abstract

We revisit the majority problem in the population protocol communication model, as first studied by Angluin et al. (Distributed Computing 2008). We consider a more general version of this problem known as plurality consensus, which has already been studied intensively in the literature. In this problem, each node in a system of n nodes, has initially one of k different opinions, and they need to agree on the (relative) majority opinion. In particular, we consider the important and intensively studied model of Undecided State Dynamics.

Our main contribution is an almost tight lower bound on the stabilization time: we prove that there exists an initial configuration, even with bias $\Delta = \omega(\sqrt{n \log n})$, where stabilization requires $\Omega(kn \log \frac{\sqrt{n}}{k \log n})$ interactions, or equivalently, $\Omega(k \log \frac{\sqrt{n}}{k \log n})$ parallel time for any $k = o(\frac{\sqrt{n}}{\log n})$. This bound is tight for any $k \leq n^{\frac{1}{2}-\epsilon}$, where $\epsilon > 0$ can be any small constant, as Amir et al. (PODC'23) gave a $O(k \log n)$ parallel time upper bound for $k = O(\frac{\sqrt{n}}{\log^2 n})$.

CCS Concepts

• Theory of computation → Distributed algorithms.

Keywords

Population protocols, k -majority problem, undecided state dynamics, lower bounds

ACM Reference Format:

Antoine El-Hayek, Robert Elsässer, and Stefan Schmid. 2025. An Almost Tight Lower Bound for Plurality Consensus with Undecided State Dynamics in the Population Protocol Model. In *ACM Symposium on Principles of Distributed Computing (PODC '25)*, June 16–20, 2025, Huatulco, Mexico. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3732772.3733505>

1 Introduction

Population protocols are a simple and natural computational framework, in which n anonymous nodes (also called agents) communicate and interact with each other to solve a predefined problem in a distributed manner. In the underlying communication model,

a scheduler selects in each discrete time step two nodes for interaction. The selected nodes exchange their current states and each of them changes its own state according to the transition function defined by the population protocol. The set of nodes is expected to eventually stabilize in a final configuration that is given by the problem definition. In the plurality consensus problem, as considered in this paper, every node has at the beginning an opinion assigned from a set of k different opinions, and in the final configuration all nodes agree on one of the initial opinions.

In the model introduced by Angluin et al. [7] the population is represented by the nodes of a graph and a scheduler can only choose nodes connected by an edge for interaction. The usual complexity measures in which one is interested are the cardinality of the state space of the transition function and the time needed for the population to stabilize. The time is defined as the number of interactions until a stable configuration is reached. As in previous work, we are also interested in the so-called parallel time, which corresponds to the number of interactions divided by the population size n .

Population protocols have various applications. In their original paper, Angluin et al. [7] motivated population protocols in the context of sensor networks where nodes perform simple computations. Other motivating examples are processes in chemical reaction networks [29]. Population protocols can also be implemented at the level of DNA molecules, as shown in [16]. Furthermore, Cardelli and Csiksz-Nagy [15] considered similarities between biochemical regulatory processes in living cells and population protocols. Population protocols highly influenced the way certain aspects of distributed computing evolved in recent years, for which in 2020 the original paper by Angluin et al. [7] has been awarded the Edsger W. Dijkstra Prize in Distributed Computing.

In the vast majority of the related papers, as in our paper as well, it is assumed that the graph modeling the population is a clique and a random scheduler is in place. That is, in each discrete time step two nodes are selected uniformly at random for interaction. One of the most prominent protocols for the plurality consensus problem is the Undecided State Dynamics, on which we focus in this paper. In its original, unconditional version, this protocol is simple and only uses $k+1$ different states (see the problem definition below for details). This protocol and variants of it have extensively been analyzed in different communication models. Although very recently an upper bound of $O(k \log n)$ on the parallel time has been derived [6], the question of whether this bound is tight is still open.

1.1 Model and Problem

A population protocol and its time performance can be formalized as follows (cf. [12]). Let V denote the set of agents in the population,



This work is licensed under a Creative Commons Attribution International 4.0 License. *PODC '25, Huatulco, Mexico*

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1885-4/25/06

<https://doi.org/10.1145/3732772.3733505>

and let $n = |V|$. Let Σ be the set of states of the protocol, whose cardinality may grow with n . Two interacting nodes change their states according to a deterministic function $f : \Sigma^2 \mapsto \Sigma^2$. That is, $f(q', q'') = (r', r'')$ describes the following transition: if a node in state q' interacts with a node in state q'' , then the first node changes its state to r' and the second node to r'' . A population protocol is also assigned an *output function* $\gamma : \Sigma \rightarrow \Gamma$, which maps every state to an output value. The set Γ may be the same as Σ .

In this paper we analyze plurality consensus (aka the k -majority problem): Each agent i has as an input an initial opinion $s_i \in [k]$, and the goal of the agents is to decide which opinion was the (relative) majority at the start of the computation. The transition function we consider is given by the (unconditional) Undecided State Dynamics: Σ consists of $k+1$ states, the original k opinions, and an extra one, \perp , that represents the state of being undecided. Then $f(s_1, s_2) = (\perp, \perp)$ if $s_1 \neq s_2$, and $s_1, s_2 \in [k]$, $f(s, \perp) = (s, s)$ for any $s \in [k]$. Otherwise, f is just the identity function. In other words, when two agents with different opinions meet, they both become undecided, but when a decided agent meets an undecided one, the latter takes on the opinion of the former. Note that in the case of the Undecided State Dynamics the set $\Gamma = \Sigma$ and the output function γ is the identity. We assume that the interaction graph is a clique, and at each time step, two nodes are selected for interaction, which are chosen uniformly at random (without replacement), independently of the other time steps. We ask how long it takes for the system to stabilize, in the particular case where the majority opinion starts with an initial additional bias of $\Omega(\sqrt{n \log n})$. Note that such a bias is probably needed in order to guarantee that w.h.p. the opinion with the largest (relative) initial support wins (cf. [6, 9]).

1.2 Related Work

Computations over dynamic networks (such as sensor networks) have already been studied intensively in the literature for various models [1, 23, 26]. Probably the most studied problems in the framework of population protocols are leader election and (exact) majority. For two recent surveys, which focus on these problems, we refer the reader to [4] and [20]. As in this paper we analyze plurality consensus, we only present the most relevant results on the related majority problem and adapt the description of [20] to outline the related work on majority in population protocols. In this problem, at the beginning every node is in one of two states (called e.g. A and B). In exact majority, the final opinion has to be the one with the largest initial support, even if at the beginning the difference between the support of the two was just 1. In approximate majority this requirement is weaker: the initial majority should only win with high probability if there is a sufficiently large initial bias between the two opinions (usually this bias is of order $\Omega(\sqrt{n \log n})$).

For the majority problem (as well as leader election), there are a number of results, which present lower bounds on the number of states under certain time requirements, or bound the (stabilization) time under specific assumptions w.r.t. the number of states. Furthermore, several majority and leader election algorithms have been derived, which upper bound the number of states as well as the stabilization time.

Two early papers by Draief and Vojnović [19] and Mertzios et al. [25] consider population protocols for exact majority. They studied (almost) the same four-state protocol, which has a polynomial stabilization time (with high probability¹ and in expectation) on any graph.

In an early paper, Angluin et al. [8] presented population protocols with a constant number of states for several different functions. The protocols they propose are only correct with high probability and they assume that a designated leader is available from the start of the computation, which synchronizes the nodes. Their exact majority protocol has a w.h.p. stabilization time of $O(\log^2 n)$, and it alternates between so-called cancellation and duplication phases, an idea used in many subsequent papers.

Note that any protocol for exact majority which uses a constant number of states (as the four-state protocol described above) is in general slow. However, if the initial imbalance between the support of A and B is large, then the four-state protocol stabilizes fast. In order to increase the initial imbalance, Alistarh et al. [5] multiplied the opinion on each node by some (large) integer. Based on this idea, Alistarh et al. [2] achieved a stabilization time of $O(\log^3 n)$, w.h.p. and in expectation, by utilizing $O(\log^2 n)$ states.

Bilke et al. [14] extended the cancellation-duplication framework from [8] to the leaderless case, provided that the agents have enough states to store the number of interactions they performed so far. The stabilization time of their majority protocol is $O(\log^2 n)$ w.h.p. and in expectation, and it utilizes $O(\log^2 n)$ states.

On the lower bound side, Alistarh et al. [3] proved that any exact majority protocol with expected stabilization time $O(n^{1-\epsilon})$ (ϵ can be any positive constant), which satisfies two natural properties called *monotonicity* and *output dominance*, requires $\Omega(\log n)$ states. They also presented an algorithm with $\Theta(\log n)$ states and $O(\log^2 n)$ stabilization time (w.h.p. and in expectation). Monotonic protocols have the property that their running time does not increase if they are run on a smaller number of agents. Output dominance means that “if the positive counts of states in a stable configuration are changed, then the protocol will stabilize to the same output” (cf. [20]). The (w.h.p. and expected) stabilization time has subsequently been improved to $O(\log^{5/3} n)$ in [12], to $O(\log^{3/2} n)$ in [11], and finally to $O(\log n)$ in [18], by keeping the number of states at $O(\log n)$.

All the results above were derived for *stabilization* time, and the lower bounds do not hold for the so-called *convergence* time. The convergence time is the time required by the protocol to reach a configuration with the correct output property; however, the system may leave this configuration with a small probability. In contrast, if the system stabilizes, then the output of the system does not change anymore.² Kosowski and Uznański [22] and Berenbrink et al. [13] derived algorithms with polylogarithmic *convergence* time, which use $o(\log n)$ states. As outlined in [12], in [22] the authors presented a programming framework that leads to protocols which require only $O(1)$ states and converge in polylogarithmic time (in expectation), but they are only correct w.h.p. These protocols can be changed so that they are always correct by either

¹With high probability or w.h.p. means with probability at least $1 - n^{-\Omega(1)}$, where n is the number of agents.

²See e.g. [13] for details. In the Undecided State Dynamics, convergence and stabilization are equivalent.

allowing $O(\log \log n)$ states, while the convergence time still remains polylogarithmic, or by allowing $O(n^\epsilon)$ convergence time, while keeping the number of states constant. In [13] the authors presented an always correct protocol with a w.h.p. convergence time of $O(\log^2 n / \log s)$ and $\Theta(s + \log \log n)$ states, and an always correct protocol with w.h.p. stabilization time of $O(\log^2 n / \log s)$ and $O(s \cdot \log n / \log s)$ states, where $s \in [2, n]$.

One research direction in plurality consensus focuses on the state complexity (regardless of the time complexity) of protocols, which are required to *always* determine the plurality opinion. While clearly at least k states are required to encode k opinions, [27] shows that always correct plurality consensus needs even $\Omega(k^2)$ states. The protocol of [21] utilizes $O(k^{11})$ states, which can be improved to $O(k^6)$ provided that a total ordering among the opinions exists. Clearly, the lower bound of $\Omega(k^2)$ only holds, if it is required that the correct plurality opinion is determined with probability 1. If such strong guarantees are not required, then the number of states can be much smaller. In [9] a synchronized variant of the Undecided State Dynamics has been presented that reaches consensus w.h.p. in $O(\log^2 n)$ parallel time using $O(k \log n)$ states. However, this protocol solves *approximate* plurality consensus, i.e., if the initial bias is $\Omega(\sqrt{n \log n})$, then w.h.p. the opinion with the initially largest support wins, otherwise a so-called significant opinion wins w.h.p. In the case $k = 2$, the unconditional Undecided State Dynamics has a w.h.p. and expected stabilization time of $O(\log n)$ [17]. Recently, Amir et al. [6] analyzed the unconditional Undecided State Dynamics for the plurality consensus problem and showed that their protocol stabilizes w.h.p. within $O(k \log n)$ parallel time for any initial configuration as long as $k = O(\sqrt{n} / \log^2 n)$. The question of whether this bound on the stabilization time is tight is still open.

We should note that the unconditional Undecided State Dynamics has extensively been analyzed in the Gossip communication model. In this model (which can be seen as a synchronous variant of the population protocol model) in each discrete time step, every node randomly chooses another node for interaction to perform a state transition. Becchetti et al. defined the concept of *monochromatic distance* $\text{md}(c)$ of a configuration c and showed that in this model the time needed to reach a final configuration is $O(\text{md}(c) \log n)$ w.h.p., where c is the initial configuration in the population [10]. They also derived a lower bound that is asymptotically tight up to a $\log n$ factor.

As described by Amir et al. [6], the differences in how nodes are scheduled for interaction in the population protocol model and the Gossip model, respectively, cause the Undecided State Dynamics to “exhibit significant qualitative differences when run in either setting, even in the case when $k = 2$ ”. One of the reasons for these differences is the fact that while in the Gossip model in each step a node may change its opinion only once, and each node is selected for interaction, in the population protocol model a node may change its opinion up to $\Omega(\log n)$ many times in n consecutive interactions (which corresponds to one parallel round) while a constant fraction of nodes is not even selected for interaction. Thus, there are so far no general analysis techniques that allow us to transfer results from one model to the other one (cf. [17]).

1.3 Contribution

Our main contribution is an almost tight lower bound on the stabilization time of Undecided State Dynamics for plurality consensus in population protocols. Specifically, we show that the time needed to stabilize is $\Omega\left(kn \log \frac{\sqrt{n}}{k \log n}\right)$ interactions, or in $\Omega\left(k \log \frac{\sqrt{n}}{k \log n}\right)$ parallel time³, in the case where $k = o\left(\frac{\sqrt{n}}{\log n}\right)$. In our proofs we assume that our initial configuration has bias at most $O(\sqrt{n} / (k \log n))^{1/4} \sqrt{n \log n}$. Interestingly, this includes even an initial bias of $\omega(\sqrt{n \log n})$. This bound is tight for any $k \leq n^{\frac{1}{2}-\epsilon}$, where $\epsilon > 0$ can be any small constant, as Amir, Aspnes, Berenbrink, Biermeier, Hahn, Kaaser, Lazarsfeld [6] gave a $O(k \log n)$ parallel time upper bound. For larger values of k , as any initial configuration that is valid for k_0 is also valid for $k \geq k_0$, we can simply plug in $k_0 = \frac{\sqrt{n}}{\log n \log \log n}$ to get a $\Omega\left(\frac{\sqrt{n} \log \log \log n}{\log n \log \log n}\right)$ parallel time lower bound.

Our analysis is based on a precise characterization of the number of undecided nodes over time, using drift analysis and random walks. Our technical approach and its novelty will be discussed in more details in the next section.

1.4 Notation and Organization

Throughout the paper, \mathbf{x} represents a configuration the system can be in: $\mathbf{x} = (x_1, \dots, x_k, u)$, where x_i is the number of agents with opinion i for all i , and u is the number of undecided agents. We denote by $\mathbf{x}(t)$ the configuration of the system after the t -th interaction, and $\mathbf{x}(0)$ is the initial configuration. $x_i(t)$ is the number of agents with opinion i after interaction t . Similarly, $u(t)$ is the number of undecided nodes after t interactions. Accordingly, $x_i(0)$ is the initial number of nodes with opinion i (clearly, $u(0) = 0$). We assume that the opinions are initially ordered from most common opinion to least common, i.e. $x_1(0) \geq x_2(0) \geq \dots \geq x_k(0)$.

The remainder of this paper is organized as follows. We first provide a technical overview and an empirical motivation in Section 2. Our lower bound analysis is presented in Section 3. We conclude and discuss open research questions in Section 4. Some technical details are deferred to the Appendix.

2 Motivation and Technical Overview

Before we delve into the technical details of our lower bound proof, let us provide some intuition about how agent opinions propagate across the distributed system, and how majority and minority opinions evolve accordingly. Based on this intuition and the observed challenges, we will then give an overview of our analytical approach and its key ideas.

To illustrate the evolution of opinions and the stabilization behavior, Figure 1 (left) plots a typical simulation run over parallel time. In this example, all minority opinions are initially equally frequent, while the majority one has an initial additional bias of $\sqrt{n \log n}$. The majority opinion is plotted in red. The minority opinions are plotted in yellow, except for a random one which we select as an example and which we highlight in orange; for better visibility we multiplied the number of each minority opinion by k .

³Recall that the parallel time is equal to the number of interactions divided by n .

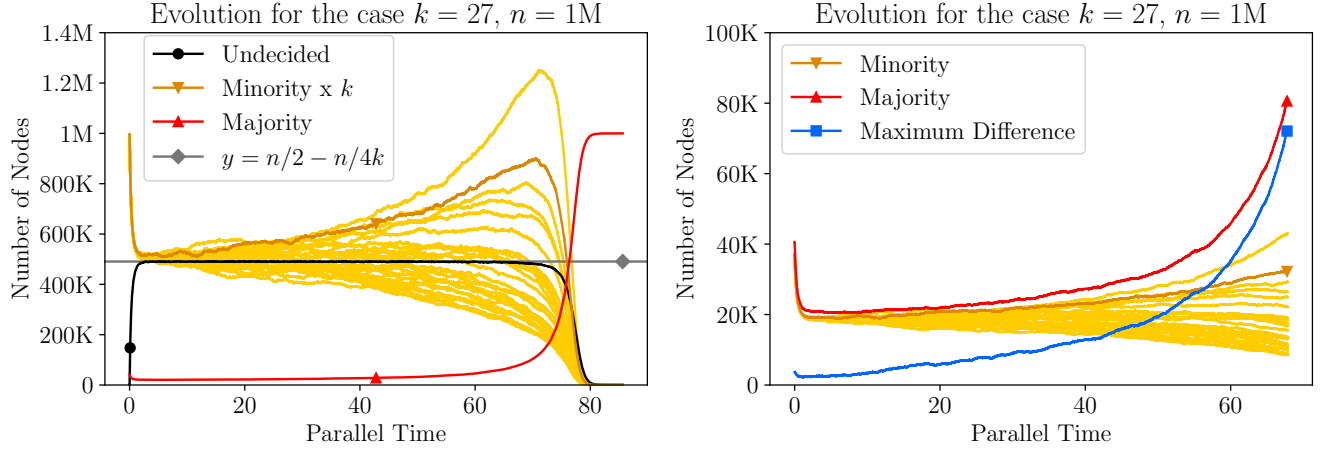


Figure 1: Intuition for stabilization of plurality consensus in undecided state dynamics. On the left figure, we scale up the minority opinions by a k factor for visibility. In this experiment, $n = 1,000,000$ agents and $k = \frac{\sqrt{n}}{\log n \log \log n}$. In the initial configuration $k-1$ opinions had the same support, while opinion 1 had an additive advantage of $\sqrt{n \log n}$. Out of all the minority opinions, we plot one with a darker color for visibility.

We make several observations which highlight parts of the complexity of the problem. First, note that different minority opinions evolve differently. In particular, not all minority opinions are strictly decreasing over time, but many are actually increasing over a long time period. In the example in Figure 1, one opinion even surpasses its initial count. On the other hand, the majority opinion in this example (which is typical for many runs) remains low for a long time during stabilization, but then increases quickly toward the end. We can also see that the number of undecided opinions stays close to $\frac{n}{2} - \frac{n}{4k}$ throughout the execution.

In Figure 1 (right), we zoom in on the time period for which it takes x_1 to double from its initial number of nodes. We can see that it takes most of the stabilization time to reach this configuration, as the system needs around 70 parallel time to do so, after which only 20 more rounds are required to fully stabilize.

These simulations also provide some intuitive motivation for our choices for analyzing the problem theoretically in the next section. We will prove that the number of undecided nodes does not substantially exceed $\frac{n}{2} - \frac{n}{4k}$ with high probability, and we only consider the interactions before x_1 reaches $2x_1(0)$, which we will approximate as $2\frac{n}{k}$. One particular ingredient we consider in our analysis is the *maximum* difference between the majority opinion and minority ones (also plotted in Figure 1 (right)): $\max_{j \geq 2} \{x_1 - x_j\}$. The idea is that as long as this difference stays small, the system is very slow to change.

More concretely, for our analysis, we proceed as follows.

We first observe that a precise characterization of the number of undecided nodes $u(t)$ is key to understand the state dynamics as, intuitively, for any opinion i , the larger the number of undecided nodes, the more new nodes can the i -opinionated nodes “convert” to the opinion i , hoping to compensate for the i -opinionated nodes that meet nodes with other opinions and thus become themselves undecided. In fact, in their analysis of the stabilization time, Amir,

Aspnes, Berenbrink, Biermeier, Hahn, Kaaser and Lazarsfeld [6] derived an upper bound and a lower bound on the number of undecided nodes throughout the interactions: $\frac{n}{2} - \frac{x_1}{2} \leq u(t) \leq \frac{n}{2}$ for any t after the first $n \log n$ interactions. More precisely, the evolution of the number of nodes in opinion i is driven by the number of undecided nodes. For each opinion, there exists a value u_i of the number of undecided nodes that acts as a threshold. That is, if the number of undecided nodes u is above u_i , then the corresponding opinion x_i increases, whereas if u is below the threshold, then x_i decreases. The threshold is a decreasing function in the number of nodes in opinion i : the larger x_i is, the smaller u_i is. One can see this in action in Figure 1 on the left. When $u(t)$ is still very small in the beginning, all of the opinions decrease quickly. Then, once $u(t)$ settles around $\frac{n}{2} - \frac{n}{4k}$, some opinions start to steadily increase (even minority ones) while others decrease. One can understand this the following way: in the initial steps, where the number of undecided nodes increase quickly, some opinions, due to randomness, end up having an advantage on the others. Towards the end, the number of undecided nodes starts dropping, it thus goes below all thresholds but the one of the majority opinion, which means all opinions drop quickly apart from the majority one.

In our case, for the study of the lower bound, we fully control the initial configuration of the opinions, and we are able to give a better upper bound on $u(t)$ over time, which happens to be very close to the threshold of any opinion i . To do so, we use drift analysis (for a nice introduction, we refer to Lengler [24]). More specifically, we use a result by Oliveto and Witt [28]. Here is a quick intuition on drift analysis: assuming we know the current configuration of the system, we compute the expectation of the change in the number of agents in each state after the next interaction. The idea is, if a number changes in expectation by a value α at each interaction, and our goal is to show that the actual number does not wander off by more than β from its original value, this takes at least $\Omega(\beta/\alpha)$

many interactions. A few more hypotheses are needed to ensure that this result holds with high probability, and to ensure that the probabilities are well-behaved, but this is the main idea.

It turns out that $u(t)$ settles around $\frac{n}{2} - \frac{n}{4k}$, and to show that $u(t)$ never substantially exceeds this value, we prove that if it slightly exceeds this value, then at each interaction, in expectation, $u(t)$ decreases by at least $\sqrt{\log n/n}$. Drift analysis tools then allow us to ensure that $u(t)$ will drift no more than (roughly) $\theta(\sqrt{n \log n})$ away from that position.

The second step we make after giving an upper bound on $u(t)$, is to show that in $\theta(kn)$ interactions, no opinion can go from $\frac{3n}{2k}$ nodes to $\frac{2n}{k}$ agents. Here, classic drift analysis fails to help, and that for one main reason: drift analysis essentially looks at how much the expectation increases or decreases over time, but sometimes, the expectation gets beaten by the variance of the process, and drift analysis fails to capture the lack of concentration of the process. To understand that, consider a random walk starting at 0 and which at each step either increases or decreases by 1, each with probability $1/2$. The expectation of its increase is 0 so in expectation, after m steps, it is still at 0. However, we know that with high probability it will have reached $\theta(\sqrt{m})$ at some point during those m steps, as one can see the position after m steps as the result of a binomial distribution with parameters $(1/2, m)$, whose standard deviation is of the order of \sqrt{m} .

In our case, we avoid this problem by looking more carefully at the probabilities involved. By considering the evolution of x_i , while its expectation increases slowly over time, it is not overtaken by the variance. To see that, think again of our example of our ± 1 random walk: Imagine that now, the random walk increases with probability $p/2$, decreases with probability $p/2$, and stays put with probability $1-p$. Look now at what happens after m steps. The walk actually moved for pm out of those steps, and thus the standard deviation is now around \sqrt{pm} . If $p = o(1)$, this has a significant impact.

This is exactly what happens in our case, where for x_i to change, in an interaction, a node with opinion i must have been chosen. If at most $2\frac{n}{k}$ nodes with this opinion exist, then the probability of this event happening is of the order of $1/k = o(1)$.

While we could have chosen to stop here for our analysis, and thus give a $\Omega(k)$ lower bound, we go one step further: there is a weakness in this analysis, i.e., we overestimate the number of nodes in opinion i , by stating that the initial count is at most $\frac{3n}{2k}$, while in practice, most opinions stay close to $\frac{n}{2k}$ for most of the process. This (high) estimate of $2\frac{n}{k}$ then gives a too low threshold on u for x_i to increase, and thus the upper bound on the rate at which x_i increases is not tight enough. To improve the result, we would need to give a better estimate on the initial x_i , and show that it is close to $\frac{n}{2k}$. This is not straightforward, as the initial count is close to $\frac{n}{k}$, and we would thus need to analyze the initial phase where u increases sharply while all the x_i decrease.

We find a workaround for this difficulty: while it might take effort to accurately approximate x_i after the first quick-changing phase, a value that is easier to estimate is $\Delta_{ij} = x_i - x_j$. Here, instead of requiring a good estimate on x_i and x_j to analyze the evolution of Δ_{ij} , it turns out that we only need their order of magnitude, as well as a good estimate of Δ_{ij} , which we have. Therefore, knowing that

$x_i \leq 2\frac{n}{k}$ for $\theta(kn)$ iterations allows us to show that the maximum Δ_{ij} needs at least $\theta(kn)$ iterations to double. However, if all Δ_{ij} are small enough (if they are all $o(\frac{n}{k})$), then it means that no opinion drifts far from the others, and thus $x_i \leq \frac{3n}{2k}$ after those $\theta(kn)$ interactions.

Hence, our high level argument works as follows: First give a good estimate on u , then on the number of nodes in each opinion, and on the maximum difference between opinions: As long as $x_i \leq \frac{3n}{2k}$ at a time t and $\Delta_{ij} = o(\frac{n}{k})$ for all i, j , then the order of x_i does not change in the next $\theta(kn)$ interactions, which in turn is used to show that Δ_{ij} does not double during those same interactions. This in turn shows that $x_i \leq \frac{3n}{2k}$, which allows us to restart another iteration of the induction. The induction holds for $\log\left(\frac{\sqrt{n}}{k \log n}\right)$ iterations, which then gives the lower bound.

3 Lower Bound

In this section, we will give a lower bound on the number of rounds needed for the protocol to stabilize. For that, we analyze the situation where all minority opinions start with the same number of nodes, that is, $x_i(0) = x_j(0)$ for any $i, j \in [2, k]$, while the majority opinion starts with an initial gap of $x_1(0) - x_2(0) = O\left((\sqrt{n}/(k \log n))^{1/4} \sqrt{n \log n}\right)$. We also require that $k = o\left(\frac{\sqrt{n}}{\log n}\right)$. Since a $\Omega(\log n)$ lower bound for the problem is trivial (in $o(\log n)$ parallel time, w.h.p. there are nodes that have not interacted at all) which implies a $\Omega(k \log n)$ lower bound for $k = O(1)$, we focus on the case $k = \omega(1)$.

We first begin by giving an upper bound on the number of undecided nodes.

Lemma 3.1. *For any $\tau \leq n^4$, and any initial configuration, it holds with probability at least $1 - n^{-4}$ that $u(\tau) \leq \frac{n}{2} - \frac{n}{4k} + \frac{10n}{(k-1)^2} + (20 \cdot 132 + 1)\sqrt{n \log n}$.*

PROOF. To prove the lemma, we will use Theorem A.1. We model $u(t)$ to be a random walk over integers, and first compute the expectation of $u(t+1)$ conditioned on the configuration of the system at time t , which is denoted by $\mathbf{x} = (x_1, \dots, x_k, u)$ (for ease of notation, we write x_i instead of $x_i(t)$ and u instead of $u(t)$). Clearly, the number of undecided nodes can either decrease by 1, if a decided node interacts with an undecided node, or it can increase by two, if two nodes of different opinions interact. The probability to decrease by one is $2\frac{u}{n} \cdot \frac{n-u}{n-1}$, while the probability to increase by two is $\frac{\sum_{i \in [k]} x_i}{n} \cdot \frac{\sum_{j \neq i} x_j}{n-1}$. Note that $\frac{1}{n-1} = \frac{1}{n} + O(\frac{1}{n^2})$. We therefore have:

$$\begin{aligned} \mathbb{E}[u(t+1) | \mathbf{x}(t) = \mathbf{x}] &= u - 2\frac{u}{n} \cdot \frac{n-u}{n} + 2\frac{\sum_{i \in [k]} x_i}{n} \cdot \frac{\sum_{j \neq i} x_j}{n} + O\left(\frac{1}{n}\right) \\ &= u - 2\frac{u}{n} + 2\frac{u^2}{n^2} + 2\frac{\sum_{i \in [k]} x_i(n-u-x_i)}{n^2} + O\left(\frac{1}{n}\right) \end{aligned}$$

For the second equality we used $\sum_{j \neq i} x_j = n - u - x_i$. Splitting $x_i(n - u - x_i)$ into $x_i(n - u)$ and $-x_i^2$, and writing $\sum_{i \in [k]} x_i = n - u$ we obtain

$$\begin{aligned} &= u - 2 \frac{u}{n} + 2 \frac{u^2}{n^2} + 2 \frac{n^2 - 2un + u^2}{n^2} - 2 \frac{\sum_{i \in [k]} x_i^2}{n^2} + O\left(\frac{1}{n}\right) \\ &= u - 6 \frac{u}{n} + 4 \frac{u^2}{n^2} + 2 - 2 \frac{\sum_{i \in [k]} x_i^2}{n^2} + O\left(\frac{1}{n}\right) \end{aligned}$$

Clearly, if we fix $n - u$, the sum $\sum_{i \in [k]} x_i^2$ is minimized if all x_i are equal. Then we have

$$\begin{aligned} &\mathbb{E}[u(t+1) | \mathbf{x}(t) = \mathbf{x}] \\ &\leq u - 6 \frac{u}{n} + 4 \frac{u^2}{n^2} + 2 - 2 \frac{(n-u)^2}{kn^2} + O\left(\frac{1}{n}\right) \\ &= u - 6 \frac{u}{n} + 4 \frac{u^2}{n^2} + 2 - \frac{2}{k} + \frac{4u}{nk} - \frac{2u^2}{n^2 k} + O\left(\frac{1}{n}\right) \end{aligned}$$

Let us assume that $u = \frac{n}{2} - \frac{n}{4k} + \frac{10n}{(k-1)^2} + c\sqrt{n \log n}$ for some constant c . Then,

$$\begin{aligned} &\mathbb{E}[u(t+1) | \mathbf{x}(t) = \mathbf{x}] \\ &\leq u - 6 \frac{\frac{n}{2} - \frac{n}{4k} + \frac{10n}{(k-1)^2} + c\sqrt{n \log n}}{n} \\ &\quad + 4 \frac{(\frac{n}{2} - \frac{n}{4k} + \frac{10n}{(k-1)^2} + c\sqrt{n \log n})^2}{n^2} + 2 - \frac{2}{k} \\ &\quad + \frac{4(\frac{n}{2} - \frac{n}{4k} + \frac{10n}{(k-1)^2} + c\sqrt{n \log n})}{nk} \\ &\quad - \frac{2(\frac{n}{2} - \frac{n}{4k} + \frac{10n}{(k-1)^2} + c\sqrt{n \log n})^2}{n^2 k} + O\left(\frac{1}{n}\right) \\ &\leq u - 3 + \frac{3}{2k} - \frac{60}{(k-1)^2} - 6c\sqrt{\frac{\log n}{n}} \\ &\quad + 1 + \frac{1}{4k^2} - \frac{1}{k} + \frac{40}{(k-1)^2} \\ &\quad + 4c\sqrt{\frac{\log n}{n}} + 2 - \frac{2}{k} + \frac{2}{k} - \frac{1}{k^2} - \frac{1}{2k} + \frac{1}{2k^2} \\ &\quad + O\left(\frac{1}{k^3}\right) + o\left(\sqrt{\frac{\log n}{n}}\right) \\ &\leq u - c\sqrt{\frac{\log n}{n}} \end{aligned}$$

for any $c \geq 1$ if k is large enough. Define $\tilde{u} = \frac{n}{2} - \frac{n}{4k} + \frac{10n}{(k-1)^2}$.

We can now apply Theorem A.1 with $X_t = -u(t)$, $X_0 = 0$, $a = -\tilde{u} - \sqrt{n \log n} - 20 \cdot 132\sqrt{n \log n}$, $b = -\tilde{u} - \sqrt{n \log n}$, $\ell = 20 \cdot 132\sqrt{n \log n}$, $\epsilon = \sqrt{\frac{\log n}{n}}$ and $r = \sqrt{5}$. Then, we obtain that T^* , the first time so that $u(T^*) \geq \frac{n}{2} - \frac{n}{4k} + \frac{10n}{(k-1)^2} + \sqrt{n \log n} + 20 \cdot 132\sqrt{n \log n}$, satisfies:

$$\mathbb{P}[T^* \leq \exp(4 \log n)] \leq O(\exp(-4 \log n))$$

This implies that $\mathbb{P}[T^* \leq n^4] \leq O(n^{-4})$. Therefore, with high probability $u(t)$ is less than $\frac{n}{2} - \frac{n}{4k} + \frac{10n}{(k-1)^2} + \sqrt{n \log n} + 20 \cdot 132\sqrt{n \log n}$ for all $0 \leq t \leq n^4$. \square

We will need the following lemma to continue our analysis. It is a simplified version of Theorem 20 in [24], and provides the tool we need for Lemmas 3.3 and 3.4.

Lemma 3.2. *Let $Y(t)$, $t \geq 0$, be a random walk defined over the set of integers as follows: $Y(0) = 0$, $Y(t+1) = Y(t)$ with probability $1-p(t)$, $Y(t+1) = Y(t)+1$ with probability $\frac{p(t)+q(t)}{2}$, and $Y(t+1) = Y(t)-1$ with probability $\frac{p(t)-q(t)}{2}$. Assume there are values $p > 0$ and $q > 0$ such that $0 \leq p(t) \leq p$ and $-p(t) \leq q(t) \leq q$. For any T such that $T \geq 32 \left(\frac{p-q^2}{2q} + \frac{2}{3} \right) \log n$, with probability at least $1 - n^{-2}$, $Y(t) < T$ for every $t < \min\{\frac{T}{2q}, n^2\}$ steps.*

PROOF. To prove this lemma, we will need Bernstein's inequality (cf. Theorem A.2). We first introduce the random variables $\tilde{Y}(t)$, coupled to $Y(t)$ as follows: $\tilde{Y}(t)$ is a random walk such that with probability $1-p(t)$, $\tilde{Y}(t+1) = \tilde{Y}(t)$, with probability $\frac{p(t)+q}{2}$, $\tilde{Y}(t+1) = \tilde{Y}(t)+1$, and with probability $\frac{p(t)-q}{2}$, $\tilde{Y}(t+1) = \tilde{Y}(t)-1$. Also, with probability 1, $Y(t+1) = Y(t)$ iff $\tilde{Y}(t+1) = \tilde{Y}(t)$ and if $Y(t+1) \geq Y(t)$, then $\tilde{Y}(t+1) \geq \tilde{Y}(t)$.

This can easily be done by sampling a random number $r(t) \in (0, 1)$ uniformly at random for every t . If $r(t) \leq 1-p(t)$, we set $Y(t+1) = Y(t)$ and $\tilde{Y}(t+1) = \tilde{Y}(t)$. If $1-p(t) \leq r(t) \leq 1-p(t) + \frac{p(t)+q(t)}{2}$, we set $Y(t+1) = Y(t)+1$ and $\tilde{Y}(t+1) = \tilde{Y}(t)+1$. If $1-p(t) + \frac{p(t)+q(t)}{2} \leq r(t) \leq 1-p(t) + \frac{p(t)+q}{2}$, we set $Y(t+1) = Y(t)-1$ and $\tilde{Y}(t+1) = \tilde{Y}(t)+1$. Else, we set $Y(t+1) = Y(t)-1$ and $\tilde{Y}(t+1) = \tilde{Y}(t)-1$.

Essentially, if $Y(t)$ increases, so does $\tilde{Y}(t)$. If $Y(t)$ stays put, so does $\tilde{Y}(t)$. If $Y(t)$ decreases, $\tilde{Y}(t)$ might either decrease or increase. This coupling ensures that $\tilde{Y}(t) \geq Y(t)$ for all t . To prove the theorem, it thus suffices to prove the theorem for $\tilde{Y}(t)$.

We will now use Bernstein's inequality with $X_i = \tilde{Y}(i+1) - \tilde{Y}(i) - q$. We use $M = 2$ and $\mathbb{E}[X_i^2] = (1-p(t))q^2 + \left(\frac{p(t)+q}{2}\right)(1-q)^2 + \left(\frac{p(t)-q}{2}\right)(-1-q)^2 = p(t) - q^2 \leq p - q^2$. Hence for any N , $\sum_{i \in [N]} \mathbb{E}[X_i^2] \leq N(p - q^2)$. With $N \leq \frac{T}{2q}$:

$$\begin{aligned} \mathbb{P}(\tilde{Y}(N) \geq T) &= \mathbb{P}\left(\sum_{i=1}^N X_i \geq T - qN\right) \leq \mathbb{P}\left(\sum_{i=1}^N X_i \geq \frac{T}{2}\right) \\ &\leq \exp\left(-\frac{\frac{1}{8}T^2}{\sum_{i \in [N]} \mathbb{E}[X_i^2] + \frac{2T}{3}}\right) \\ &\leq \exp\left(-\frac{\frac{1}{8}T^2}{N(p - q^2) + \frac{2T}{3}}\right) \leq \exp\left(-\frac{\frac{1}{8}T}{\frac{p-q^2}{2q} + \frac{2}{3}}\right) \\ &\leq n^{-4} \end{aligned}$$

Using a union bound over the first $\min\{\frac{T}{2q}, n^2\}$ steps, we get that the probability that T is reached within the first $\min\{\frac{T}{2q}, n^2\}$ steps is at most n^{-2} . \square

We now prove that w.h.p. any given opinion can not increase too much in $\theta(kn)$ interactions, if it starts with fewer than $\frac{3n}{2k}$ nodes.

Lemma 3.3. *Let $i \in \{1, \dots, k\}$ be an arbitrary but fixed opinion. Let $\tau_{3n/2k} \leq n^2$ be a time where $x_i(\tau_{3n/2k}) \leq \frac{3n}{2k}$ and $\tau_{2n/k}$ the*

random variable denoting the time at which $x_i(\tau_{2n/k}) = \frac{2n}{k}$. Then, with probability at least $1 - O(n^{-2})$ we have $\tau = \tau_{2n/k} - \tau_{3n/2k} \geq \frac{kn}{25}$.

PROOF. Let us consider the evolution of $x_i(t)$ for some $t \in \{t_0, \dots, t_0 + kn/25\}$, where $t_0 = \tau_{3n/2k}$. For the analysis, we create a new random process y similar to x as follows: $y(0) = x(0)$, and for every t , $y(t+1) = x(t+1)$ if $y(t) = x(t)$ and $u(t+1) \leq \tilde{u} + (20 \cdot 132 + 1)\sqrt{n \log n}$. If $u(t+1)$ does not satisfy the condition, we halt y (we do not define $y(t+1)$) and say that y fails at time $t+1$. By Lemma 3.1, y fails with probability $O(n^{-4})$ in the first n^4 rounds. The modification described above enforces that in y the number of undecided nodes never surpasses $\frac{n}{2} - \frac{n}{4k} + \frac{10n}{(k-1)^2} + (20 \cdot 132 + 1)\sqrt{n \log n}$. Clearly, if during the execution of x the number of undecided nodes does not reach $\frac{n}{2} - \frac{n}{4k} + \frac{10n}{(k-1)^2} + \sqrt{n \log n} + 20 \cdot 132\sqrt{n \log n}$, which happens with probability at least $1 - n^{-4}$ in the first n^4 interactions, then x and y behave identically.

Let us now consider the evolution of $x_i(t)$ in y . $x_i(t)$ increases by one if a node of opinion i meets with an undecided node, which happens with probability

$$\mathbb{P}(+1) := \mathbb{P}(x_i(t+1) - x_i(t) = 1 | y(t) = y) = 2 \frac{x_i}{n} \frac{u}{n-1}$$

Similarly, the probability that it decreases by one is

$$\mathbb{P}(-1) := \mathbb{P}(x_i(t+1) - x_i(t) = -1 | y(t) = y) = 2 \frac{x_i}{n} \frac{n-u-x_i}{n-1}$$

The goal is to use Lemma 3.2. We thus need to compute $p(t) = \mathbb{P}(+1) + \mathbb{P}(-1)$ and $q(t) = \mathbb{P}(+1) - \mathbb{P}(-1)$.

As long as $x_i \leq 2\frac{n}{k}$, we have that $\mathbb{P}(+1) + \mathbb{P}(-1) = 2 \frac{x_i}{n} \frac{n-x_i}{n-1} = 2 \frac{x_i}{n} (1 + o(1)) \leq \frac{5}{k}$. The difference between the two is

$$\begin{aligned} \mathbb{P}(+1) - \mathbb{P}(-1) &= 2 \frac{x_i}{n-1} \left(\frac{u}{n} - \frac{n-u-x_i}{n} \right) = 2 \frac{x_i}{n(n-1)} (2u - n + x_i) \\ &\leq 2 \frac{x_i}{n^2} \left(n - \frac{n}{2k} + O\left(\frac{n}{k \log n}\right) - n + \frac{2n}{k} \right) \\ &= \frac{x_i}{n^2} (3 + O(\frac{1}{\log n})) \frac{n}{k} = 3 \frac{x_i}{n} \frac{1+o(1)}{k} \leq \frac{6}{k^2} (1+o(1)) \end{aligned}$$

We apply Lemma 3.2, with $p = \frac{5}{k}$, $q = \frac{6.25}{k^2}$, $T = \frac{n}{2k}$. It now suffices to check the conditions of Lemma 3.2.

$$\frac{p-q^2}{2q} \leq \frac{\frac{5}{k}}{2 \frac{6.25}{k^2}} = O(k) = o(k \log n)$$

and

$$T = \frac{n}{2k} \geq \frac{k^2 \log^2 n}{2k} = \omega(k \log^2 n)$$

Therefore, we have that $T \geq 8(\frac{p-q^2}{2q} + \frac{2}{3}) \log n$. Thus, $x_i(t_0 + \tau) < \frac{2n}{k}$ for any $\tau \leq \frac{kn}{25}$ with probability at least $1 - n^{-2}$ in y . As x and y behave identically with probability $1 - n^{-4}$ in the first n^4 interactions, the lemma follows. \square

Now that we have a good upper bound on u and a good estimate on the order of magnitude of all of the x_i , we can show that it takes $\theta(kn)$ interactions for the difference between two opinions to double.

Lemma 3.4. Assume that at some time step $t_0 \leq nk \log^2 n$ the difference between any two opinions is at most $\frac{\alpha}{2} = \omega(\sqrt{n \log n})$, where $\alpha = o(\frac{n}{k})$. Furthermore, assume that $x_i(t_0) \leq \frac{3n}{2k}$ for every $i \in [k]$. Then, with probability at least $1 - O(k^2/n^2)$, the difference between any two opinions does not exceed α before interaction $t_0 + \tau$, where $\tau = \frac{1}{24}kn$.

PROOF. Consider two arbitrary but fixed opinions i and j . Similarly to the proof of Lemma 3.3 we analyze the evolution of $\Delta_{ij}(t) = x_i(t) - x_j(t)$ for a time step $t \in \{t_0, \dots, t_0 + \tau\}$. For the analysis, we create a new random process y similar to x as follows: $y(0) = x(0)$, and for every t , $y(t+1) = x(t+1)$ if $y(t) = x(t)$ and $u(t+1) \leq \tilde{u} + (20 \cdot 132 + 1)\sqrt{n \log n}$ and $x_i(t+1) \leq \frac{2n}{k}$ for every i . If either $u(t+1)$ or one of the x_i does not satisfy the condition, we halt y (we do not define $y(t+1)$) and say that y fails at time $t+1$.

Clearly, x and y behave identically with probability $1 - O(kn^{-2})$ according to Lemmas 3.1 and 3.3 (by union bound over the k different opinions).

We consider now the evolution of $\Delta_{ij}(t)$ in y . For ease of notation, we set $x_i = x_i(t)$, $x_j = x_j(t)$, and $u = u(t)$. $\Delta_{ij}(t)$ increases by one if x_i increases by one, which happens when a node with opinion i interacts with an undecided node, hence with probability $2 \frac{x_i}{n} \frac{u}{n-1}$. It also increases by 1 if x_j decreases by one, but not x_i , which happens when a node with opinion j interacts with a node of opinion in $[k] \setminus \{i, j\}$, hence with probability $2 \frac{x_j}{n} \frac{n-u-x_i-x_j}{n-1}$. Therefore, the probability that $\Delta_{ij}(t+1) > \Delta_{ij}(t)$ is:

$$\begin{aligned} \mathbb{P}(+1) &:= \mathbb{P}(\Delta_{ij}(t+1) - \Delta_{ij}(t) = 1 | y(t) = y) \\ &= 2 \frac{x_i}{n} \frac{u}{n-1} + 2 \frac{x_j}{n} \frac{n-u-x_i-x_j}{n-1} \end{aligned}$$

Similarly, the probability that it decreases by one is

$$\begin{aligned} \mathbb{P}(-1) &:= \mathbb{P}(\Delta_{ij}(t+1) - \Delta_{ij}(t) = -1 | y(t) = y) \\ &= 2 \frac{x_j}{n} \frac{u}{n-1} + 2 \frac{x_i}{n} \frac{n-u-x_i-x_j}{n-1} \end{aligned}$$

The goal is to use Lemma 3.2. We thus need to compute $p(t) = \mathbb{P}(+1) + \mathbb{P}(-1)$ and $q(t) = \mathbb{P}(+1) - \mathbb{P}(-1)$. Expanding the terms, we get that the sum of $\mathbb{P}(+1)$ and $\mathbb{P}(-1)$ is:

$$\begin{aligned} \mathbb{P}(+1) + \mathbb{P}(-1) &= 2 \frac{(x_i + x_j)u}{n(n-1)} + 2 \frac{(x_i + x_j)}{n(n-1)} (n-u-x_i-x_j) \\ &\leq 2 \frac{4 \frac{n}{k}}{n(n-1)} (n-x_i-x_j) = 8 \frac{1}{k} (1+o(1)) \end{aligned}$$

The difference between the two is:

$$\begin{aligned}
& \mathbb{P}(+1) - \mathbb{P}(-1) \\
&= \frac{2}{n(n-1)} ((ux_i + nx_j - ux_j - x_i x_j - x_j^2) \\
&\quad + (-ux_j - nx_i + ux_i + x_i^2 + x_i x_j)) \\
&= \frac{2}{n(n-1)} (2u(x_i - x_j) - n(x_i - x_j) \\
&\quad + (x_i - x_j) + (x_i - x_j)(x_i + x_j)) \\
&= \frac{2}{n(n-1)} (x_i - x_j)(2u - n + x_i + x_j) \\
&\leq \frac{2}{n(n-1)} (x_i - x_j)(n - \frac{n}{2k} + o(\frac{n}{k}) - n + 2\frac{n}{k}) \\
&= \frac{3}{n-1} \Delta_{ij} (\frac{1}{k} + o(\frac{1}{k}))
\end{aligned}$$

If the difference at a time t is $x_i(t) - x_j(t) \leq \alpha$, then $\mathbb{P}(+1) - \mathbb{P}(-1) \leq \frac{3\alpha}{nk} (1 + o(1))$.

We apply Lemma 3.2, with $p = \frac{9}{k}$, $q = \frac{6\alpha}{nk}$, and $T = \frac{\alpha}{2}$, $Y(t) = \Delta_{ij}(t) - \frac{\alpha}{2}$. Indeed, we can see the problem as a random walk that starts at $Y(0) = 0$ (which corresponds to $\Delta_{ij} = \frac{\alpha}{2}$) and whose target is $Y(\tau) = \frac{\alpha}{2}$ (which corresponds to $\Delta_{ij} = \alpha$). As long as the target is not reached, we have that $\Delta_{ij} \leq \alpha$, and thus, $q(t) \leq \frac{6\alpha}{nk}$. It now suffices to check the conditions of Lemma 3.2.

We have $T = \frac{\alpha}{2} = \omega(\sqrt{n \log n})$, $\frac{2}{3} \log n = o(\sqrt{n \log n})$ and $\frac{p-q^2}{q} \log n \leq \frac{p}{q} \log n = O(\frac{n}{\alpha} \log n) = o(\frac{n}{\sqrt{n \log n}} \log n) = o(\sqrt{n \log n})$, and thus $32 \log n \left(\frac{p-q^2}{2q} + \frac{2}{3} \right) = o(\sqrt{n \log n}) \leq T = \omega(\sqrt{n \log n})$. We thus know that with probability at least $1 - n^{-2}$ the difference will not exceed α before $\frac{T}{2q} = \frac{1}{24} kn$.

Taking a union bound over all $k(k-1)/2$ pairs of opinions, we know that the probability that one difference exceeds α at time $t_0 + \tau$ is less than k^2/n^{-2} . Since \mathbf{x} and \mathbf{y} as described at the beginning of this proof behave identically with probability $1 - O(kn^{-2})$, we obtain the lemma. The lemma also implies that with probability $1 - O(k^2/n^2)$ the support of each opinion is less than $3n/2k$ at time $t_0 + \tau$. \square

We are now ready to prove the main theorem.

THEOREM 3.5. *For any $k = \omega(1)$, $k = o(\frac{\sqrt{n}}{\log n})$, any initial configurations where the maximum difference between two opinions is $\max_{i,j \in [k]} \{x_i(0) - x_j(0)\} = O\left((\sqrt{n}/(k \log n))^{1/4} \sqrt{n \log n}\right)$ under the Undecided State Dynamics for Plurality Consensus does not stabilize in $\frac{k}{25} \log \frac{\sqrt{n}}{k \log n}$ parallel time, with high probability.*

PROOF. Define $f: f(n) = \left(\frac{\sqrt{n}}{k \log n}\right)^{\frac{1}{4}}$.

We bunch together the interactions by groups of $\tau = \frac{kn}{25}$ interactions, and use induction on ℓ : For $\ell \leq \log \left(\frac{\frac{n^{\frac{3}{4}}}{k^{\frac{1}{2}}}}{\sqrt{n \log n f(n)}} \right)$, after the ℓ -th group of interactions, we have that $x_i(\ell\tau) \leq \frac{3n}{2k}$ for every i , and

that $\Delta(\ell\tau) := \max_{ij} \Delta_{ij}(\ell\tau) \leq 2^\ell \beta$ with probability $1 - O(\ell k^2/n^2)$, where the initial bias starts at $\Delta(0) = \beta = O\left(f(n) \sqrt{n \log n}\right)$.

The base step is trivial. To show the induction step, we first note

that for $\ell \leq \log \left(\frac{\frac{n^{\frac{3}{4}}}{k^{\frac{1}{2}}}}{\sqrt{n \log n f(n)}} \right)$, we have $2^\ell \beta \leq \frac{n^{\frac{3}{4}}}{k^{\frac{1}{2}}}$, which in turn

means that, using $k \leq \frac{n^{\frac{1}{2}}}{\log n}$, $2^\ell \beta \leq \frac{n^{\frac{3}{4}}}{k^{\frac{1}{2}}} \cdot \left(\frac{n^{\frac{1}{2}}}{k \log n} \right)^{\frac{1}{2}} = o(\frac{n}{k})$. We then

apply Lemmas 3.3 and 3.4: First, since for any i , $x_i(\ell\tau) \leq \frac{3n}{2k}$, and by our result on u (Lemma 3.1), we apply Lemma 3.3 to have that during the whole interval, with probability at least $1 - O(n^{-2})$, x_i does not exceed $2\frac{n}{k}$. By union bound over all values of i , it holds that for all i , no x_i exceeds $2\frac{n}{k}$ during the $\frac{kn}{25}$ interactions with probability at least $1 - O(kn^{-2})$. Then, since $\Delta(\ell\tau) \leq 2^\ell \beta$, $x_i(t) \leq 2\frac{n}{k}$ for any $t \leq (\ell+1)\tau$, and by our result on u , we can apply Lemma 3.4. This ensures that $\Delta((\ell+1)\tau) \leq 2^{\ell+1}\beta$ with probability at least

$1 - O((\ell+1)k^2/n^2)$. For $\ell+1 \leq \log \left(\frac{\frac{n^{\frac{3}{4}}}{k^{\frac{1}{2}}}}{\sqrt{n \log n f(n)}} \right)$, this ensures

$\Delta((\ell+1)\tau) \leq \frac{n^{\frac{3}{4}}}{k^{\frac{1}{2}}}$, which in turn means that, using $k \leq \frac{n^{\frac{1}{2}}}{\log n}$,

$\Delta((\ell+1)\tau) \leq \frac{n^{\frac{3}{4}}}{k^{\frac{1}{2}}} \cdot \left(\frac{n^{\frac{1}{2}}}{k \log n} \right)^{\frac{1}{2}} = o(\frac{n}{k})$. Thus $x_i((\ell+1)\tau) \leq \frac{3n}{2k}$ for every $i \in [k]$, as otherwise $\sum_{j \in [k]} x_j((\ell+1)\tau) \geq \frac{3}{2} \frac{n}{k} + \sum_{j \neq i} x_j((\ell+1)\tau) \geq \frac{3}{2} \frac{n}{k} + (k-1) \frac{n}{k} > n$. This proves the double induction.


Hence, with high probability, the system does not stabilize

before $\frac{kn}{25} \log \left(\frac{\frac{n^{\frac{3}{4}}}{k^{\frac{1}{2}}}}{\sqrt{n \log n f(n)}} \right) = \frac{kn}{25} \left(\frac{1}{2} \log \left(\frac{\sqrt{n}}{k \log n} \right) - \log f(n) \right) = \theta \left(kn \log \frac{\sqrt{n}}{k \log n} \right)$ interactions, concluding the proof. \square

4 Conclusion

We presented an almost tight lower bound on the stabilization time of the Undecided State Dynamics for plurality consensus in the population protocol model. While our result settles the question about the stabilization time, there are several interesting avenues for future research. In particular, it would be interesting to explore scenarios where (slightly) more memory is available at the nodes and where synchronization is possible to some extent: at which point can we break the lower bound barrier? Another open question concerns the required initial bias. While it is known from previous work that with an initial bias in the order of $O(\sqrt{n})$, the system can stabilize to a minority opinion with non-negligible probability [17], we assumed a slightly higher initial bias of $\Omega(\sqrt{n \log n})$ which ensures stabilization to the majority opinion (cf. [6]); it remains to close this small gap.

Acknowledgments

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (MoDynStruct, No. 101019564)  and the Austrian Science Fund (FWF) grant DOI 10.55776/I5862,

grant DOI 10.55776/I5982, and grant DOI 10.55776/P33775 with additional funding from the netidee SCIENCE Stiftung, 2020–2024 and the German Research Foundation (DFG), grant 470029389 (FlexNets).

A Useful Theorems

THEOREM A.1 (THEOREM 2 OF [28]). *Let $X_t, t \geq 0$, be real-valued random variables describing a stochastic process over some state space. Suppose there exist an interval $[a, b] \subseteq \mathcal{R}$ and, possibly depending on $\ell := b - a$, a drift bound $\epsilon := \epsilon(\ell) > 0$ as well as a scaling factor $r := r(\ell)$ such that for all $t \geq 0$ the following three conditions hold:*

- $\mathbb{E}[X_{t+1} - X_t | X_0, \dots, X_t; a < X_t < b] \geq \epsilon.$
- $\mathbb{P}[|X_{t+1} - X_t| \geq jr | X_0, \dots, X_t] \leq e^{-j}$ for $j \in \mathbb{N}_0.$
- $1 \leq r^2 \leq \frac{\epsilon\ell}{132 \log \frac{1}{\epsilon}}.$

Then for the first hitting time $T^ := \min\{t \geq 0 : X_t \leq a | X_0, \dots, X_t; X_0 \geq b\}$ it holds that $\mathbb{P}\left[T^* \leq \exp\left(\frac{\epsilon\ell}{132r^2}\right)\right] = O\left(\exp\left(-\frac{\epsilon\ell}{132r^2}\right)\right)$*

THEOREM A.2 (BERNSTEIN'S INEQUALITY). *Let X_1, \dots, X_n be independent zero-mean random variables. Suppose that $|X_i| \leq M$ almost surely, for all i . Then, for all positive t ,*

$$\mathbb{P}\left(\sum_{i=1}^n X_i \geq t\right) \leq \exp\left(-\frac{\frac{1}{2}t^2}{\sum_{i=1}^n \mathbb{E}[X_i^2] + \frac{1}{3}Mt}\right)$$

References

- [1] Yehuda Afek and Eli Gafni. 2013. Asynchrony from synchrony. In *Distributed Computing and Networking: 14th International Conference, ICDCN 2013, Mumbai, India, January 3–6, 2013. Proceedings 14*. Springer, 225–239.
- [2] Dan Alistarh, James Aspnes, David Eisenstat, Rati Gelashvili, and Ronald L. Rivest. 2017. Time-Space Trade-offs in Population Protocols. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'17*. 2560–2579.
- [3] 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'18*. 2221–2239.
- [4] Dan Alistarh and Rati Gelashvili. 2018. Recent Algorithmic Advances in Population Protocols. *SIGACT News* 49, 3 (2018), 63–73. <https://doi.org/10.1145/3289137.3289150>
- [5] Dan Alistarh, Rati Gelashvili, and Milan Vojnovic. 2015. Fast and Exact Majority in Population Protocols. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC'15*, Chryssis Georgiou and Paul G. Spirakis (Eds.). ACM, 47–56. <https://doi.org/10.1145/2767386.2767429>
- [6] 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 *Proceedings of the 2023 ACM Symposium on Principles of Distributed Computing, PODC 2023, Orlando, FL, USA, June 19–23, 2023*, Rotem Oshman, Alexandre Nolin, Magnús M. Halldórsson, and Alkida Balliu (Eds.). ACM, 13–23. <https://doi.org/10.1145/3583668.3594589>
- [7] Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. 2006. Computation in networks of passively mobile finite-state sensors. *Distributed Computing* 18, 4 (2006), 235–253.
- [8] Dana Angluin, James Aspnes, and David Eisenstat. 2008. Fast computation by population protocols with a leader. *Distributed Computing* 21, 3 (Sept. 2008), 183–199.
- [9] 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 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2022*. 3417–3429. <https://doi.org/10.1137/1.9781611977073.135>
- [10] Luca Becchetti, Andrea Clementi, Emanuele Natale, Francesco Pasquale, and Riccardo Silvestri. 2015. Plurality Consensus in the Gossip Model. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'15*. 371–390.
- [11] 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, Virtual Event, Italy, August 3–7, 2020*, Yuval Emek and Christian Cachin (Eds.). ACM, 191–199. <https://doi.org/10.1145/3382734.3405747>
- [12] Petra Berenbrink, Robert Elsässer, Tom Friedetzky, Dominik Kaaser, Peter Kling, and Tomasz Radzik. 2018. A population protocol for exact majority with $O(\log^{5/3} n)$ stabilization time and $O(\log n)$ states. In *Proceedings of the 32nd International Symposium on Distributed Computing, DISC'18*.
- [13] 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. Full version available at arXiv:1805.04586.
- [14] Andreas Bilke, Colin Cooper, Robert Elsässer, and Tomasz Radzik. 2017. Brief Announcement: Population Protocols for Leader Election and Exact Majority with $O(\log^2 n)$ States and $O(\log^2 n)$ Convergence Time. In *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC'17*. 451–453. Full version available at arXiv:1705.01146.
- [15] Luca Cardelli and Attila Csiksz-Nagy. 2012. The cell cycle switch computes approximate majority. *Nature Scientific Reports* 2 (2012), 656.
- [16] Yuen-Jyue Chen, Neil Dalchau, Niranjan Srinivas, Andrew Phillips, Luca Cardelli, David Soloveichik, and Georg Seelig. 2013. Programmable chemical controllers made from DNA. *Nature Nanotechnology* 8, 10 (2013), 755–762.
- [17] Andrea Clementi, Mohesen Ghaffari, Luciano Gualà, Emanuele Natale, Francesco Pasquale, and Giacomo Scornavacca. 2018. A Tight Analysis of the Parallel Undecided-State Dynamics with Two Colors. In *Proceedings of the 43rd International Symposium on Mathematical Foundations of Computer Science, MFCS'18*. 28:1–28:15.
- [18] David Doty, Mahsa Eftekhari, Leszek Gasieniec, Eric E. Severson, Przemysław 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, Denver, CO, USA, February 7–10, 2022*. IEEE, 1044–1055. <https://doi.org/10.1109/FOCS52979.2021.00104>
- [19] Moez Draief and Milan Vojnovic. 2010. Convergence Speed of Binary Interval Consensus. In *Proceedings of the 29th IEEE International Conference on Computer Communications, INFOCOM'10*. IEEE, 1792–1800. <https://doi.org/10.1109/INFCOM.2010.5461999>
- [20] Robert Elsässer and Tomasz Radzik. 2018. Recent Results in Population Protocols for Exact Majority and Leader Election. *Bull. EATCS* 126 (2018). <http://bulletin.eatcs.org/index.php/beatcs/article/view/549/546>
- [21] 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)*, Panagioti Fatourou, Ernesto Jiménez, and Fernando Pedone (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 14:1–14:14. <https://doi.org/10.4230/LIPIcs.OPODIS.2016.14>
- [22] Adrian Kosowski and Przemysław Uznanski. 2018. Brief Announcement: Population Protocols Are Fast. In *Proceedings of the 37th ACM Symposium on Principles of Distributed Computing, PODC'18*. 475–477. <https://doi.org/10.1145/3212734.3212788> Full version available at arXiv:1802.06872.
- [23] Fabian Kuhn and Rotem Oshman. 2011. Dynamic networks: models and algorithms. *ACM SIGACT News* 42, 1 (2011), 82–96.
- [24] 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
- [25] George B. Mertzios, Sotiris E. Nikolettseas, Christoforos L. Raptopoulos, and Paul G. Spirakis. 2014. Determining Majority in Networks with Local Interactions and Very Small Local Memory. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming, ICALP'14*. Springer Berlin Heidelberg, 871–882. https://doi.org/10.1007/978-3-662-43948-7_72
- [26] Othon Michail, George Skretas, and Paul G. Spirakis. 2020. Distributed computation and reconfiguration in actively dynamic networks. In *Proceedings of the 39th Symposium on Principles of Distributed Computing*. 448–457.
- [27] 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)*, Pinar Heggernes (Ed.). Springer, 323–338. https://doi.org/10.1007/978-3-030-17402-6_27
- [28] Pietro S. Oliveto and Carsten Witt. 2015. Improved time complexity analysis of the Simple Genetic Algorithm. *Theor. Comput. Sci.* 605 (2015), 21–41. <https://doi.org/10.1016/j.tcs.2015.01.002>
- [29] David Soloveichik, Matthew Cook, Erik Winfree, and Jehoshua Brook. 2008. Computation with finite stochastic chemical reaction networks. *Natural Computing* 4, 7 (2008), 615–633.