Simple Dynamics for Plurality Consensus*

L. Becchetti[1], A. Clementi[2], E. Natale[1], F. Pasquale[2], R. Silvestri[1], and L. Trevisan[3]

[1] *Sapienza* Università di Roma, `becchett@dis.uniroma1.it`, `natale@di.uniroma1.it`, `silvestri@di.uniroma1.it`

[2] Università *Tor Vergata* di Roma, `clementi@mat.uniroma2.it`, `pasquale@mat.uniroma2.it`,

[3] Stanford University, `trevisan@stanford.edu`

July 28, 2015

Abstract

We study a *Plurality-Consensus* process in which each of n anonymous agents of a communication network initially supports an opinion (a color chosen from a finite set $[k]$). Then, in every (synchronous) round, each agent can revise his color according to the opinions currently held by a random sample of his neighbors. It is assumed that the initial color configuration exhibits a sufficiently large *bias* s towards a fixed plurality color, that is, the number of nodes supporting the plurality color exceeds the number of nodes supporting any other color by s additional nodes. The goal is having the process to converge to the *stable* configuration in which all nodes support the initial plurality. We consider a basic model in which the network is a clique and the update rule (called here the *3-majority dynamics*) of the process is the following: each agent looks at the colors of three random neighbors and then applies the majority rule (breaking ties uniformly).

We prove that the process converges in time $\mathcal{O}(\min\{k, (n/\log n)^{1/3}\} \log n)$ with high probability, provided that $s \geqslant c\sqrt{\min\{2k, (n/\log n)^{1/3}\}\, n \log n}$. We then prove that our upper bound above is tight as long as $k \leqslant (n/\log n)^{1/4}$. This fact implies an exponential time-gap between the plurality-consensus process and the *median* process studied by Doerr et al. in [ACM SPAA'11].

A natural question is whether looking at more (than three) random neighbors can significantly speed up the process. We provide a negative answer to this question: In particular, we show that samples of polylogarithmic size can speed up the process by a polylogarithmic factor only.

Keywords: Plurality Consensus; Distributed Randomized Algorithms; Markov Chains.

*Partially supported by Italian MIUR-PRIN 2010-11 Project *ARS TechnoMedia* and the EU FET Project *MULTIPLEX* 317532. A preliminary version of this paper appeared in the *Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'14)*

1 Introduction

We consider a communication network in which each of n anonymous nodes supports an initial opinion (a color chosen from a finite set $[k]$). In the *Plurality Consensus* problem, it is assumed that the initial (color) configuration has a sufficiently large *bias* s towards a fixed color $m \in [k]$ - that is, the number c_m of nodes supporting the plurality color (in short, the *initial plurality size*) exceeds the number c_j of nodes supporting any other color j by an additive value s - and the goal is to design an efficient fully-distributed protocol that lets the network converge to the *plurality consensus*, i.e., to the monochromatic configuration in which all nodes support the plurality color.

Reaching plurality consensus in a distributed system is a fundamental problem arising in several areas such as Distributed Computing [7, 18], Communication Networks [19], and Social Networks [6, 16, 15]. Inspired by some recent works analyzing simple updating-rules (called *dynamics*) for this problem [1, 7], we study a discrete-time, synchronous process in which, in every round, each of the n anonymous nodes revises his color according to a (small) random sample of neighbors. We consider one of the simplest models, in which the network is a clique, and the updating rule, called here 3*-majority dynamics*, is the following simple one: Each node samples at random three neighbors, and picks the majority color among them (breaking ties uniformly at random). We remark that looking at less than three random neighbors would yield a coloring process that may converge to a *minority* color with constant probability even for $k = 2$ and large initial bias (i.e. $s = \Theta(n)$).

In [7], a tight analysis of a 3-neighbor dynamics for the *median* problem on the clique was presented: the goal there is to converge to a stable configuration where all nodes support a value which is a good approximation of the *median* of the initial configuration. It turns out that, in the binary case (i.e $k = 2$), the median problem is equivalent to plurality consensus and the 3-input dynamics for the median is equivalent to the 3-majority dynamics: As a result, they obtain, for any bias $s \geqslant c\sqrt{n \log n}$ for some constant $c > 0$, an optimal bound $\Theta(\log n)$ on the convergence time of the 3-majority dynamics for the binary case of the problem considered in this paper.

However, for any $k \geqslant 3$, it is easy to see that the two problems above differ significantly (in particular, the median may be very different from the plurality) and thus, the two dynamics are different from each other as well. Moreover, the analysis in [7] - strongly based on the properties of the median function - cannot be adapted to bound the convergence time of the 3-majority dynamics. The role of parameter $k = k(n)$ in the convergence time of this dynamics is currently unknown and, more generally, the existence of efficient dynamics reaching plurality consensus for $k \geqslant 3$ is left as an important open issue in [2, 7, 3].

Our contribution. We present a new analysis of the 3-majority dynamics in the general case (i.e. for any $k \in [n]$). Our analysis shows that, with high probability (in short, *w.h.p.*[1]), the process converges to plurality consensus within time $\mathcal{O}\left(\min\{k, (n/\log n)^{1/3}\} \log n\right)$, provided that the initial bias is $s \geqslant c\sqrt{\min\{2k, (n/\log n)^{1/3}\}\, n \log n}$, for some constant $c > 0$.

Our proof technique is accurate enough to get another interesting form of the above upper bound that does not depend on k. Indeed, when the initial plurality size c_m is larger than $n/\lambda(n)$ for any function $\lambda(n)$ such that $3 \leqslant \lambda(n) < \sqrt{n}$ and $s \geqslant \sqrt{\lambda(n)\, n \log n}$, then the process converges in time $\mathcal{O}\left(\lambda(n) \log n\right)$ w.h.p., no matter how large k is. Hence, when $c_m \geqslant n/\text{polylog}(n)$ and $s \geqslant \sqrt{n\, \text{polylog}(n)}$, the convergence time is polylogaritmic.

We then show that our upper bound is tight for a wide range of the input parameters. When $k \leqslant (n/\log n)^{1/4}$, we prove a lower bound $\Omega(k \log n)$ on the convergence time of the 3-majority dynamics starting from some configurations with bias $s \leqslant (n/k)^{1-\epsilon}$, for an arbitrarily small constant $\epsilon > 0$. Observe that this range largely includes the initial bias required by our upper bound when $k \leqslant (n/\log n)^{1/4}$. So, the *linear-in-k* dependence of the convergence time cannot be removed for a wide range of the parameter k.

Our analysis also provides a clear picture of the 3-majority dynamic process. Informally speaking, the larger the initial value of c_m is (w.r.t. n), the smaller the required initial bias s and the faster the convergence time are. On the other hand, our lower-bound argument shows, as a by-product, that the initial plurality size c_m needs $\Omega(k \log n)$ rounds just to increase from $n/k + o(n/k)$ to $2\, n/k$. Another natural issue is to analyze the process under weaker assumptions on the initial bias. We show that there are initial configurations with bias $s = O(\sqrt{kn})$ for which the bias decreases in a single round with

[1]We say that a family of events $\{\mathcal{E}_n\}_n$ holds w.h.p. if a positive constant c exists such that $\mathbf{P}\left(\mathcal{E}_n\right) \geqslant 1 - n^{-c}$ for sufficiently large n

constant probability. This implies that under initial imbalances of this magnitude, it seems unlikely that one can prove upper bounds similar to ours above, at least with high probability.

We then prove a general negative result: Under the distributed model we consider, no dynamics with at most 3 inputs (other than 3-majority) converges w.h.p. to plurality consensus starting from any initial configuration with bias $s = o(n)$. In other words, not only can we not design a 3-input dynamics that achieves convergence to plurality consensus in $o(k \log n)$ rounds, but the 3-majority dynamics is the only one that eventually achieves this goal at all, no matter how long the process takes. Rather interestingly, by comparing the $\mathcal{O}(\log n)$ bound for the median [7] to our negative results for the plurality on the same distributed model, we get an exponential time-gap between the task of computing the median and the one of computing plurality (this happens for instance when $k = n^a$, for any constant $0 < a < 1/4$).

A natural question suggested by our findings is whether (slightly) larger random samples of nodes' neighborhoods might lead to significant improvements in convergence time to plurality consensus. We provide a negative answer to this question. To this purpose, we consider h-plurality, i.e., the natural generalization of the 3-majority dynamics in which every node, in each round, updates his color according to the plurality of the colors supported by h randomly sampled neighbors. We prove a lower bound $\Omega\left(k/h^2\right)$ on the convergence time of the h-plurality dynamics, for integers k and h such that $k/h = \mathcal{O}\left(n^{1/4-\epsilon}\right)$, with ϵ an arbitrarily-small positive constant. We emphasize that scalable and efficient protocols must yield low communication complexity and small node congestion in every round. These properties are guaranteed by the h-plurality dynamics only when h is small, say $h = \mathcal{O}(\text{polylog}(n))$: In this case, our lower bound implies that the resulting speed-up is only polylogarithmic with respect to the 3-majority dynamics.

One motivation for adopting dynamics in reaching (*simple*) consensus[2] (such as the median dynamics shown in [7]) lies in their provably-good *self-stabilizing* properties against *dynamic adversary corruptions*: It turns out that the 3-majority dynamics has good self-stabilizing properties for the *plurality consensus* problem. More formally, a T-bounded adversary knows the state of every node at the end of each round and, based on this knowledge, he can corrupt the color of up to T nodes in an arbitrary way, just before the next round begins. In this case, the goal is to achieve an almost-stable phase where all but at most $\mathcal{O}(T)$ nodes agree on the plurality value. This "almost-stability" phase must have poly(n) length, with high probability. Our analysis implicitly shows that the 3-majority dynamics guarantees the self-stabilization property for plurality consensus for any k and for $T = o(s/k)$ if the initial bias is $s \geqslant c\sqrt{\min\{2k, (n/\log n)^{1/3}\}\, n \log n}$, for some constant $c > 0$.

Related work. The plurality consensus problem arises in several applications such as distributed database management, where data redundancy or replication and majority rules are used to manage the presence of unknown faulty processors [7, 18]. The goal here is to converge to the version of the data supported by the plurality of the initial distributed copies (it is reasonable that a sufficiently strong plurality of the nodes are not faulty and thus possess the correct data). Another application is distributed item ranking, in which every node initially selects some item and the goal is to agree on the most popular item according to the initial plurality opinion [19]. Further applications of majority updating rules in networks can be found in [10, 18].

The results most related to our contribution are those in [7] which have been already discussed above. Several variants of binary majority consensus have been studied in different distributed models [2, 16].

As for the *population model*, where there is only one random node-pair interaction per round (so the dynamics are strictly sequential), the binary case on the clique has been analyzed in [2] and their generalization to the multivalued case ($k \geqslant 3$) does not converge to plurality even starting with a large bias $s = \Theta(n)$.

The polling rule (a somewhat sequential-interaction version of the 1-majority dynamics) has been extensively studied on several classes of graphs (see [18]).

More expensive and complex protocols have been considered in order to speed up the process. For instance, in [12], a protocol for the sequential-interaction model is presented that requires $\Theta(\log n)$ memory per node and converges in time $\mathcal{O}(n^7)$. Other protocols for the sequential-interaction model have been analyzed in [5, 13] (with no time bound).

In [3, 8, 19], the *undecided-state* protocol on the continuous-time population model is proved to converge in $\mathcal{O}(n \log n)$ expected time only for $k = \Theta(1)$ and $s = \Theta(n)$: Even assuming such strong

[2] In the (simple) consensus problem the goal is to reach any stable monochromatic configuration (any color is accepted) starting from any initial configuration.

restrictions, the bound does not hold in "high probability" and, moreover, their analysis, based on real-valued differential-equations, do not work for the discrete-time parallel model considered in this paper. The simple rule of the *undecided-state* [2, 19] is to "add" one extra state to somewhat account for the "previous" opinion supported by an agent.

In a recent work [4] (appeared after the conference version of this paper), the undecided-state protocol has been analyzed on the discrete-time parallel model for any $k = \mathcal{O}((n/\log n)^{1/3})$ and for initial configurations $\mathbf{c} = (c_1, \ldots, c_k)$ such that the (multiplicative) bias is $c_m/c_j = \Omega(1)$. There, it is shown that this dynamics has a convergence time which is w.h.p. linear in the *monochromatic* distance of the initial configuration $\mathbf{c}$. The monochromatic distance of a configuration $\mathbf{c} = (c_1, \ldots, c_k)$ is defined as

$$\sum_{i=1}^{k} \left(\frac{c_i}{c_m}\right)^2 .$$

It turns out that there are initial configurations (in particular, those having "almost all" nodes supporting only a polylogarithmic number of colors) from which the undecided-state protocol is exponentially faster than the 3-majority. On the other hand, in addition to the above condition on the multiplicative bias, we note that, the undecided-state protocol may fail to reach consensus when $k = \omega(\sqrt{n})$.

Finally, protocols for specific network topologies and some "social-based" communities have been studied in [1, 8, 15, 19].

Roadmap. Section 2 formalizes the basic concepts and gives some preliminary results. Section 3 is devoted to the proofs of the upper bounds on the convergence time of the 3-majority dynamics. In Section 4, the lower bounds for the studied dynamics are described. Section 5 discusses some interesting open questions such as the tightness of the initial bias.

2 Preliminaries

A *(k-color) configuration* (for short *k-cd*) is any k-tuple $\mathbf{c} = (c_1, \ldots, c_k)$ such that c_js are non negative integers and $\sum_{j=1,\ldots,k} c_j = n$. In what follows, we will always assume wlog $c_1 \geqslant c_2 \geqslant \cdots \geqslant c_k$. So c_1 is the *plurality color* and $s(\mathbf{c}) = c_1 - c_2$ is the *bias* of $\mathbf{c}$.

The 3-majority protocol works as follows:

> *At every round, every node picks three nodes uniformly at random (including itself and with repetitions) and recolors itself according to the majority of the colors it sees. If it sees three different colors, it chooses the first one.*

Clearly, in the case of three different colors, choosing the second or the third one would not make any difference. The same holds even if the choice would be uniformly at random among the three colors.

For any round t and for any $j \in [k]$, let $C_j^{(t)}$ be the r.v. counting the number of nodes colored j at round t and let $\mathbf{C}^{(t)} = (C_1^{(t)}, \ldots, C_k^{(t)})$ denote the random variable indicating the k-cd at time t of the execution of the 3-majority protocol.

For every $j \in [k]$ let $\mu_j(\mathbf{c})$ be the expected number of nodes with color j at the next round when the current k-cd is $\mathbf{c}$, i.e. $\mu_j(\mathbf{c}) = \mathbb{E}\left[C_j^{(t+1)} \,\middle|\, \mathbf{C}^{(t)} = \mathbf{c}\right]$.

Lemma 2.1 (Next expected coloring) *For any k-cd $\mathbf{c}$ and for every color $j \in [k]$, it holds that*

$$\mu_j(\mathbf{c}) = c_j \left[1 + \frac{1}{n^2}\left(n c_j - \sum_{h \in [k]} c_h^2\right)\right].$$

Proof. According to the 3-majority protocol, a node i gets color j if it chooses three times color j, or if it chooses two times j and one time a different color, or if it chooses the first time color j and then, the second and third time, two different distinct colors. Hence if we name name $X_{i,j}^{(t)}$ the indicator random veriable of the event "Node i gets color j at time t", we have that

$$P\left(X_{i,j}^{(t+1)} = 1 \,|\, \mathbf{C}^{(t)} = \mathbf{c}\right) = \left(\frac{c_j}{n}\right)^3 + 3\left(\frac{c_j}{n}\right)^2\left(\frac{n-c_j}{n}\right) + \left(\frac{c_j}{n}\right)\left[1 - \left(\frac{\sum_{h=1}^k c_h^2}{n^2} + 2\left(\frac{c_j}{n}\right)\left(\frac{n-c_j}{n}\right)\right)\right]$$
$$= \left(\frac{c_j}{n^3}\right)\left(n^2 + c_j n - \sum_{h=1}^k c_h^2\right).$$

□

Lemma 2.2 (Next expected bias) *For any k-cd $\mathbf{c}$ and for every color $j \in [k]$ with $j \neq 1$, it holds that*

$$\mu_1 - \mu_j \geqslant s(\mathbf{c})\left(1 + \frac{c_1}{n}\left(1 - \frac{c_1}{n}\right)\right). \tag{1}$$

Proof. Observe that, when we assume $c_1 \geqslant c_2 \geqslant \cdots \geqslant c_k$, we can give the following upper bound on the sum of squares in Lemma 2.1

$$\sum_{h\in[k]} c_h^2 = c_1^2 + \sum_{h=2}^k c_h^2 \leqslant c_1^2 + c_2 \sum_{h=2}^k c_h \leqslant c_1^2 + nc_2. \tag{2}$$

From Lemma 2.1 it thus follows that, for any $j \neq 1$,

$$\mu_1 - \mu_j \geqslant \mu_1 - \mu_2 = (c_1 - c_2) + \frac{(c_1^2 - c_2^2)}{n} - \frac{c_1 - c_2}{n^2}\sum_{h\in k} c_h^2$$
$$= s(\mathbf{c})\left(1 + \frac{c_1 + c_2}{n} - \frac{1}{n^2}\sum_{h\in k} c_h^2\right)$$
$$\geqslant s(\mathbf{c})\left(1 + \frac{c_1 + c_2}{n} - \frac{c_1^2 + nc_2}{n^2}\right)$$
$$= s(\mathbf{c})\left(1 + \frac{c_1}{n}\left(1 - \frac{c_1}{n}\right)\right),$$

where in the inequality we used (2) and the fact that $c_1 - c_2 \geqslant 0$. □

3 Upper bounds for 3-majority

In this section, we provide the following upper bound on the convergence time of the 3-majority dynamics which clarifies the roles played by the plurality color and by the initial bias.

Theorem 3.1 (the general upper bound) *Let λ be any value such that $\lambda < \sqrt[3]{n}$ and let $\mathbf{c}$ be any initial k-cd, with $c_1 \geqslant n/\lambda$ and $s(\mathbf{c}) \geqslant 24\sqrt{2\lambda n \log n}$. Then the 3-majority protocol converges to the plurality color in $\mathcal{O}(\lambda \log n)$ time w.h.p.*

The next three corollaries of Theorem 3.1 address three relevant special cases. Corollary 3.2 is obtained by setting $\lambda = \min\left\{2k,\ \sqrt[3]{n/\log n}\right\}$ and it provides a bound which does not assume any condition on c_m.

Corollary 3.2 *Let $\mathbf{c}$ be any initial k-cd with*

$$s(\mathbf{c}) \geqslant 22\sqrt{\min\left\{2k,\ \sqrt[3]{\frac{n}{\log n}}\right\} n \log n}.$$

Then, the 3-majority protocol converges to the plurality color in $\mathcal{O}\left(\min\left\{2k,\ \sqrt[3]{n/\log n}\right\}\log n\right)$ time w.h.p.

Corollaries 3.3 and 3.4 are obtained by setting $\lambda = \mathrm{poly}\log(n)$ and $\lambda = \Theta(1)$, respectively. They provide sufficient conditions for a polylogarithmic convergence time.

Corollary 3.3 *Let* $\mathbf{c}$ *be any initial* k*-cd with* $c_1 \geqslant n/\log^{\ell} n$ *and* $s(\mathbf{c}) \geqslant 22\sqrt{n \log^{\ell+1} n}$*. Then, the* 3*-majority protocol converges to the plurality color in* $\mathcal{O}(\log^{\ell+1} n)$ *time w.h.p.*

Corollary 3.4 *Let* $\mathbf{c}$ *be any* k*-cd with* $c_1 \geqslant n/\beta$ *and* $s(\mathbf{c}) \geqslant 22\sqrt{\beta n \log n}$*, for some constant* $\beta \geqslant 3$*. Then, the* 3*-majority protocol converges to the plurality color in* $\mathcal{O}(\log n)$ *time w.h.p.*

In order to prove Theorem 3.1, we need the following three technical lemmas that essentially characterize three different phases of the process analysis. Each of them concerns a different range assumed by the plurality c_1. The first lemma considers configurations in which c_1 is under a suitable constant fraction of n: in this case, it shows that the bias between the plurality size and the size of any other color increases by a factor $1 + \Omega(c_1/n) = 1 + \Omega(1/\lambda)$.

Lemma 3.5 (from plurality to majority) *Let* $\mathbf{c}$ *be any* k*-cd with* $n/\lambda \leqslant c_1 \leqslant 2n/3$ *and* $s(\mathbf{c}) \geqslant \alpha\sqrt{\lambda n \log n}$ *where* $\lambda < \sqrt[3]{n}$ *and* α *is a sufficiently large constant. Then, for any other color* $j \neq 1$ *it holds that*

$$\mathbf{P}\left(C_1^{(t+1)} - C_j^{(t+1)} \geqslant s(\mathbf{c})\left(1 + \frac{c_1}{4n}\right) \,\middle|\, \mathbf{C}^{(t)} = \mathbf{c} \right) \geqslant 1 - \frac{1}{n^3}. \tag{3}$$

Proof. Conditional on any configuration $\mathbf{c}$, from the Chernoff bound w.h.p. it holds that

$$\begin{aligned} C_j &\leqslant \max\left\{\mu_j + \alpha\sqrt{\mu_j \log n}, \log n\right\} \qquad (4)\\ C_1 &\geqslant \mu_1 - \alpha\sqrt{\mu_1 \log n} \end{aligned}$$

Thus, if $\mu_j + \alpha\sqrt{\mu_j \log n} \geqslant \log n$, w.h.p. it holds that[3]

$$C_1 - C_j \geqslant \mu_1 - \mu_j - \alpha\sqrt{\mu_1 \log n} - \alpha\sqrt{\mu_j \log n} \geqslant \mu_1 - \mu_j - 2\alpha\sqrt{\mu_1 \log n}. \tag{5}$$

Otherwise, if $\mu_j + \alpha\sqrt{\mu_j \log n} < \log n$, then w.h.p. it holds that

$$\begin{aligned} C_1 - C_j &\geqslant \mu_1 - \alpha\sqrt{\mu_1 \log n} - 2\log n\\ &\geqslant \mu_1 - \frac{4}{3}\alpha\sqrt{\mu_1 \log n}\\ &\geqslant \mu_1 - \frac{5}{3}\alpha\sqrt{\mu_1 \log n} - \mu_j - \alpha\sqrt{\mu_j \log n}\\ &\geqslant \mu_1 - \mu_j - 2\alpha\sqrt{\mu_1 \log n}. \qquad (6) \end{aligned}$$

From Lemma 2.2 and the hypothesis $c_1 \leqslant 2n/3$ we get that

$$\mu_1 - \mu_j \geqslant (c_1 - c_j)\left(1 + \frac{c_1}{3n}\right).$$

Thus, in (5) and (6) we get

$$\begin{aligned} \mu_1 - \mu_j - 2\alpha\sqrt{\mu_1 \log n} &\geqslant (c_1 - c_j)\left(1 + \frac{c_1}{3n}\right) - 2\alpha\sqrt{2c_1 \log n}\\ &\geqslant (c_1 - c_j)\left(1 + \frac{c_1}{3n}\right) - 2\alpha\sqrt{2c_1 \log n}\\ &\geqslant (c_1 - c_j)\left(1 + \frac{c_1}{3n} - \frac{2\alpha\sqrt{2c_1 \log n}}{(c_1 - c_j)}\right)\\ &\geqslant (c_1 - c_j)\left(1 + \frac{c_1}{3n} - \frac{\sqrt{c_1/n}}{12\sqrt{\lambda}}\right)\\ &\geqslant (c_1 - c_j)\left(1 + \frac{c_1}{3n}\left(1 - \frac{1}{4\sqrt{c_1\lambda/n}}\right)\right)\\ &\geqslant (c_1 - c_j)\left(1 + \frac{c_1}{4n}\right), \end{aligned}$$

[3] We are using the fact that $\Pr(A \cap B) \geqslant 1 - \Pr(A^C) - \Pr(B^C)$.

concluding the proof.

□

Once c_1 becomes larger than $2n/3$ the negative occurrence of c_1 in (1) does not allow to directly show a drift towards plurality. We thus consider another useful "drift" of the process: The sum of all the other color sizes decreases exponentially w.h.p., as long as this sum is enough large to apply concentration bounds. This result is formalized in the next lemma.

Lemma 3.6 (from majority to almost all) *Let* $\mathbf{c}$ *be any* k*-cd with* $2n/3 \leqslant c_1 \leqslant n - \omega(\log n)$. *Then, it holds that*

$$\mathbf{P}\left(\sum_{i \neq 1} C_i^{(t+1)} \leqslant \frac{8}{9} \sum_{i \neq 1} c_i \,\middle|\, \mathbf{C}^{(t)} = \mathbf{c} \right) \geqslant 1 - \frac{1}{n^3}.$$

Proof. Let us define $\mu_{-1}^{(t)} = \sum_{i \neq 1} \mu_i^{(t)}$. By using (2) we have

$$\begin{aligned}
\frac{\mu_{-1}^{(t+1)}}{n} &= \sum_{i \neq 1} \frac{c_i}{n} \left(1 + \frac{c_i}{n} - \sum_j \left(\frac{c_j}{n} \right)^2 \right) \\
&= 1 - \frac{c_1}{n} + \sum_{i \neq 1} \left(\frac{c_i}{n} \right)^2 - \left(1 - \frac{c_1}{n} \right) \sum_j \left(\frac{c_j}{n} \right)^2 \\
&= 1 - \frac{c_1}{n} - \left(\frac{c_1}{n} \right)^2 + \frac{c_1}{n} \sum_j \left(\frac{c_j}{n} \right)^2 \\
&\leqslant 1 - \frac{c_1}{n} - \left(\frac{c_1}{n} \right)^2 + \frac{c_1}{n} \left(\left(\frac{c_1}{n} \right)^2 + \frac{c_2}{n} \left(1 - \frac{c_1}{n} \right) \right) \\
&= \left(1 - \frac{c_1}{n} \right) \left(1 - \left(\frac{c_1}{n} \right)^2 + \frac{c_1}{n} \frac{c_2}{n} \right) = \left(1 - \frac{c_1}{n} \right) \left(1 - \frac{c_1}{n} \left(\frac{c_1}{n} - \frac{c_2}{n} \right) \right).
\end{aligned}$$

Using the hypothesis $c_1/n \geqslant 2/3$ (hence $c_2/n \leqslant 1/3$) the last expression become

$$\left(1 - \frac{c_1}{n} \right) \left(1 - \frac{c_1}{n} \left(\frac{c_1}{n} - \frac{c_2}{n} \right) \right) \leqslant \left(1 - \frac{c_1}{n} \right) \left(1 - \frac{c_1}{3n} \right) \leqslant \left(1 - \frac{c_1}{n} \right) = \frac{7}{9} \frac{\mu_{-1}^{(t)}}{n}. \tag{7}$$

Now observe that, from the Chernoff bound, as long as $\mu_{-1}^{(t+1)} \in \omega(\log n)$, w.h.p. it holds

$$\sum_{i \neq 1} C_i^{(t+1)} \leqslant \mu_{-1}^{(t+1)} + \sqrt{\mu_{-1}^{(t+1)} \log n} = \mu_{-1}^{(t+1)} \left(1 + \sqrt{\frac{\log n}{\mu_{-1}^{(t+1)}}} \right) = \mu_{-1}^{(t+1)} \left(1 + o(1) \right). \tag{8}$$

Moreover, from Lemma 2.1 it follows that

$$\mu_1 \leqslant 2 c_1. \tag{9}$$

Thus, by replacing (9) and (7) in (8), we get that w.h.p. it holds

$$\sum_{i \neq 1} C_i^{(t+1)} \leqslant \frac{7}{9} \mu_{-1}^{(t+1)} \left(1 + o(1) \right) \leqslant \frac{8}{9} \sum_{i \neq 1} \mu_i^{(t+1)},$$

concluding the proof.

□

Finally, when the sum of all the minority colors is not larger than a polylogarithmic function, the probability that they all disappear in one round is high. This is shown in the new, final lemma.

Lemma 3.7 (the last step) *Let* $\alpha > 0$ *and let* $\mathbf{c}$ *be any* k*-cd with* $c_1 \geqslant n - \log^{\alpha} n$. *Then, it holds that*

$$\mathbf{P}\left(\sum_{i \neq 1} C_i^{(t+1)} = 0 \,|\, \mathbf{C}^{(t)} = \mathbf{c} \right) \geqslant 1 - \frac{3 \log^{2\alpha} n}{n}. \tag{10}$$

Proof. As in the previous proof let $\mu_{-1} = \sum_{i \neq 1} \mu_i$. Note that $c_1 \geqslant n - \log^\alpha n$ implies $\sum_{i \neq 1} c_i \leqslant \log^\alpha n$. Thus, from Lemma 2.1 we have

$$\begin{aligned}
\mu_{-1} &= \sum_{i \neq 1} c_i \left(1 + \frac{c_i}{n} - \sum_j \left(\frac{c_j}{n} \right)^2 \right) \\
&\leqslant \sum_{i \neq 1} c_i \left(1 + \frac{c_i}{n} - \left(\frac{c_1}{n} \right)^2 \right) \\
&= \sum_{i \neq 1} c_i \left(1 + \frac{c_i}{n} - \left(1 - \sum_{j \neq 1} \frac{c_j}{n} \right)^2 \right) \\
&\leqslant \sum_{i \neq 1} c_i \left(\frac{c_i}{n} + 2 \sum_{j \neq 1} \frac{c_j}{n} \right) \\
&\leqslant \sum_{i \neq 1} c_i \left(\frac{3 \log^\alpha n}{n} \right) = \frac{3 \log^{2\alpha} n}{n}.
\end{aligned}$$

Finally, (10) follows from Markov's inequality on the event "$\sum_{i \neq 1} C_i^{(t+1)} \geqslant 1$" and, since $\sum_{i \neq 1} C_i^{(t+1)}$ is a non-negative integer-valued r.v., this is equivalent as "$\sum_{i \neq 1} C_i^{(t+1)} > 0$".

□

Proof of Theorem 3.1 From Lemma 3.5, we prove that, as long as the number of nodes with the plurality color c_1 is smaller than a constant fraction of n, the bias between c_1 and c_2 increases by a factor $(1 + \lambda^{-1})$, w.h.p. (ii) In Lemma 3.6 we prove that, when the plurality color reaches a suitable constant fraction of n, then the number of nodes with non-plurality colors decreases at exponential rate, w.h.p. Finally, (iii) in Lemma 3.7 we consider separately the last step of the protocol, where all colors but the plurality one disappear w.h.p. □

Plurality consensus in an adversarial model. Let us consider a dynamic adversary that can change the color of up to T nodes at the beginning of each round and assume $T = o(s/\lambda)$. Then, Theorem 3.1 still holds since the impact of such a T-*bounded* adversary is negligible in the growth of the bias s. Indeed, from Lemma 3.5 w.h.p. it holds

$$C_1^{(t+1)} - C_j^{(t+1)} \geqslant s(\mathbf{c}) + \frac{s(\mathbf{c})}{4\lambda}$$

Then, for any $T = o(s/\lambda)$, w.h.p. we have that

$$C_1^{(t+1)} - C_j^{(t+1)} \geqslant s(\mathbf{c}) + \frac{s(\mathbf{c})}{4\lambda} - T \geqslant s(\mathbf{c}) + \Theta\left(\frac{s(\mathbf{c})}{\lambda} \right)$$

For instance, from Corollary 3.2, when $k \leqslant 2 \sqrt[3]{n/\log n}$, the resilience of the 3-majority dynamics is $T = o(s/k)$.

4 Lower bounds

This section is organized in four subsections: in the first one, we prove a lower bound on the convergence time of the 3-majority dynamics; in the second subsection, we show that 3-majority is essentially the only 3-input dynamics that converges to plurality consensus; in the third subsection, we provide a lower bound on the convergence time of the h-plurality dynamics for $h > 3$; finally, in the fourth subsection we show that our assumption on the magnitude of the initial bias is in a sense (almost) tight if one wants to prove the bounds of Section 3 with high probability.

4.1 Lower bound for 3-majority

In this section we show that if the 3-majority dynamics starts from a sufficiently balanced configuration (i.e., at the beginning there are $n/k \pm o(n/k)$ nodes of every color) then it will take $\Omega(k \log n)$ rounds w.h.p. to reach one of the absorbing configurations where all nodes have the same color. In what follows, all events and random variables thus concern the Markovian process yielded by the 3-majority dynamics.

In the next lemma we show that if there are at most $n/k + b$ nodes of a specific color, where b is smaller than n/k, then at the next round there are at most $n/k + (1 + 3/k)b$ nodes of that color w.h.p.

Lemma 4.1 *Let the number of colors k be such that $k \leqslant (n/\log n)^{1/4}$, let b be any number with $k\sqrt{n \log n} \leqslant b \leqslant n/k$, and let $\mathbf{c} = (c_1, \dots, c_k)$ be a configuration. If $c_j = n/k + a$ for some color $j \in [k]$ and for some $a \leqslant b$, then the number of nodes with color j at the next round are at most $n/k + (1+3/k)b$ w.h.p.; more precisely, for any $a \leqslant b$ and for any configuration $\mathbf{c}$ such that $c_j = n/k + a$ it holds that*

$$\mathbf{P}\left(C_j^{(t+1)} \geqslant \frac{n}{k} + \left(1 + \frac{3}{k}\right) b \;\middle|\; \mathbf{C}^{(t)} = \mathbf{c} \right) \leqslant \frac{1}{n^2}.$$

Proof. For any configuration $\mathbf{c} = (c_1, \dots, c_k)$ with $\sum_{j=1}^k c_j = n$ and any color $j \in [k]$, the expected value of the number of nodes colored j at round $t+1$ conditional on $\left\{\mathbf{C}^{(t)} = \mathbf{c}\right\}$ is (see Lemma 2.1)

$$\mathbf{E}\left[C_j^{(t+1)} \mid \mathbf{C}^{(t)} = \mathbf{c}\right] = c_j \left(1 + \frac{c_j}{n} - \frac{1}{n^2}\sum_{j=1}^{k} c_j^2\right).$$

Observe that, since $\sum_{j=1}^k c_j = n$, from Jensen inequality (see Lemma A.2) it follows that $(1/n^2)\sum_{j=1}^k c_j^2 \geqslant 1/k$. Hence, we can give an upper bound on the expectation of $C_j^{(t+1)}$ that depends only on c_j and not on the whole configuration $\mathbf{c}$ at round t, namely

$$\mathbf{E}\left[C_j^{(t+1)} \mid \mathbf{C}^{(t)}\right] \leqslant C_j^{(t)} \left(1 + \frac{C_j^{(t)}}{n} - \frac{1}{k}\right).$$

If we condition on the number of nodes of color j being $c_j = n/k + a$ in configuration $\mathbf{c}$, for some $a \leqslant b$, we get

$$\begin{aligned}\mathbf{E}\left[C_j^{(t+1)} \mid \mathbf{C}^{(t)} = \mathbf{c}\right] &\leqslant \left(\frac{n}{k} + a\right)\left(1 + \frac{n/k + a}{n} - \frac{1}{k}\right) = \frac{n}{k} + \left(1 + \frac{1}{k}\right) a + \frac{a^2}{n} \\ &\leqslant \frac{n}{k} + \left(1 + \frac{1}{k}\right) b + \frac{b^2}{n} \leqslant \frac{n}{k} + \left(1 + \frac{2}{k}\right) b,\end{aligned}$$

where in the last two inequalities we used that $a \leqslant b$ and $b \leqslant n/k$.[4] Since $C_j^{(t+1)}$ conditional on $\left\{\mathbf{C}^{(t)} = \mathbf{c}\right\}$ can be written as a sum of n independent Bernoulli random variables, from the Chernoff bound (see Lemma A.1) we thus get that for every $a \leqslant b$ it holds that

$$\mathbf{P}\left(C_j^{(t+1)} \geqslant \frac{n}{k} + \left(1 + \frac{3}{k}\right) b \;\middle|\; \mathbf{C}^{(t)} = \mathbf{c} \right) \leqslant e^{-2(b/k)^2/n} \leqslant \frac{1}{n^2},$$

where in the last inequality we used that $b \geqslant k\sqrt{n \log n}$. □

Let us say that a configuration $\mathbf{c} = (c_1, \dots, c_k) \in \{0, 1, \dots, n\}^k$ with $\sum_{j=1}^k c_j = n$ is *monochromatic* if there is an $j \in [k]$ such that $c_j = n$. In the next theorem we show that if we start from a sufficiently *balanced* configuration, then the 3-majority protocol takes $\Omega(k \log n)$ rounds w.h.p. to reach a monochromatic configuration.

Theorem 4.2 *Let $\tau = \inf\{t \in \mathbb{N} : \mathbf{C}^{(t)}$ is monochromatic$\}$ be the random variable indicating the first round such that the system is in a monochromatic configuration. If the initial number of colors is $k \leqslant (n/\log n)^{1/4}$ and the initial configuration is $\mathbf{c} = (c_1, \dots, c_k)$ with $\max\{c_j : j = 1, \dots, k\} \leqslant \frac{n}{k} + \left(\frac{n}{k}\right)^{1-\varepsilon}$ then $\tau = \Omega(k \log n)$ w.h.p.*

[4] Notice that the inequality holds in particular for negative a as well

Idea of the proof. For a color $j \in [k]$ let us denote the difference $C_j - n/k$ as the *positive unbalance*. In Lemma 4.1 we proved that, as long as the positive unbalance of a color is smaller than n/k, this will increase by a factor smaller than $(1+3/k)$ at every round (w.h.p.). Hence, if a color starts with a positive unbalance smaller than $(n/k)^{1-\varepsilon}$, then it will take $\Omega(k \log n)$ rounds to reach an unbalance of n/k w.h.p. By union bounding on all the colors, we can get the stated lower bound. □

Proof. Observe that if $T \leqslant c\, k \log n$, for a suitable positive constant c, then $(1-3/k)^T (n/k)^{1-\varepsilon}$ is smaller than n/k. Since in the initial configuration $\mathbf{c}$ for any color $j \in [k]$ we have that $c_j \leqslant n/k + (n/k)^{1-\varepsilon}$, for $T \leqslant c\, k \log n$ it holds that

$$\mathbf{P}\left(C_j^{(t)} = n \;\middle|\; \mathbf{C}^{(0)} = \mathbf{c} \right) \leqslant \mathbf{P}\left(C_j^{(T)} \geqslant \frac{n}{k} + \left(1 + \frac{3}{k}\right)^T \left(\frac{n}{k}\right)^{1-\varepsilon} \;\middle|\; \mathbf{C}^{(0)} = \mathbf{c} \right), \tag{11}$$

Since $c_j \leqslant n/k + (n/k)^{1-\varepsilon}$, if we also have $C_j^{(T)} \geqslant n/k + (1+3/k)^T (n/k)^{1-\varepsilon}$, then a round t with $0 \leqslant t \leqslant T-1$ must exist such that $C_j^{(t)} \leqslant n/k + b$ and $C_j^{(t+1)} \geqslant n/k + (1+3/k)b$ for some value b, with $k\sqrt{n \log n} \leqslant b \leqslant n/k$, thus

$$\mathbf{P}\left(C_j^{(T)} \geqslant \frac{n}{k} + \left(1 + \frac{3}{k}\right)^T \left(\frac{n}{k}\right)^{1-\varepsilon} \;\middle|\; \mathbf{C}^{(0)} = \mathbf{c} \right) \tag{12}$$

$$\leqslant \mathbf{P}\left(\exists 0 \leqslant t \leqslant T-1 \;:\; C_j^{(t)} \leqslant \frac{n}{k} + b \text{ and } C_j^{(t+1)} \geqslant \frac{n}{k} + \left(1 + \frac{3}{k}\right) b \;\middle|\; \mathbf{C}^{(0)} = \mathbf{c} \right) \tag{13}$$

$$\leqslant \sum_{t=0}^{T-1} \mathbf{P}\left(C_j^{(t)} \leqslant \frac{n}{k} + b_t \text{ and } C_j^{(t+1)} \geqslant \frac{n}{k} + \left(1 + \frac{3}{k}\right) b_t \;\middle|\; \mathbf{C}^{(0)} = \mathbf{c} \right) \tag{14}$$

where the inequality from (12) to (13) holds for some b with $k\sqrt{n \log n} \leqslant b \leqslant n/k$, and the inequality from (13) to (14) holds for some $b_0, \ldots, b_{T-1}$ with $k\sqrt{n \log n} \leqslant b_t \leqslant n/k$ for every $t = 0, \ldots, T-1$. Now observe that

$$\begin{aligned}
&\mathbf{P}\left(C_j^{(t)} \leqslant \frac{n}{k} + b_t \text{ and } C_j^{(t+1)} \geqslant \frac{n}{k} + \left(1 + \frac{3}{k}\right) b_t \;\middle|\; \mathbf{C}^{(0)} = \mathbf{c} \right) \\
&= \sum_{a \leqslant b_t} \mathbf{P}\left(C_j^{(t)} = \frac{n}{k} + a \text{ and } C_j^{(t+1)} \geqslant \frac{n}{k} + \left(1 + \frac{3}{k}\right) b_t \;\middle|\; \mathbf{C}^{(0)} = \mathbf{c} \right) \\
&= \sum_{a \leqslant b_t} \mathbf{P}\left(C_j^{(t+1)} \geqslant \frac{n}{k} + \left(1 + \frac{3}{k}\right) b_t \;\middle|\; C_j^{(t)} = \frac{n}{k} + a \text{ and } \mathbf{C}^{(0)} = \mathbf{c} \right) \cdot \mathbf{P}\left(C_j^{(t)} = \frac{n}{k} + a \;\middle|\; \mathbf{C}^{(0)} = \mathbf{c} \right) \\
&\leqslant \frac{1}{n^2} \sum_{a \leqslant b_t} \mathbf{P}\left(C_j^{(t)} = \frac{n}{k} + a \;\middle|\; \mathbf{C}^{(0)} = \mathbf{c} \right) \leqslant \frac{1}{n^2},
\end{aligned} \tag{15}$$

where in the last line we used Lemma 4.1.

By combining (11), (14), and (15) we get that, for every color $j \in [k]$, if the initial number of nodes colored j is $c_j \leqslant n/k + (n/k)^{1-\varepsilon}$ at any round $T \leqslant c\, k \log n$ the probability that all nodes are colored j is at most T/n^2. The probability that $\mathbf{C}^{(T)}$ is monochromatic is thus at most $(kT)/n^2 \leqslant n^{-\alpha}$ for some positive constant α. □

It may be worth noticing that what we actually prove in Theorem 4.2 is that $\Omega(k \log n)$ rounds are required in order to go from a configuration where the majority color has at most $n/k + (n/k)^{1-\varepsilon}$ nodes to a configuration where it has $2n/k$ colors.

4.2 A negative result for 3-input dynamics

In order to prove that dynamics that differ from the majority ones do not solve plurality consensus, we first give some formal definitions of the dynamics we are considering.

Definition 4.3 ($\mathcal{D}_h(k)$ protocols) *An h*-dynamics *is a synchronous protocol where at each round every node picks h random neighbors (including itself and with repetition) and recolors itself according to some deterministic rule that depends only on the colors it sees. Let $\mathcal{D}_h(k)$ be the class of h-dynamics and observe that a dynamics $\mathcal{P} \in \mathcal{D}_h$ can be specified by a function*

$$f \,:\, [k]^h \rightarrow [k],$$

such that $f(x_1, \dots, x_h) \in \{x_1, \dots, x_h\}$. Where $f(x_1, \dots, x_h)$ is the color chosen by a node that sees the (ordered) sequence $(x_1, \dots, x_h)$ of colors.

In the class $\mathcal{D}_3(k)$, there is a subset $\mathcal{M}^3$ of equivalent protocols called 3-majority dynamics having two key-properties described below: the clear-majority and the uniform one.

Definition 4.4 (clear-majority property) *Let $(x_1, x_2, x_3) \in [k]^3$ be a triple of colors. We say that (x_1, x_2, x_3) has a* clear majority *if at least two of the three entries have the same value. A dynamics $\mathcal{P} \in \mathcal{D}_3(k)$ has the* clear-majority *property if whenever its f sees a clear majority it returns the majority color.*

Given any 3-input dynamics function $f(x_1, x_2, x_3)$, for any triple of distinct colors $r, g, b \in [k]$, let $\Pi(r, g, b)$ be the subset of permutations of the colors r, g, b and define the following "counters":

$$\begin{aligned} \delta_r &= |\{(z_1, z_2, z_3) \in \Pi(r, g, b),\ s.t.\ f(z_1, z_2, z_3) = r\}|, \\ \delta_g &= |\{(z_1, z_2, z_3) \in \Pi(r, g, b),\ s.t.\ f(z_1, z_2, z_3) = g\}|, \\ \delta_b &= |\{(z_1, z_2, z_3) \in \Pi(r, g, b),\ s.t.\ f(z_1, z_2, z_3) = b\}|. \end{aligned}$$

Observe that for any 3-inputs dynamics it must hold $\delta_g + \delta_r + \delta_b = 6$.

Definition 4.5 (uniform property) *A dynamics $\mathcal{P} \in \mathcal{D}_3(k)$ has the* uniform *property if, for any triple of distinct colors $r, g, b \in [k]$, it holds that $\delta_r = \delta_g = \delta_b$ $(= 2)$.*

Informally speaking, the clear-majority and the uniform properties provide a clean characterization of those dynamics that are good solvers for plurality consensus. This fact is formalized in the next definitions and in the final theorem.

Definition 4.6 (3-majority dynamics) *A protocol $\mathcal{P} \in \mathcal{D}_3(k)$ belongs to the class $\mathcal{M}^3 \subset \mathcal{D}_3(k)$ of* 3-majority dynamics *if its function $f(x_1, x_2, x_3)$ has the clear-majority and the uniform properties.*

Definition 4.7 ((s, ε)-plurality consensus solver) *We say that a protocol $\mathcal{P}$ is an* (s, ε)-solver *(for the plurality consensus problem) if for every initial s-biased configuration* $\mathbf{c}$, *when running $\mathcal{P}$, with probability at least $1 - \varepsilon$ there is a round t by which all nodes gets the plurality color of c.*

Let us observe that, by definition of h-dynamics, any monochromatic configuration is an absorbing state of the relative Markovian process. Moreover, the smaller s and ε the better an (s, ε)-solver is; in other words, if a dynamics is an (s, ε)-solver then it is also an (s', ε')-solver for every $s' \geqslant s$ and $\varepsilon' \geqslant \varepsilon$. In Section 3, we showed that any dynamics in $\mathcal{M}^3$ is a $(\Theta(\sqrt{\min\{2k, (n/\log n)^{1/3}\} n \log n}), \Theta(1/n))$-solver in $\mathcal{D}_3$. We can now state the main result of this section.

Theorem 4.8 (properties of good solvers) *Given a protocol $\mathcal{P}$, the following hold:*

(a) If $\mathcal{P}$ is an $(n/4, 1/4)$-solver in $\mathcal{D}_3$, then its f must have the clear-majority property.

(b) A constant $\eta > 0$ exists such that, if $\mathcal{P}$ is an $(\eta \cdot n, 1/4)$-solver, then its f must have the uniform property.

The above theorem also provides the clear reason why some dynamics can solve consensus but cannot solve plurality consensus in the non-binary case. A relevant example is the *median* dynamics studied in [7]: it has the clear-majority property but not the uniform one.

For readability sake, we split the proof of the above theorem in two technical lemmas: in the first one, we show the first claim about clear majority while in the second lemma we show the second claim about the uniform property.

Lemma 4.9 (clear majority) *If a protocol $\mathcal{P} \in \mathcal{D}_3$ is an $(n/4, 1/4)$-solver, then it chooses the majority color every time there is a triple with a clear majority.*

Proof. For every triple of colors $(x_1, x_2, x_3) \in [k]^3$ that has a clear majority, let us define $\delta(x_1, x_2, x_3)$ to be 1 if protocol $\mathcal{P}$ behaves like the majority protocol over triple (x_1, x_2, x_3) and 0 otherwise. Consider an initial configuration with only two colors, say red (r) and blue (b), with c_r red nodes and $c_b = n - c_r$ blue nodes. Let us define Δ_r and Δ_b as follows

$$\Delta_r = \delta(r, r, b) + \delta(r, b, r) + \delta(b, r, r),$$
$$\Delta_b = \delta(b, b, r) + \delta(b, r, b) + \delta(r, b, b).$$

We can write the probability that a node chooses color red as

$$\begin{aligned} p(r) &= \left(\frac{c_r}{n}\right)^3 + \left(\frac{c_r}{n}\right)^2 \frac{c_b}{n} \cdot \Delta_r + \left(\frac{c_b}{n}\right)^2 \frac{c_r}{n} (3 - \Delta_b) \\ &= \frac{c_r}{n^3} \left(c_r^2 + c_b \left(c_r \Delta_r - c_b \Delta_b\right) + 3c_b^2\right). \end{aligned} \tag{16}$$

Observe that for a majority protocol we have that $\Delta_r = \Delta_b = 3$. In what follows we show that if this is not the case then there are configurations where the majority color does not increase in expectation. We distinguish two cases, case $\Delta_r \neq \Delta_b$ and case $\Delta_r = \Delta_b$.

Case $\Delta_r \neq \Delta_b$: Suppose w.l.o.g. that $\Delta_r < \Delta_b$, and observe that since they have integer values it means $\Delta_r \leqslant \Delta_b - 1$. Now we show that, if we start from a configuration where the red color has the majority of nodes, the number of red nodes decreases in expectation. By using $\Delta_r \leqslant \Delta_b - 1$ in (16) we get

$$p(r) \leqslant \frac{c_r}{n^3} \left(c_r^2 + c_b(c_r - c_b)\Delta_b - c_r c_b + 3c_b^2\right). \tag{17}$$

If the majority of nodes is red then $c_r - c_b$ is positive, and since Δ_b can be at most 3 from (17) we get

$$p(r) \leqslant \frac{c_r}{n^3} \left(c_r^2 + 2c_r c_b\right). \tag{18}$$

Finally, if we put $c_r = n/2 + s$ and $c_b = n/2 - s$, for some positive s, in (18), we get that

$$p(r) \leqslant \frac{c_r}{n^3} \left(\frac{3}{4}n^2 + (n - s)s\right) \leqslant \frac{c_r}{n}. \tag{19}$$

Case $\Delta_r = \Delta_b$: When $\Delta_r = \Delta_b$, observe that if the protocol is not a majority protocol then it must be $\Delta_r = \Delta_b \leqslant 2$. Hence, if we start again from a configuration where $c_r \geqslant c_b$, from (16) we get that

$$p(r) \leqslant \frac{c_r}{n^3} \left(c_r^2 + 2c_b(c_r - c_b) + 3c_b^2\right) = \frac{c_r}{n}. \tag{20}$$

In both cases, for any protocol $\mathcal{P}$ that does not behave like a majority protocol on triples with a clear majority, if we name X_t the random variable indicating the number of red nodes at round t, from (19) and (20) we get that $\mathbf{E}\left[X_{t+1} \,|\, X_t\right] \leqslant X_t$, hence X_t is a supermartingale. Now let τ be the random variable indicating the first time the chain hits one of the two absorbing states, i.e.

$$\tau = \inf\{t \in \mathbb{N} \,:\, X_t \in \{0, n\}\}.$$

Since $\mathbf{P}\left(\tau < \infty\right) = 1$ and all X_t's have values bounded between 0 and n, from the martingale stopping theorem[5] we get that $\mathbf{E}\left[X_\tau\right] \leqslant \mathbf{E}\left[X_0\right]$. If we start from a configuration that is $n/4$-unbalanced in favor of the red color, we have that $X_0 = n/2 + n/8$, and if we call ε is the probability that the process ends up with all blue nodes we have that $\mathbf{E}\left[X_\tau\right] = (1 - \varepsilon)n$. Hence it must be $(1 - \varepsilon)n \leqslant n/2 + n/8$ and the probability to end up with all blue nodes is $\varepsilon \geqslant 5/8 > 1/4$. Thus the protocol is not a $(n/4, 1/4)$-solver. □

Lemma 4.10 (uniform property) *A constant $\eta > 0$ exists such that, if $\mathcal{P}$ is an $(\eta n, 1/4)$-solver, then its f must have the uniform property.*

[5] See e.g. Chapter 17 in [14] for a summary of martingales and related results

Proof. Thanks to the previous lemma, we can assume that f has the clear-majority property but a triple (r, g, b) exists such that $\delta_r < \max\{\delta_g, \delta_b\}$. Let us start the process with the following initial configuration having only the above 3 colors and then show that the process w.h.p. will not converge to the plurality color r:

$$\mathbf{c} = (c_r, c_g, c_b) = (n/3 + s,\ n/3,\ n/3 - s) \quad \text{where } s = \Theta(\sqrt{n \log n}).$$

We consider the "hardest" case where $\delta_r = 1$: the case $\delta_r = 0$ is simpler since in this case, no matter how the other $\delta's$ are distributed, it is easy to see that the r.v. c_r will decrease exponentially to 0 starting from the above configuration.

- Case $\delta_r = 1$**,** $\delta_g = 3$**, and** $\delta_b = 2$ (and color-symmetric cases). Starting from the above initial configuration, we can compute the probability $p(r) = \mathbf{P}\left(X_v = r \mid C = \mathbf{c}\right)$ that a node gets the color r.

$$\begin{aligned} p(r) &= \left(\frac{c_r}{n}\right)^3 + 3\left(\frac{c_r}{n}\right)^2 \frac{n - c_r}{n} + \frac{c_r c_g c_b}{n^3} \\ &= \frac{n+3s}{3n^3}\left(\left(\frac{n}{3}+s\right)^2 + 3\left(\frac{n}{3}+s\right)\left(\frac{2}{3}n - s\right) + \left(\frac{n}{3}\right)\left(\frac{n}{3}-s\right)\right). \end{aligned}$$

After some easy calculations, we get

$$p(r) = \frac{8}{27}\left(1 + O\left(\frac{s}{n}\right)\right).$$

As for $p(g)$, by similar calculations, we obtain the following bound

$$p(g) = \frac{10}{27}\left(1 - O\left(\frac{s^2}{n^2}\right)\right).$$

From the above two equations, we get the following bounds on the expectation of the r.v.'s X^r and X^g counting the nodes colored with r and g, respectively (at the next round).

$$\begin{aligned} \mathbb{E}\left[X^r \mid \mathbf{C} = \mathbf{c}\right] &\leqslant \frac{8}{27}n + O(s) \quad \text{and} \\ \mathbb{E}\left[X^g \mid \mathbf{C} = \mathbf{c}\right] &\geqslant \frac{10}{27}n - O\left(\frac{s^2}{n}\right). \end{aligned}$$

By a standard application of Chernoff's Bound, we can prove that, if $s \leqslant \eta n$ for a sufficiently small $\eta > 0$, the initial value c_r will w.h.p. decreases by a constant factor, going much below the new plurality c_g. Then, by applying iteratively the above reasoning we get that the process will not converge to r, w.h.p.

- Case $\delta_r = 1$**,** $\delta_g = 4$**, and** $\delta_b = 1$ (and color-symmetric cases). In this case it is even simpler to show that w.h.p., starting from the same initial configuration considered in the previous case, the process will not converge to color r. □

4.3 A lower bound for h-plurality

In Subsection 4.1, we have shown that the 3-majority protocol takes $\Theta(k \log n)$ rounds w.h.p. to converge in the worst case. A natural question is whether by using the h-plurality protocol, with h slightly larger than 3, it is possible to significantly speed-up the process. We prove that this is not the case.

Let us consider a set of n nodes, each node colored with one out of k colors. The h-plurality protocol works as follows:

> *At every round, every node picks h nodes uniformly at random (including itself and with repetitions) and recolors itself according to the plurality of the colors it sees (breaking ties u.a.r.)*

Let $j \in [k]$ be an arbitrary color, in the next lemma we prove that, if the number of j-colored nodes is smaller than $2n/k$ and if $k/h = \mathcal{O}(n^{(1-\varepsilon)/4})$, then the probability that the number of j-nodes increases by a factor $(1 + h^2/k)$ is exponentially small.

Lemma 4.11 *Let $\mathbf{c} = (c_1, \ldots, c_k)$ be a configuration and let $j \in [k]$ be a color such that $(n/k) \leqslant c_j \leqslant 2(n/k)$. If $k/h = \mathcal{O}(n^{(1-\varepsilon)/4})$ then it holds that*

$$\mathbf{P}\left(C_j^{(t+1)} \geqslant \left(1 + \frac{h^2}{k}\right) c_j \,\middle|\, \mathbf{C}^{(t)} = \mathbf{c} \right) \leqslant e^{-\Theta(n^{\varepsilon})}.$$

Proof. Consider a specific node, say $u \in [n]$, let N_j be the number of j-colored nodes picked by u during the sampling stage of the t-th round and let Y be the indicator random variable of the event that node u chooses color j at round $t+1$. We give an upper bound on the probability of the event $Y = 1$ by conditioning it on $N_j = 1$ and $N_j \geqslant 2$ (observe that if $N_j = 0$ node u cannot choose j as its color at the next round)

$$\mathbf{P}(Y_u = 1) \leqslant \mathbf{P}(Y_u = 1 \mid N_j = 1)\, \mathbf{P}(N_j = 1) + \mathbf{P}(N_j \geqslant 2). \tag{21}$$

Now observe that

- $\mathbf{P}(Y_u = 1 \mid N_j(u) = 1) \leqslant 1/h$ since it is exactly $1/h$ if all other sampled nodes have distinct colors and it is 0 otherwise;
- $\mathbf{P}(N_j = 1) \leqslant hc_j/n$ since it can be bounded by the probability that at least one of the h samples gives color j;
- $\mathbf{P}(N_j \geqslant 2) \leqslant \binom{h}{2} c_j^2/n^2$ since it is the probability that a pair of sampled nodes exist with the same color j.

Hence, in (21) we have that

$$\mathbf{P}(Y = 1) \leqslant \frac{c_j}{n} + \frac{h^2}{2} \cdot \frac{c_j^2}{n^2}.$$

Thus, for the expected number of j-colored nodes at the next round we get

$$\mathbf{E}\left[C_j^{(t+1)} \mid \mathbf{C}^{(t)} = \mathbf{c}\right] \leqslant c_j + \frac{h^2}{2n} c_j^2 = c_j\left(1 + \frac{h^2}{2n} c_j\right) \leqslant c_j \left(1 + \frac{h^2}{k}\right),$$

where in the last inequality we used the hypothesis $c_j \leqslant 2(n/k)$. Since $C_j^{(t+1)}$ conditional on $\{\mathbf{C}^{(t)} = \mathbf{c}\}$ is a sum of n independent Bernoulli random variables, from the Chernoff bound (Lemma A.1 with $\lambda = c_j h^2/k$), we finally get

$$\begin{aligned}\mathbf{P}\left(C_j^{(t+1)} \geqslant c_j\left(1 + 2\frac{h^2}{k}\right) \,\middle|\, \mathbf{C}^{(t)} = \mathbf{c} \right) &\leqslant \exp\left(-\frac{2(c_j h^2/k)^2}{n}\right) \\ &\leqslant \exp\left(-\Omega(n^{\varepsilon})\right),\end{aligned}$$

where in the last inequality we used $c_j \geqslant n/k$ and $k/h = \mathcal{O}(n^{(1-\varepsilon)/4})$. □

By adopting a similar argument to that used for proving Theorem 4.2, we can get a lower bound $\Omega(k/h^2)$ on the completion time of the h-plurality.

Theorem 4.12 *Let $\mathbf{C}^{(t)}$ be the random variable indicating the configuration at round t according to the h-plurality protocol and let $\tau = \inf\{t \in \mathbb{N} : \mathbf{C}^{(t)}$ is monochromatic$\}$. If the initial configuration $\mathbf{c} = (c_1, \ldots, c_k)$ is such that $\max\{c_j : j = 1, \ldots, k\} \leqslant 3n/(2k)$ then $\tau = \Omega(k/h^2)$ w.h.p.*

Proof. Since in the initial configuration for any color $j \in [k]$ we have that $c_j \leqslant 3n/(2k)$, from Lemma 4.11 it follows that the number of nodes supporting the plurality color increases at a rate smaller than $(1 + 2h^2/k)$ with probability exponentially close to 1. This easily implies a recursive relation of the form $C_j^{(t+1)} \leqslant \left(1 + 2h^2/k\right) C_j^{(t)}$ which, in turn, gives

$$C_j^{(t)} \leqslant \left(1 + \frac{2h^2}{k}\right)^t C_j^{(0)} \leqslant \left(1 + \frac{2h^2}{k}\right)^t \frac{3n}{2k}.$$

Thus, for $t < k/h^2 \log(4/3)$, w.h.p. we have that

$$C_j^{(t)} \leqslant \frac{3n}{2k}\left(1 + \frac{2h^2}{k}\right)^t < \frac{2n}{k},$$

concluding the proof. □

4.4 On the initial bias

In this section, we show that there are initial configurations with bias $s = \mathcal{O}(\sqrt{kn})$ for which the bias decreases in a single round with constant probability. This shows that under initial imbalances of this magnitude, it seems unlikely that one can prove bounds as those shown in Section 3, at least with high probability.

Lemma 4.13 *Assume $k \geqslant 4$. For any value $s \leqslant \sqrt{kn}/6$ of the initial bias, there are initial configurations* $\mathbf{c}$ *such that, for any fixed color $j \neq 1$ we have:*

$$\mathbf{P}\left(C_1^{(1)} - C_j^{(1)} < s | \mathbf{C}^{(0)} = \mathbf{c}\right) \geqslant \frac{1}{16e}.$$

Proof. We consider an initial configuration $\mathbf{c}$ in which we have k colors. Let $x = (n-s)/k$ (we neglect integer parts for the sake of the analysis). We let $c_1^{(0)} = x + s$ and $c_j^{(0)} = x$, for $j \neq 1$ and we further assume that $s \leqslant x$. Considered any fixed $j \neq 1$, we next prove that $C_1^{(1)} - C_j^{(1)} < s$ with constant probability.

The outline of the proof is as follows: We first show that $\mathbf{E}\left[C_1^{(1)} \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right] - \mathbf{E}\left[C_j^{(1)} \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right] \leqslant s + 3xs/n$, then we observe that with constant probablity $\mathbf{E}\left[C_1^{(1)} \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right]$ is not above its average. Finally, we prove that $C_j^{(1)} > \mathbf{E}\left[C_j^{(1)} \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right] + 3xs/n$ with constant probability, whenever $s \leqslant \sqrt{kn}/6$, which concludes the proof of the lemma.

To begin with, from Lemma 2.1, we easily get the following derivations:

$$\begin{aligned}
\mathbf{E}\left[C_1^{(1)} \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right] &= x + s + \frac{x^2}{n} + \frac{2xs + s^2}{n} - \frac{x+s}{n^2}\gamma \quad \text{and} \\
\mathbf{E}\left[C_j^{(1)} \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right] &= x + \frac{x^2}{n} - \frac{x}{n^2}\gamma,
\end{aligned}$$

where $\gamma = \sum_h x_h^2 = nx + xs + s^2$.

Next, from the definition of γ we get

$$\begin{aligned}
\mathbf{E}\left[C_1^{(1)} \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right] &= x + s + \frac{x^2}{n} + \frac{2xs + s^2}{n} - \frac{x+s}{n^2}(nx + xs + s^2) \\
&= x + s + \frac{xs}{n} + \frac{s^2}{n} - \frac{s}{n^2}(x+s)^2 \\
&\leqslant x + s + \frac{2xs}{n} - \frac{s}{n^2}(x+s)^2 < x + s + \frac{2xs}{n},
\end{aligned}$$

where the third inequality follows by assuming $s \leqslant x$. Analogously we have:

$$\begin{aligned}
\mathbf{E}\left[C_j^{(1)} \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right] &= x + \frac{x^2}{n} - \frac{x}{n^2}(nx + xs + s^2) \\
&= x - \frac{xs}{n} \cdot \frac{x+s}{n} \geqslant x - \frac{2x^2 s}{n^2},
\end{aligned} \tag{22}$$

where to derive the last inequality we again use $s \leqslant x$. As a consequence we can write:

$$\mathbf{E}\left[C_1^{(1)} \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right] - \mathbf{E}\left[C_j^{(1)} \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right] \leqslant s + \frac{2xs}{n} - \frac{2x^2 s}{n} \leqslant s + \frac{3xs}{n},$$

where the second inequality holds whenever $x \leqslant n/2$, which is our case.

For convenience sake, let us name for any $j \in [k]$

$$\mu_j := \mathbf{E}\left[C_j^{(1)} \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right]$$

We note that

$$\mathbf{P}\left(C_1^{(1)} - C_j^{(1)} < s \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right) \geqslant \mathbf{P}\left(C_1^{(1)} < \mu_1 \wedge C_j^{(1)} \geqslant \mu_j + \frac{3xs}{n} \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right).$$

Fact 1 *The following holds:*

$$\mathbf{P}\left(C_1^{(1)} < \mu_1 \wedge C_j^{(1)} \geqslant \mu_j + \frac{3xs}{n} \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right) \geqslant \mathbf{P}\left(C_1^{(1)} < \mu_1 \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right) \mathbf{P}\left(C_j^{(1)} \geqslant \mu_j + \frac{3xs}{n} \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right).$$

Proof. We have:

$$\begin{aligned}
&\mathbf{P}\left(C_1^{(1)} < \mu_1 \wedge C_j^{(1)} \geqslant \mu_j + \frac{3xs}{n} \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right) \\
&= \mathbf{P}\left(C_1^{(1)} < \mu_1 \,\Big|\, C_j^{(1)} \geqslant \mu_j + \frac{3xs}{n}\right) \mathbf{P}\left(C_j^{(1)} \geqslant \mu_j + \frac{3xs}{n} \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right) \\
&= \left(1 - \mathbf{P}\left(C_1^{(1)} \geqslant \mu_1 \,\Big|\, C_j^{(1)} \geqslant \mu_j + \frac{3xs}{n}\right)\right) \mathbf{P}\left(C_j^{(1)} \geqslant \mu_j + \frac{3xs}{n} \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right) \\
&= \mathbf{P}\left(C_j^{(1)} \geqslant \mu_j + \frac{3xs}{n} \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right) - \mathbf{P}\left(C_1^{(1)} \geqslant \mu_1 \wedge C_j^{(1)} \geqslant \mu_j + \frac{3xs}{n} \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right) \\
&\geqslant \mathbf{P}\left(C_j^{(1)} \geqslant \mu_j + \frac{3xs}{n} \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right) - \mathbf{P}\left(C_1^{(1)} < \mu_1 \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right) \mathbf{P}\left(C_j^{(1)} \geqslant \mu_j + \frac{3xs}{n} \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right) \\
&= \mathbf{P}\left(C_j^{(1)} \geqslant \mu_j + \frac{3xs}{n} \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right) \\
&\qquad - \left(1 - \mathbf{P}\left(C_1^{(1)} < \mu_1 \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right)\right) \mathbf{P}\left(C_j^{(1)} \geqslant \mu_j + \frac{3xs}{n} \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right) \\
&= \mathbf{P}\left(C_1^{(1)} < \mu_1 \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right) \mathbf{P}\left(C_j^{(1)} \geqslant \mu_j + \frac{3xs}{n} \,\Big|\, \mathbf{C}^{(0)} = \mathbf{c}\right),
\end{aligned}$$

where the fourth inequality follows from [9, Proposition 3, claim (-OD)]. In particular, $\left\{C_1^{(1)} \geqslant \mu_1\right\}$ and $\left\{C_j^{(1)} \geqslant \mu_j + 3xs/n\right\}$ are the events that the numbers of balls thrown independently at random into two distinct bins both exceed some given thresholds. □

Fact 2 *The following holds:*

$$\mathbf{P}\left(C_1^{(1)} < \mu_1 \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right) > \frac{1}{4}.$$

Proof. Set $X = n - C_1^{(1)}$. Clearly, X is distributed as a binomial $B(n, p)$, where $p = 1 - p_1$, with p_1 the probability that the generic node recolors itself with color 1. Clearly, $p > 1/n$ as long as the number of colors is not too large (in the order of n). Then we have:

$$\mathbf{P}\left(C_1^{(1)} < \mu_1 \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right) = \mathbf{P}\left(X \geqslant \mathbf{E}[X] \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right) > \frac{1}{4},$$

where the second inequality follows from [11, Theorem 1]. □

We finally apply Theorem A.3 to $C_j^{(1)}$ and we have:

$$\mathbf{P}\left(C_j^{(1)} > \mu_j + \frac{3xs}{n} \,|\, \mathbf{C}^{(0)} = \mathbf{c}\right) \geqslant \frac{1}{4} e^{-\frac{18x^2s^2}{n^2\mu_j}} \geqslant \frac{1}{4} e^{-\frac{18xs^2}{n^2-2xs}} \geqslant \frac{1}{4e}, \tag{23}$$

where the second inequality follows from (22) and the third one holds since $s \leqslant \sqrt{kn}/6$ and recalling that $x \leqslant n/k$. Finally, from Fact 2 and (23), we get the claim. □

5 Open Questions

A general open question on the plurality consensus problem is whether an *efficient* dynamics exists that achieves plurality consensus in polylogarithmic time for any function $k = k(n)$. By *efficient* dynamics for our adopted model, we mean any dynamics that requires small (i.e. $\mathcal{O}(\log n)$) memory, small random samples, and small message size.

A more specific question about our simple distributed model is to explore the case in which the initial bias s is smaller than the lower bound assumed in our analysis (i.e. $s \geqslant c\sqrt{\min\{2k, (n/\log n)^{1/3}\}\, n \log n}$). Notice that when k is polylogarithmic, we required a bias which is only a polylogarithmic factor larger than the standard deviation $\Omega(\sqrt{n})$: the latter is a lower bound for the initial bias to converge (w.h.p.) to the plurality color. As for larger k, we did not derive any stronger bound on the required bias, however, in Subsection 4.4, we have shown some initial configurations with bias $s = \mathcal{O}(\sqrt{kn})$ for which the initial bias *decreases* in a single round with constant probability. This result implies that, when the initial bias s is "slightly" smaller than "ours", the process may be *non-monotone* w.r.t. the bias function $s(t)$. The fact that $s(t)$ is an increasing function played a key-role in the proof of our upper bound. So, under such a weaker assumption, if any upper bound similar to ours might be proved then a much more complex argument (departing from ours) seems to be necessary.

In this work, we were interested in deriving sufficient conditions under which the h-plurality dynamics converges in polylogarithmic time. A further interesting open question is to derive conditions on the parameters k, s, and h under which this dynamics converges very fast, i.e., in sublogarithmic time.

References

[1] M. A. Abdullah and M. Draief. Global majority consensus by local majority polling on graphs of a given degree sequence. *Discrete Applied Mathematics*, 180:1–10, 2015.

[2] D. Angluin, J. Aspnes, and D. Eisenstat. A Simple Population Protocol for Fast Robust Approximate Majority. *Distributed Computing*, 21(2):87–102, 2008. (Preliminary version in DISC'07).

[3] A. Babaee and M. Draief. Distributed multivalued consensus. In *Proc. of Computer and Information Sciences III*, pages 271–279. Springer, 2013.

[4] L. Becchetti, A. Clementi, E. Natale, F. Pasquale, and R. Silvestri. Plurality Consensus in the Gossip Model. In *Proc. of the 26th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'15)*, pages 371–390. SIAM, 2015.

[5] F. Bénézit, P. Thiran, and M. Vetterli. Interval consensus: from quantized gossip to voting. In *Proc. of IEEE Intern. Conf. on Acoustics, Speech and Signal Processing (ICASSP'09)*, pages 3661–3664. IEEE, 2009.

[6] A. Clementi, M Di Ianni, G. Gambosi, E. Natale, and R. Silvestri. Distributed community detection in dynamic graphs. *Theoretical Computer Science*, 584:19–41, 2015.

[7] B. Doerr, L. A. Goldberg, L. Minder, T. Sauerwald, and C. Scheideler. Stabilizing consensus with the power of two choices. In *Proc. of the 23rd Ann. ACM Symp. on Parallelism in Algorithms and Architectures (SPAA'11)*, pages 149–158. ACM, 2011.

[8] M. Draief and M. Vojnovic. Convergence speed of binary interval consensus. *SIAM J. on Control and Optimization*, 50(3):1087–1109, 2012.

[9] D. Dubhashi and D. Ranjan. Balls and bins: A study in negative dependence. *Random Structures and Algorithms*, 13(2):99–124, 1998.

[10] D. Easley and J. Kleinberg. *Networks, Crowds, and Markets.* Cambridge University Press, 2010.

[11] S. Greenberg and M. Mohri. Tight lower bound on the probability of a binomial exceeding its expectation. *Statistics & Probability Letters*, 86:91–98, 2014.

[12] M. Kearns and J. Tan. Biased voting and the democratic primary problem. In *Proc. of the 4th Workshop on Internet and Network Economics*, volume 5385 of *Lectures Notes in Computer Science*, pages 639–652. Springer, 2008.

[13] M. W. S. Land and R. K. Belew. No two-state ca for density classification exists. *Phys. Rev. Letters*, 74(25):5148–5150, 1995.

[14] D. Levin, Y. Peres, and E. L. Wilmer. *Markov Chains and Mixing Times.* AMS, 2008.

[15] E. Mossel, J. Neeman, and O. Tamuz. Majority dynamics and aggregation of information in social networks. *Autonomous Agents and Multi-Agent Systems*, 28(3):408–429, 2014.

[16] E. Mossel and G. Schoenebeck. Reaching Consensus on Social Networks. In *Proc. of the 2nd Innovations in Computer Science (ICS'10)*, pages 214–229, 2010.

[17] N. Mousavi. How tight is chernoff bound? url: https://ece.uwaterloo.ca/ nmousavi/Papers/Chernoff-Tightness.pdf.

[18] D. Peleg. Local majorities, coalitions and monopolies in graphs: a review. *Theoretical Computer Science*, 282(2):231–257, 2002.

[19] E. Perron, D. Vasudevan, and M. Vojnovic. Using three states for binary consensus on complete graphs. In *Proc. of the 28th IEEE INFOCOM*, pages 2527–2535, 2009.

A Useful Bounds

Lemma A.1 (Chernoff bounds) *Let* $X = \sum_{i=1}^{n} X_i$ *where* X_i*'s are independent Bernoulli random variables and let* $\mu = \mathbf{E}[X]$. *Then,*

1. *For any* $0 < \delta \leqslant 4$, $\mathbf{P}(X > (1+\delta)\mu) < e^{-\frac{\delta^2\mu}{4}}$;
2. *For any* $\delta \geqslant 4$, $\mathbf{P}(X > (1+\delta)\mu) < e^{-\delta\mu}$;
3. *For any* $\lambda > 0$, $\mathbf{P}(X \geqslant \mu + \lambda) \leqslant e^{-2\lambda^2/n}$.

Lemma A.2 (Jensen inequality) *Let* $\phi : \mathbb{R} \to \mathbb{R}$ *be a convex function and* $x_1, \ldots x_k \in \mathbb{R}$ *be* k *real numbers, then*

$$\phi\left(\frac{1}{k}\sum_{i=1}^{k} x_i\right) \leqslant \frac{1}{k}\sum_{i=1}^{k}\phi(x_i).$$

In Section 4.4, we use the following "reverse"-Chernoff bound [17, Theorem 2] [6]

Theorem A.3 (Reverse Chernoff bound) *Let* X *be the sum of* m *independent Bernoulli variables with probability* $p \leqslant 1/4$ *and let* $\mu = pm$. *Then, for any* $t > 0$*:*

$$\mathbf{P}(X - \mu > t) \geqslant \frac{1}{4}e^{-\frac{2t^2}{\mu}}.$$

[6] A number of pretty similar "folklore" results can be found in specialized mathematical forums, for example `http://cstheory.stackexchange.com/questions/14471/reverse-chernoff-bound`.