

On the Hierarchy of Distributed Majority Protocols

Petra Berenbrink  

Universität Hamburg, Germany

Amin Coja-Oghlan  

TU Dortmund University, Germany

Oliver Gebhard  

TU Dortmund University, Germany

Max Hahn-Klimroth  

TU Dortmund University, Germany

Dominik Kaaser  

TU Hamburg, Germany

Malin Rau  

Universität Hamburg, Germany

Abstract

We study the *consensus* problem among n agents, defined as follows. Initially, each agent holds one of two possible opinions. The goal is to reach a consensus configuration in which every agent shares the same opinion. To this end, agents randomly sample other agents and update their opinion according to a simple update function depending on the sampled opinions.

We consider two communication models: the *gossip model* and a variant of the *population model*. In the gossip model, agents are activated in parallel, synchronous rounds. In the population model, one agent is activated after the other in a sequence of discrete time steps. For both models we analyze the following natural family of majority processes called j -Majority: when activated, every agent samples j other agents uniformly at random (with replacement) and adopts the majority opinion among the sample (breaking ties uniformly at random). As our main result we show a hierarchy among majority protocols: $(j + 1)$ -Majority (for $j > 1$) converges *stochastically faster* than j -Majority for any initial opinion configuration. In our analysis we use Strassen's Theorem to prove the existence of a coupling. This gives an affirmative answer for the case of two opinions to an open question asked by Berenbrink et al. [PODC 2017].

2012 ACM Subject Classification Theory of computation → Distributed algorithms; Theory of computation → Random walks and Markov chains; Mathematics of computing → Stochastic processes

Keywords and phrases Consensus, Majority, Hierarchy, Stochastic Dominance, Population Protocols, Gossip Model, Strassen's Theorem

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2022.23

Related Version *Full Version:* <https://arxiv.org/abs/2205.08203>

Funding *Petra Berenbrink:* DFG FOR 2975

Amin Coja-Oghlan: DFG FOR 2975

Oliver Gebhard: DFG CO 646/3

Max Hahn-Klimroth: DFG FOR 2975

Malin Rau: DFG FOR 2975

1 Introduction

We consider the problem of *consensus* in a distributed system of n identical, anonymous agents. Initially each agent has one of two opinions and the goal is that all agents agree on the same opinion. Reaching consensus is a fundamental task in distributed computing



© Petra Berenbrink, Amin Coja-Oghlan, Oliver Gebhard, Max Hahn-Klimroth, Dominik Kaaser, and Malin Rau;

licensed under Creative Commons License CC-BY 4.0

26th International Conference on Principles of Distributed Systems (OPODIS 2022).

Editors: Eshcar Hillel, Roberto Palmieri, and Etienne Rivière; Article No. 23; pp. 23:1–23:19



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

with a multitude of applications, including fault tolerance in distributed sensor array, clock synchronization, control of autonomous robots, or blockchains. In computational sciences, consensus protocols model, e.g., dynamic particle systems or biological processes. In social sciences, consensus protocols have been studied in the context of opinion formation processes among social interaction systems. See [8] for a quite recent survey including references and further applications.

We study the simple and well-known class of j -Majority protocols [10, 35, 12] in two communication models, the classical *gossip model* [19, 9, 8] and a *sequential model*, a variant of the prominent *population model* [4]. In the gossip model, all agents are activated in parallel, synchronous rounds. In the sequential model, one agent is activated after the other uniformly at random. Every activated agent u considers the opinions of j agents v_1, \dots, v_j sampled uniformly at random (with replacement). It then adopts the majority opinion among the sampled opinions, breaking ties uniformly at random. We are interested in the time it takes until the protocol *converges* such that all agents share the same opinion. Setting $j = 1$ yields the so-called VOTER process [17]. A variant of 2-Majority with lazy tie-breaking is known as TWO-SAMPLE VOTING [23] or the TWOCHOICES process [35], and the 3-Majority dynamics is analyzed in [10].

The main idea of majority processes with $j > 1$ is to speed up the convergence time. For the VOTER process in the gossip model, the convergence time is linear in n (independently of the number of initial opinions) [17], whereas the convergence time of 3-Majority is $O(k \log n)$ for $k = o(n)$ possible initial opinions [35]. In [12] the authors compare the TWOCHOICES process to 3-Majority. They show a stochastic dominance of the convergence time of 3-Majority over the convergence time of VOTER and TWOCHOICES, assuming k initial opinions. For j -Majority, they conjecture a *hierarchy* of protocols (see Conjecture 6.1 in [12]). In particular, they ask whether one can couple j -Majority and $(j + 1)$ -Majority for $j \in \mathbb{N}$ such that $(j + 1)$ -Majority is stochastically faster than j -Majority.

In this paper, we settle the matter for the case of $k = 2$ opinions and prove the existence of such a hierarchy of majority protocols. Intuitively, this establishes that the processes converge *faster* (or at least *equally fast*) for larger values of j . Let T_j be the random variable for the convergence time of j -Majority. We formally prove that T_{j+1} stochastically minorizes T_j , written $T_{j+1} \preceq T_j$, assuming both processes start in the same configuration. Formally, we show that $\Pr[T_{j+1} \geq t] \leq \Pr[T_j \geq t]$ for any $t \in \mathbb{N}$. Our main technical contribution is the formal proof of this stochastic dominance.

Our proof has its foundations in quite natural observations regarding the transition properties of the j -Majority processes. Similar results for individual steps of the process have been shown, e.g., in [33]. However, formally proving and maintaining the stochastic dominance over all possible configurations requires a lot of care, and to the best of our knowledge, our result is the first proof of stochastic dominance that covers the entire execution of j -Majority for all $j \in \mathbb{N}$ in the setting with two opinions. To motivate the obstacles we have to overcome, observe that the process is influenced by opposing *forces*. Specifically, in order to make progress, an agent from the minority opinion must be activated to interact with at least $j/2$ agents from the majority opinion. Activating an agent with minority opinion becomes less likely with increasing majority, while sampling at least $j/2$ agents with majority opinion becomes more likely with increasing majority. In our analysis we carefully prove that these forces balance out in a favorable manner.

Finally, we consider 3-Majority. We show an asymptotically optimal bound in the sequential model on the convergence time of $O(n \log n)$ activations. This matches a similar result shown by Ghaffari and Lengler [35] for 3-Majority in the gossip model. Our theoretical findings are complemented by empirical results. We simulate j -Majority processes for various values of j and large numbers of agents ranging from $n = 10^2$ to $n = 10^8$.

1.1 Related Work

Consensus in the Gossip Model. A simple and natural consensus process is the so-called VOTER process [37, 43, 22, 17, 38] where every agent adopts the opinion of a single, randomly chosen agent in each round. The expected convergence time of VOTER in the gossip model is at least linear [17]. In order to speed up the process, two related protocols have been proposed, namely the TWOCHOICES process [31, 23, 24, 25] and the 3-MAJORITY dynamics [10, 35, 12]. In both processes, each agent u takes three opinions and updates its opinion to the majority among the sample. In the TWOCHOICES process, u takes its own opinion and samples two opinions u.a.r. Ties are broken towards u 's own opinion. In the 3-MAJORITY dynamics, u samples three opinions u.a.r. breaking ties randomly. In [35] the authors consider arbitrary initial configurations in the gossip model. They show that TWOCHOICES with $k = O(\sqrt{n/\log n})$ and 3-MAJORITY with $k = O(n^{1/3}/\log n)$ reach consensus in $O(k \cdot \log n)$ rounds, improving a result by Becchetti et al. [10]. For arbitrary k they show that 3-MAJORITY reaches consensus in $O(n^{2/3} \log^{3/2} n)$ rounds w.h.p., improving a result by Berenbrink et al. [12].

Schoenebeck and Yu [45] consider a generalization of multi-sample consensus protocols on complete and Erdős-Rényi graphs for two opinions. Their probabilistic model covers various consensus processes, including j -Majority, by using a so-called *update rule*, a function $f: [0, 1] \rightarrow [0, 1]$. In each round, every agent u adopts opinion a with probability $f(\alpha(u))$ for some function f , where $\alpha(u)$ is the fraction of neighbors of agent u that have opinion a . Depending on certain natural properties on f , they analyze the convergence time for complete graphs and Erdős-Rényi graphs.

Another related process is the MEDIANRULE [26], where in each round every agent adopts the median of its own opinion and two sampled opinions, assuming a total order among opinions. It reaches consensus in $O(\log k \log \log n + \log n)$ rounds w.h.p. For two opinions the MEDIANRULE is equivalent to the TWOCHOICES process, and their analysis is tight. For the case of $k > 2$ opinions we remark that assuming a total order among the opinions is a strong assumption that is not required by any of the other protocols.

Finally, considerable amount of work has been spent on analyzing the so-called undecided state dynamics introduced by Angluin et al. [5]. The basic idea is that whenever two agents with different opinions interact, they lose their opinions and become *undecided*, and undecided agents adopt the first opinion they encounter. Clementi et al. [20] study the undecided state dynamics in the gossip model. They consider two opinions and show that the protocol reaches consensus in $O(\log n)$ rounds w.h.p. If there is a so-called *bias* of order $\Omega(\sqrt{n \log n})$, the initial plurality opinion prevails. The (additive) bias is the difference between the numbers of agents holding either opinion. Becchetti et al. [9] analyze the undecided state dynamics for $k = O(n/\log n)^{1/3}$ opinions and show a convergence time of $O(k \cdot \log n)$ rounds w.h.p. Bankhamer et al. [36], Berenbrink et al. [16], and Ghaffari and Parter [6] consider a synchronized variant that runs in phases of length $\Theta(\log k)$. Agents can become undecided only at the start of such a phase and use the rest of the phase to obtain a new opinion. These synchronized protocols achieve consensus in $O(\log^2 n)$ rounds w.h.p. and can be further refined using more sophisticated synchronization mechanisms.

Majority and Consensus in the Population Model. In *exact* majority the goal is to identify the majority among two possible opinions, even if the bias is as small as only one [30, 41, 3, 42, 29, 1, 2, 18, 39, 14, 15, 11, 28]. The best known protocol by Doty et al. [28] solves exact majority with $O(\log n)$ states and $O(\log n)$ parallel time, both in expectation

and w.h.p. This is optimal: it takes at least $\Omega(n \log n)$ interactions until each agent interacts at least once, and any majority protocol which stabilizes in expected $n^{1-\Omega(1)}$ parallel time requires at least $\Omega(\log n)$ states (under some natural conditions, see [2]).

Approximate majority is easier: a simple 3-state protocol [5, 21] reaches consensus w.h.p. in $O(\log n)$ parallel time and correctly identifies the initial majority w.h.p. if an initial bias of order $\Omega(\sqrt{n \log n})$ is present. Condon et al. [21] also consider a variant of the 3-Majority process in (a variant of) the gossip model where three randomly chosen agents interact. They show a parallel convergence time of $O(k \log n)$ w.h.p., provided a sufficiently large initial bias is present. Furthermore, Kosowski and Uznanski [39] mention a protocol which determines the exact majority in $O(\log^2 n)$ parallel time w.h.p. using only constantly many states.

Less is known about population protocols that solve consensus among more than two opinions. One line of research considers only the required number of states to eventually identify the opinion with the largest initial support correctly. For this problem, Natale and Ramezani [44] show a lower bound of $\Omega(k^2)$ states via an indistinguishability argument. The currently best known protocol uses $O(k^6)$ states if there is an order among the opinions and $O(k^{11})$ states otherwise [34]. Sacrificing the strong guarantees of always-correct exact plurality consensus, Bankhamer et al. [6] achieve *approximate consensus* in $O(\log^2 n)$ parallel time w.h.p. using only $O(k \log n)$ states. If there is an initial bias of order $\Omega(\sqrt{n \log n})$, the initial plurality opinion wins w.h.p. In [7] another variant of the population model is considered where agents are activated by random clocks. At each clock tick, every agent may open communication channels to constantly many other agents chosen uniformly at random or from a list of at most constantly many agents contacted in previous steps. In this model, opening communication channels is subject to a random delay. The authors show that consensus is reached by all but a $1/\text{poly } \log n$ fraction of agents in $O(\log \log_\alpha k \log k + \log \log n)$ parallel time w.h.p., provided a sufficiently large bias is present.

1.2 Models and Results

Gossip Model. In the gossip model [19, 9, 8] all agents are activated simultaneously in synchronous rounds. In each round every agent u opens a communication channel to j agents v_1, \dots, v_j chosen independently and uniformly at random with replacement. (For simplicity we also allow that $v_i = u$ and assume that the v_i are sampled with replacement.) The running time (or *convergence time*) of a majority protocol is measured in the numbers of rounds until all agents agree on the same opinion.

Sequential Model. The population model was introduced by Angluin et al. [4] to model systems of resource limited mobile agents that perform a computation via a sequence of pairwise interactions. We consider a variant where in each time step one agent u is chosen uniformly at random to interact with j randomly sampled agents v_1, \dots, v_j . (As before, we do not rule out that $u = v_i$ for some i). When u is activated it updates its opinion according to the random sample. The running time is measured in the number of interactions. To allow for a comparison with the (inherently) parallel gossip model, the so-called *parallel time* is defined as the number of interactions divided by the number of agents n . Note that our processes do not *halt*: agents do not know that consensus has been reached (see also the impossibility result in [27]).

j -Majority Processes. In the following we use P_j to denote the j -Majority process. When executing process P_j , the system transitions through a sequence of *configurations* $(C_t)_{t \in \mathbb{N}_0}$. At time $t \in \mathbb{N}_0$ the *configuration* $C_t \in \{a, b\}^n$ assigns each agent an opinion in $\{a, b\}$. In

our analysis we are interested in the number of agents with majority opinion. We will always assume w.l.o.g. that a is the majority opinion and we denote a *state* X_t as the number of agents with majority opinion in configuration C_t . The configuration C_0 at time 0 is called the *initial configuration* and the corresponding state X_0 is called the *initial state*. The *convergence time* $T_j(C_0)$ is defined as the first time where all agents have the same opinion when starting process P_j in initial configuration C_0 . Note that the convergence time only depends on the number of agents with majority opinion since two agents with the same opinion are not distinguishable. Hence we write $T_j(X_0)$ in the following. Formally, $T_j(X_0)$ is a stopping time defined as $T_j(X_0) = \min \{ t \in \mathbb{N}_0 \mid X_t = n \}$. Each transition of the system is done according to the following update rule.

► **Definition (Process P_j).** *Agents are activated according to either the gossip model or the sequential model. In process P_j each activated agent u samples j agents with replacement and adopts the majority opinion among the sample, breaking ties uniformly at random.*

Note that tie-breaking is not required in process P_{2j+1} (i.e., when every agent samples an odd number of agents). Since we have $k = 2$ opinions we are guaranteed to have a clear majority in this case.

Stochastic Dominance. Before we formally present our result, it remains to define stochastic dominance.

► **Definition (Stochastic Dominance).** *Let \mathcal{E} be a Polish space¹ with a partial ordering $\leq_{\mathcal{E}}$. Let $\mu, \nu \in \mathcal{P}(\mathcal{E})$ be probability measures on \mathcal{E} . If, for every $x \in \mathcal{E}$, we have*

$$\mu(\{y \in \mathcal{E} : y \geq_{\mathcal{E}} x\}) \geq \nu(\{y \in \mathcal{E} : y \geq_{\mathcal{E}} x\}),$$

we say that μ stochastically dominates ν . In this case we also say that μ majorizes ν (written as $\nu \succeq \mu$) or ν minorizes μ ($\nu \preceq \mu$).

We now formally state our main result which applies for both communication models, the gossip model and the sequential model.

► **Theorem 1 (Main Result).** *Let $T_j(X_0)$ be the convergence time of process P_j with initial state X_0 in either the gossip model or the sequential model. Then*

$$T_{j+1}(X_0) \preceq T_j(X_0) \quad \text{for any } j > 1.$$

Furthermore, for all $j > 1$,

$$\mathbb{E}[T_{2j+2}(X_0)] = \mathbb{E}[T_{2j+1}(X_0)] < \mathbb{E}[T_{2j}(X_0)].$$

In our second result we show that 3-Majority P_3 converges in $O(n \log n)$ time w.h.p.² To the best of our knowledge, this is the first analysis of 3-Majority with sequential updates. Our proof is similar to the proof by Condon et al. [21] for the convergence time of approximate majority in tri-molecular chemical reaction networks. We emphasize that Theorem 1 implies that all j -Majority processes with $j > 3$ converge in $O(n \log n)$ time w.h.p.

¹ A Polish space is a complete metric space with a countable dense subset.

² The expression *with high probability (w.h.p.)* refers to a probability of $1 - n^{-\Omega(1)}$.

► **Theorem 2.** Let $T_3(X_0)$ be the convergence time of the 3-Majority process P_3 in the sequential model with initial configuration X_0 .

1. It holds that $T_3(X_0) \leq O(n \log n)$ w.h.p.
2. If $X_0 \geq n/2 + \zeta \sqrt{n \log n}$ for some sufficiently large constant $\zeta > 0$ then the initial majority opinion wins w.h.p.

We remark that the convergence time of $O(n \log n)$ is asymptotically tight. Indeed, for any number of time steps in $o(n \log n)$ there is a constant probability that two agents with opposing opinions are not activated even once.

2 Analysis

In this section we formally prove our theorems. We prove Theorem 1 in Section 2.1 and Section 2.2 for the sequential model and the gossip model, respectively. Theorem 2 is then shown in Section 2.3. All technical details for the rigorous proofs can be found in the full version [13].


2.1 Sequential Model

We start our analysis with a comparison of one step of the processes P_j and P_{j+1} at time t when starting in an identical state X_t . We are able to express the differences in the probabilities of increasing the majority opinion, decreasing it or remaining in the same state for the both processes. To this end, we visualize a possible coupling by a decision tree that incorporates all the different possibilities. We will observe that, within this one step, we can couple both processes such that the supposedly faster process increases the majority opinion with probability one if the supposedly slower process increases this opinion. This coupling will be guaranteed by an application of *Strassen's Theorem*.

The proof of the main result will be conducted inductively. We start both processes in the same initial state and assume that there is a majority opinion a . Now, the aforementioned coupling ensures that, after the first step, the supposedly faster process will have at least as many agents of opinion a than the supposedly slower process. Now, we show a kind of *monotony* in the studied processes. Assume we have two instances of the same process, one in state $X_t = s$ and one in state $X'_t = s'$ where X_t, X'_t denote the number of agents with opinion a after t steps. If $s > s'$, then the random variable X_{t+1} will stochastically dominate X'_{t+1} , formally $X'_{t+1} \preceq X_{t+1}$. This observation is crucial. It allows us to show that in the second step, we can again construct a coupling such that, if the supposedly slower process moves, the supposedly faster process does as well almost surely. Indeed, either both processes are in the same state, then we find the stochastic dominance by the decision trees, or the fast process has more agents of opinion a . But as stochastic dominance is transitive, we can construct a coupling via the triangle inequality.

Finally, we will describe the overall coupling of the two processes as the *path-coupling* along those couplings per step which will prove the first part of Theorem 1. The second part will follow analogously as we can show via the decision trees that in the comparison of P_{2j-1} and P_{2j} , the chance to obtain the same state in the next step is equal under both processes while in the comparison of P_{2j} and P_{2j+1} those decision trees show that the probability of increasing the majority opinion is larger in P_{2j+1} .

► **Observation 3.** The processes P_1 and P_3 have, almost surely, a finite stopping time.



■ **Figure 1** Decision tree comparing P_{2j} and P_{2j+1} .

For the sequential process, we show this for P_3 in Section 2.3, while for P_1 this follows by the results of Schoenebeck and Yu [45]. For the gossip model, this is proven in [35]. In this setting, Strassen’s Theorem guarantees the existence of a coupling $\gamma \in \mathcal{P}(\mathcal{E}^2)$ of μ and ν with the following property.

► **Theorem 4** (Strassen’s Theorem [46]). *Let μ, ν be probability measures on a Polish space endowed with a partial ordering \preceq such that μ stochastically dominates ν . Let $X \sim \mu$ and $Y \sim \nu$, then there is a coupling γ of μ and ν such that, if $(\hat{X}, \hat{Y}) \sim \gamma$, we have*

$$X \stackrel{d}{=} \hat{X}, \quad Y \stackrel{d}{=} \hat{Y} \quad \text{and} \quad \Pr[\hat{Y} \preceq \hat{X}] = 1.$$

In words, this means that if μ stochastically dominates ν , X is sampled from μ , and Y is sampled from ν , there is a coupling under which $X \leq Y$ almost surely (with probability 1).

► **Lemma 5.** *We find for P_k the following. Let X_t denote the number of agents with majority opinion at time t . If $s > s'$, then for all $d \in \{0, 1, \dots, n\}$*

$$\Pr[X_{t+1} \geq d \mid X_t = s] \geq \Pr[X_{t+1} \geq d \mid X_t = s'].$$

We provide the detailed calculation in the full version [13] and get the following corollary.

► **Corollary 6.** *For any two processes P, P' , we find the following stochastic dominance. Let X_t denote the number of agents with opinion a with respect to process P at time t and let X'_t be the analogous quantity with respect to P' . Assume that for any $d \in [n]$*

$$\Pr[X_{t+1} \geq d \mid X_t = s] \geq \Pr[X'_{t+1} \geq d \mid X'_t = s],$$

then we have also


$$\Pr[X_{t+1} \geq d \mid X_t = s + t'] \geq \Pr[X'_{t+1} \geq d \mid X'_t = s],$$

for any $d \in [n]$ and $t' > 0$ such that $s + t' \leq n$.

Let $X_t^{(k)}$ denote the number of agents with majority opinion after step t of process P_k for any $k \in \mathbb{N}$. Furthermore, for a given agent x we denote by $x_t^{(k)}$ its opinion in process P_k at time t . In the following we compare two processes with each other. The comparisons of P_{2j} (even) to P_{2j+1} (odd) and P_{2j-1} (odd) to P_{2j} (even) require slightly different calculations. Therefore, we have to show two similar lemmas for these two cases, Lemma 7 for the former case and Lemma 9 for the latter case.

First we compare two successive processes P_{2j} and P_{2j+1} . The following lemma states that in process P_{2j+1} it is more likely for an agent with opinion b to change to a while in process P_{2j} it is more likely that an agent with opinion a changes to opinion b than in the other process respectively.

23:8 On the Hierarchy of Distributed Majority Protocols



■ **Figure 2** Decision tree comparing P_{2j-1} and P_{2j} .

► **Lemma 7.** Let x be an agent that is updated in the next step, $x_t^{(k)}$ its opinion in process P_k at time t , $s \in [n]$ and $\alpha = \frac{s}{n}$. It holds that

$$\begin{aligned} & \Pr \left[x_{t+1}^{(2j+1)} = a \mid x_t^{(2j+1)} = b, X_t^{(2j+1)} = s \right] \\ &= \Pr \left[x_{t+1}^{(2j)} = a \mid x_t^{(2j)} = b, X_t^{(2j)} = s \right] + \frac{(2\alpha - 1) \binom{2j}{j}}{2} \alpha^j (1 - \alpha)^j \end{aligned}$$

and

$$\begin{aligned} & \Pr \left[x_{t+1}^{(2j+1)} = b \mid x_t^{(2j+1)} = a, X_t^{(2j+1)} = s \right] \\ &= \Pr \left[x_{t+1}^{(2j)} = b \mid x_t^{(2j)} = a, X_t^{(2j)} = s \right] - \frac{(2\alpha - 1) \binom{2j}{j}}{2} \alpha^j (1 - \alpha)^j. \end{aligned}$$

To prove these equations it is sufficient to study the cases where the two processes have a different outcome. The probability for these cases directly reflects the difference in probability for that specific outcome. These cases are highlighted in Figure 1. We provide the detailed calculation in the full version.

This difference in probabilities allows us to prove that, given the same state, P_{2j+1} stochastically dominates the process P_{2j} in the next step:

► **Lemma 8.** For each $j \in \mathbb{N}_0$ and each $s \in \mathbb{N}$ with $s > n/2$ and any $d \in [n]$ it holds that

$$\Pr \left[X_{t+1}^{(2j+1)} \geq d \mid X_t^{(2j+1)} = s \right] \geq \Pr \left[X_{t+1}^{(2j)} \geq d \mid X_t^{(2j)} = s \right].$$

Note that this inequality follows trivially for $d \leq s - 1$ and $d > s + 1$. To prove the property for the cases $d = s$ and $d = s + 1$ we can directly use the properties from Lemma 7.

On the other hand, when comparing the processes P_{2j-1} and P_{2j} with respect to the difference in probability for an agent to change its opinion, we note that there is no difference in the probabilities given that all agents are in the same state.

► **Lemma 9.** Let x be an agent that is updated in the next step, $x_t^{(k)}$ its opinion in process P_k at time t , and $s \in [n]$. It holds that

$$\Pr \left[x_{t+1}^{(2j-1)} = a \mid x_t^{(2j-1)} = b, X_t^{(2j-1)} = s \right] = \Pr \left[x_{t+1}^{(2j)} = a \mid x_t^{(2j)} = b, X_t^{(2j)} = s \right]$$

and

$$\Pr \left[x_{t+1}^{(2j-1)} = b \mid x_t^{(2j-1)} = a, X_t^{(2j-1)} = s \right] = \Pr \left[x_{t+1}^{(2j)} = b \mid x_t^{(2j)} = a, X_t^{(2j)} = s \right].$$

A similar statement has previously been shown by Fraigniaud and Natale [33] for a related model. The proof of Lemma 9 is analogous to the proof of Lemma 7. For completeness, it can be found in the full version

► **Lemma 10.** *For each $j \in \mathbb{N}_0$ and each $s \in \mathbb{N}$ with $s > n/2$ and any $d \in [n]$ it holds that*

$$\Pr\left[X_{t+1}^{(2j)} \geq d \mid X_t^{(2j)} = s\right] = \Pr\left[X_{t+1}^{(2j-1)} \geq d \mid X_t^{(2j-1)} = s\right].$$

The proof can be found in the full version. We are now ready to put everything together and prove our main result for the sequential model.

Proof of Theorem 1. We prove Theorem 1 by induction given the initial state X_0 and start with the case $T_{2j}(X_0) \preceq T_{2j+1}(X_0)$. Given X_0 , Lemma 8 guarantees that for all $s > 0$

$$\Pr\left[X_1^{(2j+1)} \geq s \mid X_0\right] \geq \Pr\left[X_1^{(2j)} \geq s \mid X_0\right].$$

Therefore, by Theorem 4, we find a coupling γ_1 such that, under γ_1 , $X_1^{(2j+1)} \geq X_1^{(2j)}$ almost surely. Now, assume that we constructed a coupling $\gamma^{(t)} = \gamma_1 \otimes \dots \otimes \gamma_t$ of $(X_1^{2j}, \dots, X_t^{2j})$ and $(X_1^{2j+1}, \dots, X_t^{2j+1})$, where \otimes denotes the product measure. Under $\gamma^{(t)}$ we have by induction hypothesis that

$$\Pr_{\gamma^{(t)}}\left[X_t^{2j+1} \geq X_t^{2j}\right] = 1.$$

Therefore, by Corollary 6 and Lemma 8, we find given $X_t^{(2j+1)} \geq X_t^{(2j)}$ that

$$\Pr\left[X_{t+1}^{(2j+1)} \geq s \mid X_t^{(2j+1)}\right] \geq \Pr\left[X_{t+1}^{(2j)} \geq s \mid X_t^{(2j)}\right].$$

Thus, Theorem 4 implies, given $X_t^{(2j+1)} \geq X_t^{(2j)}$ the existence of a coupling γ_{t+1} such that

$$\Pr[\gamma_{t+1}]X_{t+1}^{2j+1} \geq X_{t+1}^{2j} = 1.$$

We define $\gamma^{(t+1)} = \gamma^{(t)} \otimes \gamma_{t+1}$ and $T_{2j}(X_0) \preceq T_{2j+1}(X_0)$ follows by induction.

Next, we need to prove that $T_{2j-1}(X_0) \preceq T_{2j}(X_0)$. This follows completely analogously with Lemma 8 replaced by Lemma 10.

Finally, we need to construct the bounds on the expectation. Given the coupling $\gamma^{T(2j+1)(X_0)}$ of P_{2j+1} and P_{2j} , we find that, under this coupling, for every step $t = 1 \dots T_{2j+1}$, we have $X_t^{(2j+1)} \geq X_t^{(2j)}$ almost surely and therefore $\mathbb{E}[T_{2j+1}(X_0)] \leq \mathbb{E}[T_{2j}(X_0)]$. In more detail, due to the coupling γ , we know that

$$\Pr\left[X_{t+1}^{(2j+1)} \geq s \mid X_0\right] \geq \Pr\left[X_{t+1}^{(2j)} \geq s \mid X_0\right]$$

and hence

$$\Pr\left[X_{t+1}^{(2j+1)} < s \mid X_0\right] \leq \Pr\left[X_{t+1}^{(2j)} < s \mid X_0\right]$$

for each $s \geq n/2$ and $t \in \mathbb{N}$. Furthermore, for the first possible convergence time it holds that

$$\begin{aligned} \Pr\left[X_{n-s}^{(2j)} < n \mid X_0 = s\right] &= 1 - \Pr\left[X_{n-s}^{(2j)} = n \mid X_0 = s\right] \\ &= 1 - \prod_{i=1}^{n-s} \Pr\left[X_i^{(2j)} = s+i \mid X_{i-1}^{(2j)} = s+i-1\right] \\ &\stackrel{\text{Lemma 7}}{>} 1 - \prod_{i=1}^{n-s} \Pr\left[X_i^{(2j+1)} = s+i \mid X_{i-1}^{(2j+1)} = s+i-1\right] \\ &= 1 - \Pr\left[X_{n-s}^{(2j+1)} = n \mid X_0 = s\right] = \Pr\left[X_{n-s}^{(2j+1)} < n \mid X_0 = s\right]. \end{aligned}$$

23:10 On the Hierarchy of Distributed Majority Protocols

Therefore,

$$\begin{aligned} \mathbb{E}[T_{2j}(X_0)] &= \sum_{t \geq 0} \Pr[T_{2j}(X_0) > t] = \sum_{t \geq 0} \Pr[X_t^{(2j)} < n \mid X_0] \\ &> \sum_{t \geq 0} \Pr[X_t^{(2j+1)} < n \mid X_0] = \sum_{t \geq 0} \Pr[T_{2j+1}(X_0) > t] = \mathbb{E}[T_{2j+1}(X_0)]. \end{aligned}$$

Next we prove the equality in expectation for the convergence time of the processes P_{2j-1} and P_{2j} . To this end, we get from Lemma 10 that

$$\Pr[X_t^{(2j-1)} < n \mid X_{t-1} = s] = \Pr[X_t^{(2j)} < n \mid X_{t-1} = s].$$

Therefore, inductively,

$$\Pr[X_t^{(2j-1)} < n \mid X_0 = s] = \Pr[X_t^{(2j)} < n \mid X_0 = s].$$

But then

$$\begin{aligned} \mathbb{E}[T_{2j-1}(X_0)] &= \sum_{t \geq 0} \Pr[T_{2j-1}(X_0) > t] = \sum_{t \geq 0} \Pr[X_t^{(2j-1)} < n \mid X_0] \\ &= \sum_{t \geq 0} \Pr[X_t^{(2j)} < n \mid X_0] = \sum_{t \geq 0} \Pr[T_{2j}(X_0) > t] = \mathbb{E}[T_{2j}(X_0)]. \quad \blacktriangleleft \end{aligned}$$

2.2 Gossip Model

We now extend the previous analysis to the gossip model. Recall that in this model all agents are activated in parallel rounds. In such a round, all agents sample j other agents v_1, \dots, v_j u.a.r. Then they compute their new opinion as the majority opinion among the sample, breaking ties u.a.r. Here, the agents use the opinions of the other agents from the beginning of the round. At the end of the round (once all agents have computed the new opinion) all agents synchronously update their opinion to the new value.

Proof of Theorem 1 for the Gossip Model. In our extended analysis we use a coupling of the two parallel processes similarly to the coupling of one step of the sequential model. Observe that in process P_{2j} every agent samples $2j$ agents u.a.r., while in process P_{2j+1} every agent samples $2j + 1$ agents. Therefore, process P_{2j} makes $2j \cdot n$ random choices from $[n]$ in each round, while P_{2j+1} makes $(2j + 1) \cdot n$ random choices. We use the straight-forward coupling and define that the $2j$ choices of every agent u in P_{2j} are identical to the first $2j$ choices of agent u in process P_{2j+1} .

We now analyze the deviation of the two processes that stems from the $2j + 1$ th additional choice in process P_{2j+1} . Here we observe the following. In each round of process P_{2j} there are three disjoint sets of agents, M_a, M_b , and M_u . The sets M_a and M_b are comprised of agents that sample at least $j + 1$ agents of the majority opinion a and the minority opinion b , respectively. All other agents are in M_u . The agents in M_a will adopt opinion a at the end of the round in both processes: the $j + 1$ samples of opinion a is larger than the winning margin in both processes, which is j in P_{2j} and $(2j + 1)/2$ in P_{2j+1} . Analogously, the agents in M_b will adopt opinion b in both processes. Finally, the interesting group are the M_u agents. These agents have sampled a tie in process P_{2j} , meaning they have sampled j agents with opinion a and another j agents with opinion b . This means, in process P_{2j} all agents in M_u adopt either opinion a or opinion b with probability $1/2$ each. In process P_{2j+1} , however,

the $2j + 1$ th sample makes the decision. (Recall that in a process P_{2j+1} with an odd number of samples and $k = 2$ opinions no ties are possible.) Therefore, in process P_{2j+1} all agents in M_u adopt opinion a with probability α and opinion b with probability $(1 - \alpha)$.

Summarizing, we have the following. Due to the coupling of P_{2j} with P_{2j+1} , all agents in M_a or M_b behave exactly the same in both processes. We use $Z_a = |M_a|$ and $Z_b = |M_b|$ to denote their respective numbers. (Observe that Z_a and Z_b are the same in P_{2j} and P_{2j+1} due to the coupling.) In the following, we condition on the event that $|M_u| = m_u$. For the agents in M_u , the outcome can be described by binomial random variables: let Z_u^{2j} in process P_{2j} and Z_u^{2j+1} in process P_{2j+1} be the numbers of agents in M_u that adopt opinion a . Then

$$Z_u^{2j} \sim \text{Bin}(m_u, 1/2) \quad \text{and} \quad Z_u^{2j+1} \sim \text{Bin}(m_u, \alpha)$$

with $\alpha \geq 1/2$. Irrespective of the value of m_u we observe from well-known properties of binomial distributions that Z_u^{2j} is stochastically dominated by Z_u^{2j+1} , and hence

$$X_{t+1}^{(2j)} = Z_a + Z_u^{2j} \prec Z_a + Z_u^{2j+1} = X_{t+1}^{(2j+1)}.$$

The proof for the dominance of P_{2j} over P_{2j-1} uses similar definitions and follows analogously, with exception that M_u represents the agents that are undecided after the first $2j - 2$ draws and that Z_u^{2j-1} and Z_u^{2j} follow the same binomial distribution $\text{Bin}(m_u, \alpha)$.

The only ingredient that is left to prove is the monotonicity within one specific process. Indeed, if an analogous result as Lemma 5 in the sequential model can be proven, the path coupling argument follows the same lines as in the previous section.

► **Lemma 11.** *We find for P_{2j} and P_{2j+1} in the gossip model the following. Let X_t denote the number of agents with majority opinion at time t . If $s > s'$, then for all $d \in \{0, 1, \dots, n\}$*

$$\Pr[X_{t+1} \geq d \mid X_t = s] \geq \Pr[X_{t+1} \geq d \mid X_t = s'].$$

Proof. As before, let M_a and M_b denote the sets of agents that sample at least $j + 1$ agents of the majority opinion a and the minority opinion b , respectively. If $s > s'$, the monotonicity of the binomial distribution yields

$$|M_a|_{X_t=s} \succeq |M_a|_{X_t=s'} \quad \text{and} \quad |M_b|_{X_t=s} \preceq |M_b|_{X_t=s'}.$$

Therefore, the lemma follows from Strassen's theorem. ◀

Now the path coupling follows analogously to the previous section. ◀

2.3 Analysis of 3-Majority

In this section we analyze 3-Majority in the sequential model. We start with an overview of the proof of Theorem 2. The proof consists of three parts. The first part follows along the lines of the proof by Condon et al. [21] for the related approximate majority process in tri-molecular chemical reaction networks. It shows that we preserve the initial majority (assuming a bias of $\sqrt{n \log n}$) and reach a bias of ϵn within $O(n \log n)$ time w.h.p. (Recall that the *bias* is defined as the difference of the numbers of agents supporting opinion a and opinion b .) The proof is based on the following result for gambler's ruin from [32].

► **Lemma 12** (Asymmetric one-dimensional random walk, [32, XIV.2], version from [21]). *If we run an arbitrarily long sequence of independent trials, each with success probability at least p , then the probability that the number of failures ever exceeds the number of successes by b is at most $\left(\frac{1-p}{p}\right)^b$.*

23:12 On the Hierarchy of Distributed Majority Protocols

In the second part we use a drift analysis based on [40] to show that we reach consensus on the initial majority opinion quickly once we have a bias of order $\Omega(n)$. The proof is based on a carefully conducted drift-analysis, where we use the following fairly recent result.

► **Theorem 13** (Special case of Theorem 18 of [40]). *Let $\{Y_t\}_{t \geq 0}$ be a sequence of non-negative random variables with a finite state space $S \subset \mathbb{R}_{\geq 0}$ such that $0 \in S$. Define*

$$s_{\min} = \min(S \setminus \{0\}) \quad \text{and} \quad T = \inf \{t \geq 0 \mid Y_0 = 0\}.$$

If $Y_0 = s_0$ and there is $\delta > 0$ (independent from t) such that for all $s \in S \setminus \{0\}$ and all $t \geq 0$ we have

$$\mathbb{E}[Y_t - Y_{t+1} \mid Y_t = s] \geq \delta s,$$

then, for all $r \geq 0$,

$$\Pr \left[T > \left\lceil \frac{r + \log(s_0/s_{\min})}{\delta} \right\rceil \right] \leq e^{-r}.$$

In the third part we again show that the analysis from [21] is applicable in our setting if we do not have an initial bias. All three parts together prove the first statement of our theorem. The second statement follows from part one together with part two.

Part 1. We start with the first part. We follow along the lines of [21] and use Lemma 12 to show the following statement.

► **Lemma 14.** *Let Δ_t be the additive bias at time t . With probability $1 - e^{-\Omega(\Delta_t^2/n)}$, the bias Δ_t does not drop below $\Delta_t/2$ and increases to $\min\{2\Delta_t, n\}$ within $2n$ time steps.*

Proof. Let X_t denote the number of agents with the majority opinion at time t and let $Y_t = n - X_t$ denote the number of agents with the minority opinion at time t . We analyze our process as a variant of gamblers' ruin and apply Lemma 12. We only consider *productive* steps in which the number of agents of a specific opinion changes. For $X_t \in (\frac{n}{2}, \frac{n}{2} + \varepsilon n)$ it holds that $\Pr[X_{t+1} \neq X_t] = \Omega(1)$ and hence conditioning on productive steps only increases the constants hidden in the asymptotic notation.

In each productive step, the success probability reads $p = \Pr[X_{t+1} > X_t \mid X_{t+1} \neq X_t]$ and the failure probability reads $1 - p = \Pr[X_{t+1} < X_t \mid X_{t+1} \neq X_t]$. Let $\Delta_t = X_t - \frac{n}{2}$ denote the bias at time t . We have for any Δ that

$$\frac{1-p}{p} = \frac{2(\frac{1}{2} - \frac{\Delta}{n}) - 3(\frac{1}{2} - \frac{\Delta}{n})^3 + (\frac{1}{2} - \frac{\Delta}{n})}{2(\frac{1}{2} - \frac{\Delta}{n})^4 - 5(\frac{1}{2} - \frac{\Delta}{n})^3 + 3(\frac{1}{2} - \frac{\Delta}{n})^2} < 1 - 16\frac{\Delta}{n}. \quad (1)$$

Unfortunately, the success probabilities vary over time as they depend on the bias. We proceed to bound the probabilities from below.

Let Δ_0 be the bias at time $t = 0$ and let \mathcal{R} denote the following event: during $2n$ productive steps we always have at least half of the initial bias, i.e., $\mathcal{R} = \{\forall 1 \leq i \leq 2n: \Delta_i \geq \Delta_0/2\}$. From Lemma 12 we get with $b = \Delta_0/2$ that

$$\Pr[\mathcal{R}] \geq 1 - e^{-\Omega(\Delta_0^2/n)}. \quad (2)$$

Similarly to [21], we couple the productive steps of the 3-Majority process with a biased random walk with (fixed) success probability $p > \frac{1}{2} + \frac{\Delta_0}{4n}$. As (1) is monotonously decreasing in Δ , the number of steps required by the biased random walk to increase the bias stochastically dominates the number of steps that 3-Majority requires. It follows from Chernoff bounds that the random walk reaches $2\Delta_0$ within $2n$ time steps with probability $1 - e^{-\Omega(\Delta_0^2/n)}$. Together with (2) the statement follows. ◀

We now use Lemma 14 and show that if there is a small bias of size $\sqrt{n \log n}$ then within $O(n \log n)$ rounds there will be a bias of size $\Omega(n)$ w.h.p.

► **Corollary 15.** *Assume $X_0 = \frac{n}{2} + \sqrt{n \log n}$. Then there is a time $t = O(n \log n)$ such that $X_t > \frac{1+\varepsilon}{2}n$ for some constant $\varepsilon > 0$ w.h.p. Moreover, the initial majority opinion is preserved.*

Proof. The proof follows by applying Lemma 14 $O(\log n)$ times. We remark that the initial majority opinion is preserved since the random walk modeling the bias never returns to zero. ◀

Part 2. We now show the second part, where we prove that the process converges within $O(n \log n)$ further steps once we have a bias of εn . Let Y_t denote the number of agents of the *minority* opinion at time t and assume that $Y_0 \leq \frac{n}{2} - \varepsilon n$. In a first step, we claim that the process will not improve the minority opinion severely if only $Cn \log n$ steps are conducted for some large constant C .

► **Lemma 16.** *Assume $Y_0 \leq \frac{n}{2} - \varepsilon n$. Then there is a time $t = O(n \log n)$ such that $Y_t = 0$ w.h.p. Moreover, $Y_{t'} \leq \frac{1-\varepsilon}{2}n$ for all $t' \leq t$.*

Proof. We start the proof by showing the following claim:

▷ **Claim.** $Y_{t'} \leq \frac{1-\varepsilon}{2}n$ for all $t' = O(n \log n)$ w.h.p.

This is an immediate consequence of the following coupling. Let R_t be the (unbiased) random walk on \mathbb{Z} . It is a well known fact that after T steps the random walk R_t has distance at most $O(\log^2 n \cdot \sqrt{n})$ from the origin w.h.p. By construction, $R_t \preceq Y_t$ and the claim follows.

We now calculate $\mathbb{E}[Y_t - Y_{t+1} \mid Y_t = s]$ for P_3 in the sequential model. Given $Y_t = s$, let $p_s(a, b)$ be the probability to increase the minority opinion by one and let $p_s(b, a)$ be the probability to decrease the minority opinion by one. Then,

$$\mathbb{E}[Y_t - Y_{t+1} \mid Y_t = s] = p_s(b, a) - p_s(a, b).$$

We observe

$$p_s(a, b) = \frac{n-s}{n} \Pr\left[\text{Bin}\left(3, \frac{s}{n}\right) \geq 2\right], \quad p_s(b, a) = \frac{s}{n} \Pr\left[\text{Bin}\left(3, \frac{s}{n}\right) \leq 1\right],$$

and therefore

$$\frac{p_s(b, a) - p_s(a, b)}{b} = \frac{2s^2 - 3sn + n^2}{n^3}.$$

We define $\delta_s = \frac{p_s(b, a) - p_s(a, b)}{s}$ and observe

$$\delta_s - \delta_{s-1} = \frac{4s - 3n - 2}{n^3} < 0 \quad \text{if } s \leq 0.75n.$$

Therefore, since $s \leq \frac{1-\varepsilon}{2}n$ by the previous claim, δ_s is monotonously decreasing in s . Furthermore,

$$\delta_{\frac{1-\varepsilon}{2}n} = \frac{(1 + \varepsilon/2)\varepsilon}{4n} \quad \text{and} \quad \delta_1 = n^{-1} + O(n^{-2}).$$

Thus, we apply Theorem 13 with

$$\delta = \frac{(1 + \varepsilon/2)\varepsilon}{4n}, \quad s_0 = \left(\frac{1}{2} - \varepsilon\right)n, \quad r = \log n, \quad \text{and} \quad s_{\min} = 1$$

and the statement follows. ◀

Part 3. It remains to show the third part of the proof. We observe the following. We use the same *checkpoint states* g_j as in [21] where $g_0 = 0$ and $g_j = 2^{j+3} \cdot \sqrt{n}$. A checkpoint state can be intuitively described as follows. We let P_3 run in packages of $2n$ productive update steps and monitor the majority opinion. Suppose we are in checkpoint state $g_1 = 8\sqrt{n}$. After $2n$ productive updates, Lemma 14 guarantees that with probability at least $1 - 1/(2^j + O(1))$ the majority opinion exceeds g_2 . Now we interpret this process as a (biased) random walk on the checkpoint states $\{g_j\}_j$ in which every conducted step consists of $2n$ productive update steps of 3-Majority. Analogously to the analysis of [21], it holds that

1. the transition between checkpoint states g_0 and g_1 has probability $\Omega(1)$, and
2. for $j \geq 1$ the transition between checkpoint states g_j and $g_j + 1$ has probability at least $1 - 1/(2^j + O(1))$.

As in [21], the first statement follows from a coupling with an unbiased random walk, and the second statement follows from Lemma 14. It follows from the analysis in [21, Section 3.2] that 3-Majority reaches a bias of $\sqrt{n \log n}$ within $O(n \log n)$ time. This proof is based on a careful trade-off between the geometrically increasing success probability $1 - 1/(2^j + O(1))$ to get into the next checkpoint state and the number of trials that are necessary to indeed reach the next state instead of falling back.

With all three parts, we are now ready to put everything together and prove Theorem 2.

Proof of Theorem 2. Assume there is no bias. From the analysis in [21] we obtain (see above) that we reach a bias of size $\sqrt{n \log n}$ within $O(n \log n)$ time w.h.p. From Corollary 15 we obtain that within further $O(n \log n)$ time the bias is amplified to ϵn for some constant $\epsilon > 0$ w.h.p. Finally, from the drift analysis in Lemma 16 we get that we converge in further $O(n \log n)$ time once we have a constant-factor bias w.h.p. Together, this shows the first part of the theorem.

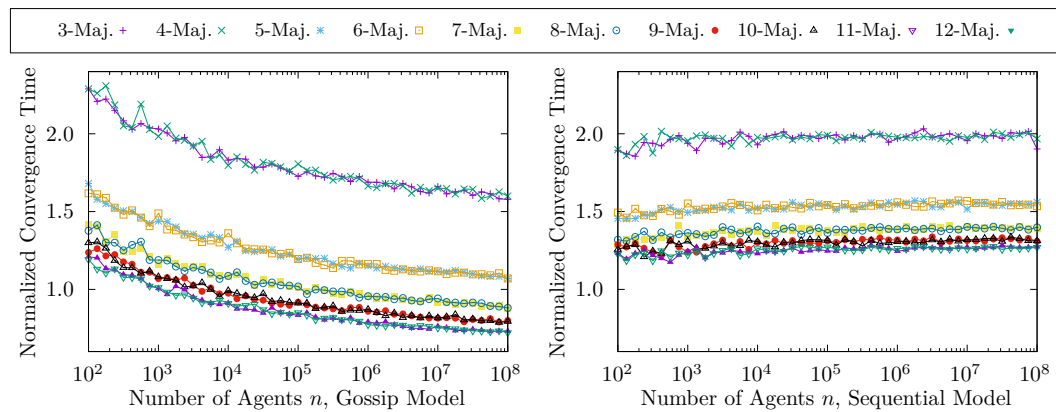
The second part of the theorem follows from Lemma 14 and Lemma 16, where we observe that the initial majority opinion is preserved w.h.p. This concludes the proof. ◀

3 Empirical Analysis

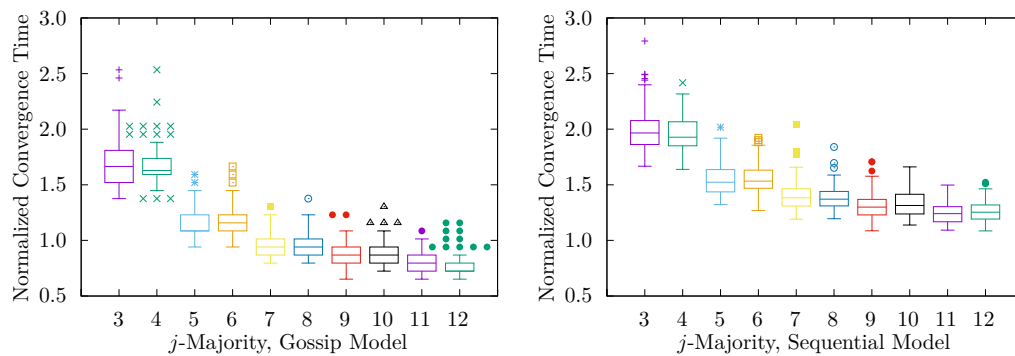
In this section we present simulation results to support our theoretical findings. Our simulation software is implemented in the C++ programming language. As a source of randomness it uses the Mersenne Twister `mt19937_64` provided by the C++11 `<random>` library. Our simulations have been carried out on machines with two Intel(R) Xeon(R) E5-2630 v4 CPUs and 128 GiB of memory each running the Linux 5.13 kernel. The simulation software and all required tools to reproduce our plots are publicly available in our Github repository.

In Figure 3 we plot the required number of rounds until j -Majority converges when each opinion is initially supported $n/2$ agents. The data show the average convergence time over 100 independent simulation runs for $j = 3, \dots, 12$. The number of agents n is shown on the x -axis, and the normalized convergence time is shown on the y -axis. The left plot shows the data for the gossip model, where the normalization means that the required number of rounds is divided by $\log n$. The right plot shows the data for the sequential model, where the normalization means that the required number of interactions is divided by $n \log n$.

Our empirical data confirm our theoretical findings. In particular, we observe that the processes exhibit a running time of $\Theta(\log n)$ rounds (gossip model) or $\Theta(n \log n)$ interactions (sequential model) for the values of j we consider. Furthermore, we clearly see that $\mathbb{E}[T_{2j+2}(X_0)] = \mathbb{E}[T_{2j+1}(X_0)]$ (i.e., 3-Majority converges as quickly as 4-Majority, 5-Majority converges as quickly as 6-Majority, and so on) and $\mathbb{E}[T_{2j+1}(X_0)] \leq \mathbb{E}[T_{2j}(X_0)]$ (i.e.,



■ **Figure 3** Average convergence time of j -Majority without initial bias and $j = 3, \dots, 12$ normalized over $\log n$ (gossip model) or $n \log n$ (sequential model). Each data point shows the average of 100 independent runs. The left plot shows the gossip model and the right plot shows the sequential model.



■ **Figure 4** Boxplots for the normalized convergence time of j -Majority without initial bias. The plots show details of the distribution of the same data as in Figure 3 for $n = 10^6$.

5-Majority is faster than 4-Majority, 7-Majority is faster than 6-Majority, and so on). This empirically confirms our results from Theorem 1 for both models, and it shows that the known results from the gossip model for 3-Majority [35] carry over to the sequential model as predicted in Theorem 2.

In the left plot in Figure 3 for the gossip model we additionally observe that the required number of rounds to reach consensus is slightly larger for smaller values of n . This appears to be a consequence of the discrete rounds in the synchronous model: the observed deviation scales as $O(1/\log n)$, which is of the same size as the rounding error that arises when reporting the running time in discrete rounds of n interactions each.

Finally, in Figure 4 we show additional detail for the distribution of the convergence times of the j -Majority processes with $n = 10^6$ and $j = 3, \dots, 12$. Our boxplots show that the running times are strongly concentrated around the mean, and the constants hidden in the asymptotic analysis are small: the running time is less than $3 \log n$ rounds in the gossip model and less than $3n \log n$ interactions in the sequential model. The small constants hint at the practical applicability of the simple 3-Majority process.

4 Conclusions and Open Problems

We analyze the family of j -Majority processes in two communication models with parallel and sequential activations. In both models our results affirmatively answer an open question from [12] for the case of two opinions and prove the existence of a *hierarchy*: our results show the stochastic dominance of the convergence time of the $(j + 1)$ -Majority process over the j -Majority process. For 3-Majority in the sequential model we show an asymptotically optimal bound of $O(n \log n)$ sequential activations. This matches the well-known bounds for the corresponding process in the gossip model.

An open question is whether a similar hierarchy exists for *lazy* processes where agents keep their previous opinion if there is a tie among the sampled opinions. A coupling between 3-Majority and the (lazy) TWOCHOICES process was analyzed in [12]. However, their general framework cannot be adapted to lazy processes for larger value of j : their analysis requires so-called *AC-Processes* in which the next state of an agent depends only on the global opinion distribution but not on the agent's current state. This is obviously not the case for lazy processes. Note that our analysis also cannot be applied to lazy processes directly: Lemmas 8 and 10 do not hold for lazy processes.

Another interesting open question considers the *communication complexity* of a protocol instead which counts the number of interactions. Note that in j -Majority each activated agent interacts with j agents. It would be interesting to rigorously analyze the trade-off between the convergence time and the communication complexity.

Finally, the most interesting open question is whether similar results can be shown for more than two opinions. Unfortunately, our majoritization-based approach does not generalize to $k > 2$. The main reason is that natural monotonicity properties do not hold: the probability to increase the majority opinion does not only depend on the size of the majority opinion itself but instead on the entire opinion distribution. This aligns well with a conjecture from [12] that states that counterexamples exist for any majoritization attempt that uses a total order on opinion state vectors. We believe that in order to show a hierarchy of majority protocols for more than two opinions different techniques will be needed.

References

- 1 Dan Alistarh, James Aspnes, David Eisenstat, Rati Gelashvili, and Ronald L. Rivest. Time-Space Trade-offs in Population Protocols. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017*, pages 2560–2579. SIAM, 2017. doi:10.1137/1.9781611974782.169.
- 2 Dan Alistarh, James Aspnes, and Rati Gelashvili. Space-Optimal Majority in Population Protocols. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*, pages 2221–2239. SIAM, 2018. doi:10.1137/1.9781611975031.144.
- 3 Dan Alistarh, Rati Gelashvili, and Milan Vojnovic. Fast and Exact Majority in Population Protocols. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC 2015*, pages 47–56. ACM, 2015. doi:10.1145/2767386.2767429.
- 4 Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Comput.*, 18(4):235–253, 2006. doi:10.1007/s00446-005-0138-3.
- 5 Dana Angluin, James Aspnes, and David Eisenstat. A simple population protocol for fast robust approximate majority. *Distributed Comput.*, 21(2):87–102, 2008. doi:10.1007/s00446-008-0059-z.
- 6 Gregor Bankhamer, Petra Berenbrink, Felix Biermeier, Robert Elsässer, Hamed Hosseinpour, Dominik Kaaser, and Peter Kling. Fast Consensus via the Unconstrained Undecided State Dynamics. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022*, pages 3417–3429. SIAM, 2022. doi:10.1137/1.9781611977073.135.

- 7 Gregor Bankhamer, Robert Elsässer, Dominik Kaaser, and Matjaz Krnc. Positive Aging Admits Fast Asynchronous Plurality Consensus. In *PODC '20: ACM Symposium on Principles of Distributed Computing*, pages 385–394. ACM, 2020. doi:10.1145/3382734.3406506.
- 8 Luca Becchetti, Andrea E. F. Clementi, and Emanuele Natale. Consensus Dynamics: An Overview. *SIGACT News*, 51(1):58–104, 2020. doi:10.1145/3388392.3388403.
- 9 Luca Becchetti, Andrea E. F. Clementi, Emanuele Natale, Francesco Pasquale, and Riccardo Silvestri. Plurality Consensus in the Gossip Model. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015*, pages 371–390. SIAM, 2015. doi:10.1137/1.9781611973730.27.
- 10 Luca Becchetti, Andrea E. F. Clementi, Emanuele Natale, Francesco Pasquale, Riccardo Silvestri, and Luca Trevisan. Simple dynamics for plurality consensus. *Distributed Comput.*, 30(4):293–306, 2017. doi:10.1007/s00446-016-0289-4.
- 11 Stav Ben-Nun, Tsvi Kopelowitz, Matan Kraus, and Ely Porat. An $O(\log^{3/2} n)$ Parallel Time Population Protocol for Majority with $O(\log n)$ States. In *PODC '20: ACM Symposium on Principles of Distributed Computing*, pages 191–199. ACM, 2020. doi:10.1145/3382734.3405747.
- 12 Petra Berenbrink, Andrea E. F. Clementi, Robert Elsässer, Peter Kling, Frederik Mallmann-Trenn, and Emanuele Natale. Ignore or Comply?: On Breaking Symmetry in Consensus. In *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC 2017*, pages 335–344. ACM, 2017. doi:10.1145/3087801.3087817.
- 13 Petra Berenbrink, Amin Coja-Oghlan, Oliver Gebhard, Max Hahn-Klimroth, Dominik Kaaser, and Malin Rau. On the Hierarchy of Distributed Majority Protocols. *CoRR*, abs/2205.08203, 2022. doi:10.48550/arXiv.2205.08203.
- 14 Petra Berenbrink, Robert Elsässer, Tom Friedetzky, Dominik Kaaser, Peter Kling, and Tomasz Radzik. A Population Protocol for Exact Majority with $O(\log^{5/3} n)$ Stabilization Time and $\Theta(\log n)$ States. In *32nd International Symposium on Distributed Computing, DISC 2018*, pages 10:1–10:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.DISC.2018.10.
- 15 Petra Berenbrink, Robert Elsässer, Tom Friedetzky, Dominik Kaaser, Peter Kling, and Tomasz Radzik. Time-space trade-offs in population protocols for the majority problem. *Distributed Comput.*, 34(2):91–111, 2021. doi:10.1007/s00446-020-00385-0.
- 16 Petra Berenbrink, Tom Friedetzky, George Giakkoupis, and Peter Kling. Efficient Plurality Consensus, Or: the Benefits of Cleaning up from Time to Time. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016*, pages 136:1–136:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.ICALP.2016.136.
- 17 Petra Berenbrink, George Giakkoupis, Anne-Marie Kermarrec, and Frederik Mallmann-Trenn. Bounds on the Voter Model in Dynamic Networks. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016*, pages 146:1–146:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.ICALP.2016.146.
- 18 Andreas Bilke, Colin Cooper, Robert Elsässer, and Tomasz Radzik. Brief Announcement: Population Protocols for Leader Election and Exact Majority with $O(\log^2 n)$ States and $O(\log^2 n)$ Convergence Time. In *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC 2017*, pages 451–453. ACM, 2017. doi:10.1145/3087801.3087858.
- 19 Keren Censor-Hillel, Bernhard Haeupler, Jonathan A. Kelner, and Petar Maymounkov. Global computation in a poorly connected world: fast rumor spreading with no dependence on conductance. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012*, pages 961–970. ACM, 2012. doi:10.1145/2213977.2214064.
- 20 Andrea E. F. Clementi, Mohsen Ghaffari, Luciano Gualà, Emanuele Natale, Francesco Pasquale, and Giacomo Scornavacca. A Tight Analysis of the Parallel Undecided-State Dynamics with Two Colors. In *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2018*, pages 28:1–28:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.MFCS.2018.28.

- 21 Anne Condon, Monir Hajiaghayi, David G. Kirkpatrick, and Ján Manuch. Approximate majority analyses using tri-molecular chemical reaction networks. *Nat. Comput.*, 19(1):249–270, 2020. doi:10.1007/s11047-019-09756-4.
- 22 Colin Cooper, Robert Elsässer, Hirotaka Ono, and Tomasz Radzik. Coalescing random walks and voting on graphs. In *ACM Symposium on Principles of Distributed Computing, PODC '12*, pages 47–56. ACM, 2012. doi:10.1145/2332432.2332440.
- 23 Colin Cooper, Robert Elsässer, and Tomasz Radzik. The Power of Two Choices in Distributed Voting. In *Automata, Languages, and Programming – 41st International Colloquium, ICALP 2014*, pages 435–446. Springer, 2014. doi:10.1007/978-3-662-43951-7_37.
- 24 Colin Cooper, Robert Elsässer, Tomasz Radzik, Nicolas Rivera, and Takeharu Shiraga. Fast Consensus for Voting on General Expander Graphs. In *Distributed Computing – 29th International Symposium, DISC 2015*, pages 248–262. Springer, 2015. doi:10.1007/978-3-662-48653-5_17.
- 25 Colin Cooper, Tomasz Radzik, Nicolas Rivera, and Takeharu Shiraga. Fast Plurality Consensus in Regular Expanders. In *31st International Symposium on Distributed Computing, DISC 2017*, pages 13:1–13:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.DISC.2017.13.
- 26 Benjamin Doerr, Leslie Ann Goldberg, Lorenz Minder, Thomas Sauerwald, and Christian Scheideler. Stabilizing consensus with the power of two choices. In *SPAA 2011: Proceedings of the 23rd Annual ACM Symposium on Parallelism in Algorithms and Architectures*, pages 149–158. ACM, 2011. doi:10.1145/1989493.1989516.
- 27 David Doty and Mahsa Eftekhari. Efficient Size Estimation and Impossibility of Termination in Uniform Dense Population Protocols. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019*, pages 34–42. ACM, 2019. doi:10.1145/3293611.3331627.
- 28 David Doty, Mahsa Eftekhari, Leszek Gasieniec, Eric E. Severson, Przemyslaw Uznanski, and Grzegorz Stachowiak. A time and space optimal stable population protocol solving exact majority. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021*, pages 1044–1055. IEEE, 2021. doi:10.1109/FOCS52979.2021.00104.
- 29 David Doty and David Soloveichik. Stable leader election in population protocols requires linear time. *Distributed Comput.*, 31(4):257–271, 2018. doi:10.1007/s00446-016-0281-z.
- 30 Moez Draief and Milan Vojnovic. Convergence Speed of Binary Interval Consensus. *SIAM J. Control. Optim.*, 50(3):1087–1109, 2012. doi:10.1137/110823018.
- 31 Robert Elsässer, Tom Friedetzky, Dominik Kaaser, Frederik Mallmann-Trenn, and Horst Trinker. Brief Announcement: Rapid Asynchronous Plurality Consensus. In *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC 2017*, pages 363–365. ACM, 2017. doi:10.1145/3087801.3087860.
- 32 William Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. Wiley, 3rd edition, 2008.
- 33 Pierre Fraigniaud and Emanuele Natale. Noisy rumor spreading and plurality consensus. *Distributed Comput.*, 32(4):257–276, 2019. doi:10.1007/s00446-018-0335-5.
- 34 Leszek Gasieniec, David D. Hamilton, Russell Martin, Paul G. Spirakis, and Grzegorz Stachowiak. Deterministic Population Protocols for Exact Majority and Plurality. In *20th International Conference on Principles of Distributed Systems, OPODIS 2016*, pages 14:1–14:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.OPODIS.2016.14.
- 35 Mohsen Ghaffari and Johannes Lengler. 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*, pages 305–313. ACM, 2018. doi:10.1145/3212734.3212738.
- 36 Mohsen Ghaffari and Merav Parter. A Polylogarithmic Gossip Algorithm for Plurality Consensus. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing, PODC 2016*, pages 117–126. ACM, 2016. doi:10.1145/2933057.2933097.

- 37 Yehuda Hassin and David Peleg. Distributed Probabilistic Polling and Applications to Proportionate Agreement. *Inf. Comput.*, 171(2):248–268, 2001. doi:10.1006/inco.2001.3088.
- 38 Varun Kanade, Frederik Mallmann-Trenn, and Thomas Sauerwald. 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*, pages 956–965. SIAM, 2019. doi:10.1137/1.9781611975482.59.
- 39 Adrian Kosowski and Przemyslaw Uznanski. Brief Announcement: Population Protocols Are Fast. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing, PODC 2018*, pages 475–477. ACM, 2018. doi:10.1145/3212734.3212788.
- 40 Johannes Lengler. Drift Analysis. In Benjamin Doerr and Frank Neumann, editors, *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization*, pages 89–131. Springer, 2020. doi:10.1007/978-3-030-29414-4_2.
- 41 George B. Mertzios, Sotiris E. Nikolettseas, Christoforos L. Raptopoulos, and Paul G. Spirakis. Determining Majority in Networks with Local Interactions and Very Small Local Memory. In *Automata, Languages, and Programming – 41st International Colloquium, ICALP 2014*, pages 871–882. Springer, 2014. doi:10.1007/978-3-662-43948-7_72.
- 42 Yves Mocquard, Emmanuelle Anceaume, James Aspnes, Yann Busnel, and Bruno Sericola. Counting with Population Protocols. In *14th IEEE International Symposium on Network Computing and Applications, NCA 2015*, pages 35–42. IEEE Computer Society, 2015. doi:10.1109/NCA.2015.35.
- 43 Toshio Nakata, Hiroshi Imahayashi, and Masafumi Yamashita. A probabilistic local majority polling game on weighted directed graphs with an application to the distributed agreement problem. *Networks*, 35(4):266–273, 2000. doi:10.1002/1097-0037(200007)35:4<266::AID-NET5>3.0.CO;2-4.
- 44 Emanuele Natale and Iliad Ramezani. On the Necessary Memory to Compute the Plurality in Multi-agent Systems. In *Algorithms and Complexity – 11th International Conference, CIAC 2019*, pages 323–338. Springer, 2019. doi:10.1007/978-3-030-17402-6_27.
- 45 Grant Schoenebeck and Fang-Yi Yu. Consensus of Interacting Particle Systems on Erdős-Rényi Graphs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*, pages 1945–1964. SIAM, 2018. doi:10.1137/1.9781611975031.127.
- 46 Volker Strassen. The existence of probability measures with given marginals. *The Annals of Mathematical Statistics*, 36(2):423–439, 1965.