Simple Dynamics for Plurality Consensus*

Luca Becchetti
"Sapienza" Università di Roma
becchett@dis.uniroma1.it

Andrea Clementi
Università "Tor Vergata" di Roma
clementi@mat.uniroma2.it

Emanuele Natale
"Sapienza" Università di Roma
natale@di.uniroma1.it

Francesco Pasquale
"Sapienza" Università di Roma
pasquale@di.uniroma1.it

Riccardo Silvestri
"Sapienza" Università di Roma
silvestri@di.uniroma1.it

Luca Trevisan
Stanford University
trevisan@stanford.edu

ABSTRACT

We study a *Plurality Consensus* process in which each of n anonymous agents of a communication network supports an initial opinion (a color chosen from a finite set $[k]$) and, at every time step, he can revise his color according to a random sample of neighbors. The goal (of the agents) is to let the process converge to the *stable* configuration where all nodes support the *plurality color*. It is assumed that the initial color configuration has a sufficiently large *bias* s, that is, the number of nodes supporting the plurality color exceeds the number of nodes supporting any other color by an additive value s.

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 that each agent looks at the colors of three random neighbors and then applies the majority rule (breaking ties uniformly at random).

We prove a tight bound on the convergence time which grows as $\Theta(k \log n)$ for a wide range of parameters k and n. This *linear-in-k* dependence implies an exponential time-gap between the plurality consensus process and the *median* process studied in [7].

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.

Categories and Subject Descriptors

F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity

*Partially supported by Italian MIUR-PRIN 2010-11 Project *ARS TechnoMedia* and the EU FET Project *MULTIPLEX* 317532.

http://dx.doi.org/10.1145/2612669.2612677.

Keywords

Plurality Consensus; Parallel Randomized Algorithms; Markov Chains.

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 from several areas such as Distributed Computing [7, 16], Communication Networks [17], and Social Networks [6, 15, 14]. 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, at every time step, 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 that each node samples at random three neighbors, and picks the majority color among them (breaking ties uniformly at random). Let us 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 has been presented: the goal here is to converge to a stable configuration where all nodes support a value which is a "good" approximation of the *median* of the initial color configuration. It turns out that, in the binary case (i.e $k = 2$), the median problem is equivalent to the majority consensus one 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 are different from each other (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, 1].

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 $\mathcal{O}(\lambda(n) \log n)$ w.h.p., no matter how large k is. Hence, when $c_m \geqslant n/\mathrm{polylog}(n)$ and $s \geqslant \sqrt{n\, \mathrm{polylog}(n)}$, 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 in fact 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 $2\,n/k$.

We then prove a general negative result: in the considered distributed model, there is no dynamics with at most 3 inputs (but the majority one) that w.h.p. converges to plurality consensus starting from any initial bias s such that $s = o(n)$. In other words, not only there is no hope to find a 3-input dynamics faster than $k \log n$ but the 3-majority dynamics is the only one getting the plurality consensus, no matter in how much time. 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 the task of computing the median and that of computing the plurality (this happens for instance when $k = n^a$, for any constant $0 < a < 1/4$).

A natural question arising from our results is whether a (slightly) larger random sample of neighbors might lead to a significant speed-up of the convergence time to plurality consensus. We provide a negative answer to this question. We consider the generalization of the 3-majority dynamics, the h-plurality one, where every node, at every time step, updates his color according to the plurality of the colors supported by h random 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 $k/h = \mathcal{O}\left(n^{1/4-\epsilon}\right)$, where ϵ is an arbitrarily-small positive constant. We emphasize that scalable and efficient protocols must yield low communication complexity and small node congestion at every time step. These properties are guaranteed by the h-plurality dynamics only when h is small, say $h = \mathcal{O}(\mathrm{polylog}(n))$: in this case, our lower bound says 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 starts. 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 $\mathrm{poly}(n)$ 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 works. 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, 16]. The objective here is to converge to the version of the data supported by the majority of the initial distributed copies (it is reasonable that a sufficiently large majority of the nodes are not faulty and thus have the correct data). Another application comes from the task of distributed item ranking where every node initially has ranked some item and the goal is to agree on the rank of the item based on the initial majority opinion [17]. Further applications of majority updating rules in networks can be found in [10, 16].

The results most related to our contribution are those in [7] which have been already discussed above. Several variants of the binary majority consensus have been studied in different distributed models [2, 15]. 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 multivalued case ($k \geqslant 3$) does not converge to plurality even starting from large bias $s = \Theta(n)$. The polling rule (a somewhat sequential-interaction version of the 1-majority dynamics) has been extensively studied on

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

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

several classes of graphs (see [16]). More expensive and complex protocols have been considered in order to speed up the process. For instance, in [11], 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, 12] (with no time bound). In [1, 3, 8, 17], the polling rule (with 1 more auxiliary state) 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 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. Protocols for specific network topologies and some "social-based" communities have been studied in [1, 8, 14, 17].

Roadmap of the paper. 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. In the Appendix A, we recall some standard results (such as Chernoff-Bernstein's inequalities) and provide a useful probabilistic result on Markov chains (we have not found its explicit proof in the literature). Due to lack of space, several proofs are omitted. They can be found in the full-version of the paper [4].

2. PRELIMINARIES

A *k-color distribution* (for short *k-cd*) is any k-tuple $\bar{c} = (c_1, \ldots, c_k)$ such that c_js are non negative integers and $\sum_{j=1,\ldots,k} c_j = n$. A color m is said to be a *plurality color* of $\bar{c}$ if $c_m \geqslant c_j$ for every other color $j \in [k] \setminus \{m\}$. We say that $\bar{c}$ is s-biased if a color m exists such that $c_m \geqslant c_j + s$ for every other color $j \in [k] \setminus \{m\}$.

The 3-majority protocol works as follows:

> *At every time step, 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 time step t and for any $j \in [k]$, let $C_{j,t}$ be the r.v. counting the number of nodes colored j at time step t and let $C_t = (C_{1,t}, \ldots, C_{k,t})$ 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 $\mu_j(\bar{c})$ be the expected number of nodes with color j at the next step when the current k-cd is $\bar{c}$, i.e. $\mu_j(\bar{c}) = \mathbb{E}\left[C_{j,t+1} \,|\, C_t = \bar{c}\right]$. The proof of the next lemma is a straightforward computation.

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

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

3. UPPER BOUNDS FOR 3-MAJORITY

In this section, we show an upper bound on the convergence time of the 3-majority dynamics that holds with high probability. To this aim, we need to consider the following r.v.s For a k-cd $\bar{c}$, we define

$$m(\bar{c}) = \max_{h\in[k]} c_h$$
$$M(\bar{c}) = \{ j \in [k] \mid c_j = m(\bar{c}) \}$$
$$s(\bar{c}) = \begin{cases} m(\bar{c}) - \max_{h\in[k]-M(\bar{c})} c_h & \text{if } |M(\bar{c})| = 1 \\ 0 & \text{otherwise} \end{cases}$$
$$\alpha(\bar{c}) = \frac{(n - m(\bar{c}))s(\bar{c})}{n^2}$$
$$\gamma(\bar{c}) = \frac{n \cdot m(\bar{c}) - \sum_{h\in[k]} c_h^2}{n^2} - \alpha(\bar{c})$$

The next lemma gives some useful inequalities relating the above quantities.

Lemma 3.1. *For any k-cd $\bar{c}$, the followings hold*

a) $0 \leqslant s(\bar{c}) \leqslant m(\bar{c}) - \frac{n-m(\bar{c})}{k-1}$

b) $0 \leqslant \alpha(\bar{c}) \leqslant \min\{\frac{s(\bar{c})}{n}, \frac{1}{4}\}$

c) $0 \leqslant \gamma(\bar{c}) \leqslant \frac{1}{8}$

The above lemma allows us to give a new expression for $\mu_j(\bar{c})$ that will be useful in the proofs of Lemmas 3.3 and 3.4.

Lemma 3.2. *Let $\bar{c}$ be any k-cd. Let m be any color in $M(\bar{c})$ and let $\ell \in [k]-\{m\}$ be such that $c_\ell = \max_{h\in[k]-\{m\}} c_h$.*

a) $\mu_m(\bar{c}) = c_m(1 + \gamma(\bar{c}) + \alpha(\bar{c}))$

b) $\forall j \in [k]-M(\bar{c}) \quad \mu_j(\bar{c}) = c_j\left(1 + \gamma(\bar{c}) + \alpha(\bar{c}) - \frac{m(\bar{c})-c_j}{n}\right)$

c) $\mu_\ell(\bar{c}) = c_\ell\left(1 + \gamma(\bar{c}) + \alpha(\bar{c}) - \frac{s(\bar{c})}{n}\right)$

We now evaluate the increasing rate of the bias of a k-cd during a generic step of the 3-majority dynamics.

Lemma 3.3 (increasing rate of the bias). *Let $\bar{c}$ be any k-cd such that $M(\bar{c}) = \{m\}$ for some $m \in [k]$. Then it holds that, for any $j \in [k] - m$,*

$$\mathbf{P}\left(C_{m,t+1} - C_{j,t+1} \leqslant s(\bar{c})\left(1 + \gamma(\bar{c}) + \frac{c_m\alpha(\bar{c})}{2s(\bar{c})}\right) \,\middle|\, C_t = \bar{c}\right) \leqslant \exp\left(-\frac{c_m\alpha(\bar{c})^2}{25}\right) \quad (1)$$

This is the key-lemma to get our upper bound on the convergence time so, before giving its proof, let us provide a rough but useful evaluation of Eq. 1 for a fixed setting of parameters k and s, i.e., $k = n^{1/4}$ and $s = c\sqrt{n^{3/4}\log n}$, for some constant $c > 0$. Consider the "initial phase" of the coloring process where c_m is still $\Theta(n/k) = \Theta(n^{3/4})$ and s is still $o(c_m)$. Then, by replacing the values of $\alpha(\bar{c})$ and $\gamma(\bar{c})$ in Eq. 1 (and doing some simple calculations), we get that the bias s increases by a factor $1 + \Theta(1/k)$ w.h.p. This is exactly what we need to get the upper bound $\mathcal{O}(k \log n)$ on the convergence time. The bound in Eq. 1 has a more complex, general shape since it must work for the whole process and it must lead to our stronger bound $\mathcal{O}\left(\min\{k, (n/\log n)^{1/3}\} \log n\right)$.

PROOF. (of Lemma 3.3)
In the sequel we tacitly assume that all probabilities, expected values and random variables are conditioned to "$C_t = \overline{c}$". For any fixed color $j \in [k] - \{m\}$, we consider the the random variable

$$Z = C_{m,t+1} - C_{j,t+1}$$

It holds that

$$\begin{aligned}
\mathbb{E}[Z] &= \mu_m(\overline{c}) - \mu_j(\overline{c}) \\
&= c_m(1+\gamma(\overline{c})+\alpha(\overline{c})) - c_j\left(1+\gamma(\overline{c})+\alpha(\overline{c}) - \frac{m(\overline{c})-c_j}{n}\right) \\
&\qquad \text{(from Lemma 3.2)} \\
&= (c_m - c_j)(1+\gamma(\overline{c})) + c_m\alpha(\overline{c}) + c_j\left(\frac{m(\overline{c})-c_j}{n} - \alpha(\overline{c})\right) \\
&= (c_m - c_j)(1+\gamma(\overline{c})) + c_m\alpha(\overline{c}) + c_j\left(\frac{m(\overline{c})-c_j}{n} - \alpha(\overline{c})\right) \\
&\geqslant (c_m - c_j)(1+\gamma(\overline{c})) + c_m\alpha(\overline{c}) + c_j\left(\frac{s(\overline{c})}{n} - \alpha(\overline{c})\right) \\
&\qquad \text{(since } m(\overline{c}) - c_j \geqslant s(\overline{c})\text{)} \\
&\geqslant (c_m - c_j)(1+\gamma(\overline{c})) + c_m\alpha(\overline{c}) \\
&\qquad \text{(since } \tfrac{s(\overline{c})}{n} \geqslant \alpha(\overline{c}) \text{ by Lemma 3.1)} \\
&\geqslant s(\overline{c})\left(1+\gamma(\overline{c}) + \frac{c_m\alpha(\overline{c})}{s(\overline{c})}\right) \qquad (2) \\
&\qquad \text{(since } c_m - c_j \geqslant s(\overline{c})\text{)}
\end{aligned}$$

We now introduce, for any $i \in [n]$, the random variable

$$Z_i = \begin{cases} 1 & \text{if node } i \text{ gets color } m \text{ at time } t+1 \\ -1 & \text{if node } i \text{ gets color } j \text{ at time } t+1 \\ 0 & \text{otherwise} \end{cases}$$

Clearly, the Z_i's are independent and it holds that

$$Z = \sum_{i \in [n]} Z_i$$

In order to apply the Bernstein's Inequality (Lemma A.3) to $-Z$, we firstly observe that

$$-Z_i - \mathbb{E}[-Z_i] \leqslant 2\,,$$

so we can choose $b = 2$. As for the variance σ^2 of $-Z$, we have that

$$\begin{aligned}
\sigma^2 = \mathbf{Var}[-Z] &= \sum_{i\in[n]} \mathbf{Var}[-Z_i] \\
&= \sum_{i\in[n]} (\mathbb{E}[(-Z_i)^2] - \mathbb{E}[-Z_i]^2) = \sum_{i\in[n]} (\mathbb{E}[Z_i^2] - \mathbb{E}[Z_i]^2) \\
&\leqslant \sum_{i\in[n]} \mathbb{E}[Z_i^2] = \sum_{i\in[n]} (\mathbf{P}(Z_i = 1) + \mathbf{P}(Z_i = -1)) \\
&= \mu_m(\overline{c}) + \mu_j(\overline{c}) \\
&\leqslant 2\mu_m(\overline{c}) \qquad \text{(since } \mu_j(\overline{c}) \leqslant \mu_m(\overline{c}) \text{ by Lemma 3.2)} \\
&= 2c_m(1+\gamma(\overline{c})+\alpha(\overline{c})) \\
&\leqslant 2c_m(1 + \frac{1}{8} + \frac{1}{4}) \qquad \text{(from Lemma 3.1)} \\
&\leqslant 3c_m \qquad (3)
\end{aligned}$$

For the sake of convenience, let us define

$$P = \mathbf{P}\left(Z < s(\overline{c})\left(1+\gamma(\overline{c}) + \frac{c_m\alpha(\overline{c})}{2s(\overline{c})}\right)\right)$$

Now we conclude the proof by applying the Bernstein's Inequality

$$\begin{aligned}
P &= \mathbf{P}\left(-Z > -s(\overline{c})\left(1+\gamma(\overline{c}) + \frac{c_m\alpha(\overline{c})}{s(\overline{c})}\right) + \frac{c_m\alpha(\overline{c})}{2}\right) \\
&\leqslant \mathbf{P}\left(-Z > \mathbb{E}[-Z] + \frac{c_m\alpha(\overline{c})}{2}\right) \\
&\qquad \text{(since } \mathbb{E}[-Z] \leqslant -s(\overline{c})\left(1+\gamma(\overline{c}) + \tfrac{c_m\alpha(\overline{c})}{s(\overline{c})}\right) \text{ by Ineq. 2)} \\
&\leqslant \exp\left(-\frac{\left(\frac{c_m\alpha(\overline{c})}{2}\right)^2}{2\sigma^2 + (4/3)\frac{c_m\alpha(\overline{c})}{2}}\right) \\
&\qquad \text{(from Lemma A.3 with } b = 2 \text{ and } \lambda = \tfrac{c_m\alpha(\overline{c})}{2}\text{)} \\
&\leqslant \exp\left(-\frac{c_m^2\alpha(\overline{c})^2}{24c_m + (8/3)c_m\alpha(\overline{c})}\right) \qquad \text{(from Ineq. 3)} \\
&\leqslant \exp\left(-\frac{c_m\alpha(\overline{c})^2}{25}\right) \qquad \text{(since } \alpha(\overline{c}) \leqslant 1/4 \text{ by Lemma 3.1)}
\end{aligned}$$

□

The next lemma derives from Lemmas 3.1 and 3.2.

LEMMA 3.4. *Let $\overline{c}$ be any k-cd such that $M(\overline{c}) = \{m\}$ for some $m \in [k]$. It holds that*

$$\mathbf{P}\left(C_{m,t+1} \leqslant c_m\left(1+\gamma(\overline{c}) + \frac{\alpha(\overline{c})}{2}\right) \middle| C_t = \overline{c}\right) \leqslant \exp\left(-\frac{c_m\alpha(\overline{c})^2}{11}\right)$$

PROOF. Let

$$P_m = \mathbf{P}\left(C_{m,t+1} \leqslant c_m\left(1+\gamma(\overline{c}) + \frac{\alpha(\overline{c})}{2}\right) \middle| C_t = \overline{c}\right)$$

and let

$$\delta_m = \frac{\alpha(\overline{c})}{2(1+\gamma(\overline{c})+\alpha(\overline{c}))}$$

From Lemma 3.1, $\gamma, \alpha \geqslant 0$, and thus $0 < \delta_m < 1$. Thanks to the Chernoff bound we have that

$$\begin{aligned}
P_m &= \mathbf{P}(C_{m,t+1} \leqslant (1-\delta_m)\mu_m \mid C_t = \overline{c}) \qquad \text{(from Lemma 3.2)} \\
&\leqslant \exp\left(-\frac{\delta_m^2\mu_m}{2}\right) \qquad \text{(by the Chernoff bound)} \\
&= \exp\left(-\frac{1}{2}\left(\frac{\alpha(\overline{c})}{2(1+\gamma(\overline{c})+\alpha(\overline{c}))}\right)^2 c_m(1+\gamma(\overline{c})+\alpha(\overline{c}))\right) \\
&= \exp\left(-\frac{c_m\alpha(\overline{c})^2}{8(1+\gamma(\overline{c})+\alpha(\overline{c}))}\right) \\
&\leqslant \exp\left(-\frac{c_m\alpha(\overline{c})^2}{11}\right) \\
&\qquad \text{(since } \gamma(\overline{c}) + \alpha(\overline{c}) \leqslant 3/8 \text{ by Lemma 3.1)}
\end{aligned}$$

□

We now use Lemmas 3.3 and 3.4 in order to get some bounds on the increasing rate of the bias: they will lead to a bound on convergence time that does not depend on k.

LEMMA 3.5 (LARGE PLURALITY AND LARGE BIAS). *Let $\overline{c}$ be any k-cd such that $M(\overline{c}) = \{m\}$ for some $m \in [k]$. For any value λ with $0 < \lambda \leqslant 2/3$, if $\lambda n \leqslant c_m \leqslant (2/3)n$ and $s(\overline{c}) \geqslant 22\sqrt{(1/\lambda)n\log n}$, then, for every $j \in [k] - \{m\}$,*

$$\mathbf{P}\left(C_{m,t+1} - C_{j,t+1} \leqslant s(\overline{c})\left(1+\frac{\lambda}{6}\right) \middle| C_t = \overline{c}\right) \leqslant \frac{1}{n^2}$$

and

$$\mathbf{P}\left(C_{m,t+1} \leqslant c_m \mid C_t = \bar{c}\right) \leqslant \frac{1}{n^2}$$

PROOF. From Lemma 3.3 we have that

$$\mathbf{P}\left(C_{m,t+1} - C_{j,t+1} \leqslant s(\bar{c})\left(1 + \gamma(\bar{c}) + \frac{c_m\alpha(\bar{c})}{2s(\bar{c})}\right) \,\middle|\, C_t = \bar{c}\right) \leqslant \exp\left(-\frac{c_m\alpha(\bar{c})^2}{25}\right) \quad (4)$$

It holds that

$$\frac{c_m\alpha(\bar{c})}{2s(\bar{c})} = \frac{c_m(n-c_m)}{2n^2} \geqslant \frac{\lambda n(n/3)}{2n^2} \geqslant \frac{\lambda}{6} \quad (5)$$

As regards the exponent of the probability bound of Ineq. 4 we get

$$\begin{aligned}\frac{c_m\alpha(\bar{c})^2}{25} &= \frac{c_m(n-c_m)^2 s(\bar{c})^2}{25n^4} \\ &\geqslant \frac{c_m s(\bar{c})^2}{225n^2} \qquad (\text{since } c_m \leqslant (2/3)n) \\ &\geqslant \frac{\lambda n 484(1/\lambda)n\log n}{225n^2} \\ &\qquad (\text{by the hypothesis bounds on } c_m \text{ and } s(\bar{c})) \\ &\geqslant 2\log n \qquad (6)\end{aligned}$$

By combining Ineq.s 4, 5, and 6 we obtain the first probability bound. As for the second bound, from Lemma 3.4 it holds that

$$\begin{aligned}&\mathbf{P}\left(C_{m,t+1} \leqslant c_m \mid C_t = \bar{c}\right) \\ &\leqslant \mathbf{P}\left(C_{m,t+1} \leqslant c_m\left(1 + \gamma(\bar{c}) + \frac{\alpha(\bar{c})}{2}\right) \,\middle|\, C_t = \bar{c}\right) \\ &\leqslant \exp\left(-\frac{c_m\alpha(\bar{c})^2}{11}\right)\end{aligned}$$

and, from Ineq. 6,

$$\frac{c_m\alpha(\bar{c})^2}{11} \geqslant \frac{c_m\alpha(\bar{c})^2}{25} \geqslant 2\log n$$

□

For any $m \in [k]$, let $\overline{C}_{m,t} = n - C_{m,t}$ denote the random variable counting the number of nodes with colors different from m at time t. For any k-cd $\bar{c}$ and for any $m \in [k]$, we also consider its expected value $\bar{\mu}_m(\bar{c}) = \mathbb{E}\left[\overline{C}_{m,t+1} \,\middle|\, C_t = \bar{c}\right]$, and provide the following bounds

LEMMA 3.6. *For any k-cd $\bar{c}$ and for any $m \in M(\bar{c})$, it holds that*

$$(n - c_m)\left(1 - \frac{c_m^2}{n^2}\right) \leqslant \bar{\mu}_m(\bar{c}) \leqslant (n - c_m)\left(1 - \frac{s(\bar{c})c_m}{n^2}\right)$$

We now use the above bounds to show that, when the bias of a k-cd is at least $n/3$, then the number of nodes that do not have the plurality color decreases at an exponential rate.

LEMMA 3.7 (VERY-LARGE PLURALITY). *Let $\bar{c}$ be any k-cd such that $s(\bar{c}) \geqslant n/3$ for some $m \in [k]$. If $n - c_m \geqslant \sqrt[4]{n}\log n$, then*

$$\mathbf{P}\left(\overline{C}_{m,t+1} \geqslant \frac{17}{18}(n - c_m) \,\middle|\, C_t = \bar{c}\right) \leqslant \frac{1}{n^2}$$

If $n - c_m < \sqrt[4]{n}\log n$, then

$$\mathbf{P}\left(\overline{C}_{m,t+1} > 0 \,\middle|\, C_t = \bar{c}\right) \leqslant \frac{1}{\sqrt[5]{n}}$$

$$\mathbf{P}\left(\overline{C}_{m,t+1} \geqslant \sqrt[4]{n}\log n \,\middle|\, C_t = \bar{c}\right) \leqslant \frac{1}{n^2}$$

We now exploit Lemmas 3.5 and 3.7 in order to prove the main result of this section.

THEOREM 3.8 (THE GENERAL UPPER BOUND). *Let λ be any value such that $3 \leqslant \lambda < \sqrt{n}$. If $\bar{c}$ is a k-cd such that, for some $m \in [k]$, $M(\bar{c}) = \{m\}$, $c_m \geqslant n/\lambda$, and $s(\bar{c}) \geqslant 22\sqrt{\lambda n \log n}$, then the 3-majority protocol converges to color m in $\mathcal{O}(\lambda \log n)$ time w.h.p.*

PROOF. For the sake of convenience, let

$$\Lambda = 22\sqrt{\lambda n \log n}$$

Notice that $\Lambda \leqslant 22n^{3/4}\sqrt{\log n}$. In order to make use of Lemma A.4 (see the appendix), we consider the Markov chain determined by the 3-majority protocol. The states of the Markov chain are all the possible assignments of the k colors to the n nodes. For any assignment $\mathbf{a}$, let $\bar{c}(\mathbf{a})$ denote the k-cd determined by $\mathbf{a}$ and let $c_j(\mathbf{a})$ denote any its component. Let X_t be the random variable that is the state at time t given that X_0 is a state whose k-cd is $\bar{c}$. Define

$$T_1 = \left\lfloor 1 + \frac{\log \frac{n}{3\Lambda}}{\log\left(1 + \frac{1}{6\lambda}\right)} \right\rfloor$$

$$T_2 = \left\lfloor 1 + \frac{3}{2\log(18/17)} \log \frac{n^{3/4}}{\log n} \right\rfloor$$

For any $i = 1, \ldots, T_1$, let

$$A_i = \left\{\mathbf{a} \,\middle|\, c_m(\mathbf{a}) > \frac{2}{3}n \vee \left(M(\bar{c}(\mathbf{a})) = \{m\} \wedge c_m(\mathbf{a}) \geqslant \frac{n}{\lambda} \wedge s(\bar{c}(\mathbf{a})) \geqslant \Lambda\left(1 + \frac{1}{6\lambda}\right)^{i-1}\right)\right\}$$

Observe that $X_0 \in A_1$ and $A_1 \supseteq A_2 \supseteq \cdots \supseteq A_{T_1}$. For any $i = 1, \ldots, T_2$, let

$$A_{T_1+i} = \left\{\mathbf{a} \,\middle|\, M(\bar{c}(\mathbf{a})) = \{m\} \wedge s(\bar{c}(\mathbf{a})) \geqslant \frac{n}{3} \wedge n - c_m(\mathbf{a}) \leqslant \frac{2n}{3}\left(\frac{17}{18}\right)^{i-1}\right\}$$

and let

$$A_{T_1+T_2+1} = \left\{\mathbf{a} \,\middle|\, M(\bar{c}(\mathbf{a})) = \{m\} \wedge s(\bar{c}(\mathbf{a})) \geqslant \frac{n}{3} \wedge n - c_m(\mathbf{a}) < \sqrt[4]{n}\log n\right\}$$

$$A_{T_1+T_2+2} = \{\mathbf{a} \mid c_m(\mathbf{a}) = n\}$$

It is easy to verify that $A_{T_1} \supseteq A_{T_1+1} \supseteq A_{T_1+2} \supseteq \cdots \supseteq A_{T_1+T_2} \supseteq A_{T_1+T_2+1} \supseteq A_{T_1+T_2+2}$. Thus it holds that $A_1 \supseteq \cdots \supseteq A_{T_1+T_2+2}$. Taking into account that $c_m(\mathbf{a}) > (2/3)n$ implies $s(\bar{c}(\mathbf{a})) \geqslant n/3$, from Lemma 3.5 we have that, for any $i = 1, \ldots, T_1$,

$$\mathbf{P}\left(X_t \in A_i \mid X_{t-1} \in A_i\right) \geqslant 1 - \frac{1}{n^2} \quad \text{and}$$

$$\mathbf{P}\left(X_t \in A_{i+1} \mid X_{t-1} \in A_i\right) \geqslant 1 - \frac{1}{n^2}$$

From Lemma 3.7 we get, for any $i = T_1 + 1, \dots, T_1 + T_2$

$$\mathbf{P}\left(X_t \in A_i \mid X_{t-1} \in A_i \right) \geqslant 1 - \frac{1}{n^2} \quad \text{and}$$
$$\mathbf{P}\left(X_t \in A_{i+1} \mid X_{t-1} \in A_i \right) \geqslant 1 - \frac{1}{n^2}$$

moreover

$$\mathbf{P}\left(X_t \in A_{T_1+T_2+1} \mid X_{t-1} \in A_{T_1+T_2+1} \right) \geqslant 1 - \frac{1}{n^2} \quad \text{and}$$
$$\mathbf{P}\left(X_t \in A_{T_1+T_2+2} \mid X_{t-1} \in A_{T_1+T_2+1} \right) \geqslant 1 - \frac{1}{\sqrt[5]{n}}$$

Hence, by applying Lemma A.4 with $\epsilon = 1/n^2$ and $\nu = 1/\sqrt[5]{n}$ with $\ell = 10$, we obtain

$$\mathbf{P}\left(X_{10T} \in A_T \mid X_0 \in A_1 \right) \geqslant 1 - T\left(\frac{10}{n^2} + \left(\frac{1}{\sqrt[5]{n}}\right)^{10}\right)$$
$$= 1 - \frac{11T}{n^2}$$

It easy to see that $T < n/11$. Thus in time $10T$ the 3-majority protocol converges to color m w.h.p. Now we bound T in a more precise way. It holds that

$$\begin{aligned} T &= T_1 + T_2 + 2 \\ &\leqslant 4 + \frac{\log \frac{n}{3\Lambda}}{\log\left(1 + \frac{1}{6\lambda}\right)} + \frac{3}{2\log(18/17)} \log \frac{n^{3/4}}{\log n} \\ &\leqslant 4 + 27\log n + \frac{\log \frac{n}{66\sqrt{\lambda n \log n}}}{\log\left(1 + \frac{1}{6\lambda}\right)} \\ &\leqslant 28 \log n + \frac{\log\left(\frac{1}{66}\sqrt{\frac{n}{\lambda \log n}}\right)}{\frac{1/(6\lambda)}{1+1/(6\lambda)}} \qquad \left(\text{since } \log(1+x) \geqslant \tfrac{x}{1+x}\right) \\ &\leqslant 26 \log n + 7\lambda \log(n/\lambda) \\ &\leqslant 10\,\lambda \log n \end{aligned}$$

□

Observation 3.9. *Let us consider a dynamic adversary (see the Introduction) that can change the color of up to T nodes at the beginning of each time step and assume $T = o(\lambda \cdot s)$. Then, Theorem 3.8 still holds since the impact of such a T-bounded adversary is negligible in the growth of the bias s (this can be easily seen in the proof of Lemma 3.5). For instance, when $k \leqslant 2\sqrt[3]{\frac{n}{\log n}}$, then the tolerance of the 3-majority dynamics is $T = o(s/k)$.*

The next three corollaries of Theorem 3.8 address three relevant special cases. Corollary 3.10 is obtained by setting $\lambda = \min\left\{2k,\ \sqrt[3]{\frac{n}{\log n}}\right\}$ and it provides a bound which does not assume any condition on c_m.

Corollary 3.10. *If $\bar{c}$ is a k-cd such that, for some $m \in [k]$, $M(\bar{c}) = \{m\}$ and*

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

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

Corollaries 3.11 and 3.12 are obtained by setting $\lambda = \text{poly}\log(n)$ and $\lambda = \Theta(1)$, respectively. They require some lower bounds on c_m.

Corollary 3.11. *If $\bar{c}$ is a k-cd such that, for some $m \in [k]$, $M(\bar{c}) = \{m\}$, $c_m \geqslant n/\log^\ell n$, and $s(\bar{c}) \geqslant 22\sqrt{n \log^{\ell+1} n}$, then the 3-majority protocol converges to color m in $\mathcal{O}(\log^{\ell+1} n)$ time w.h.p.*

Corollary 3.12. *If $\bar{c}$ is a k-cd such that, for some $m \in [k]$, $M(\bar{c}) = \{m\}$, $c_m \geqslant n/\beta$, and $s(\bar{c}) \geqslant 22\sqrt{\beta n \log n}$, for some constant $\beta \geqslant 3$, then the 3-majority protocol converges to color m in $\mathcal{O}(\log n)$ time w.h.p.*

4. LOWER BOUNDS

This section is organized in 3 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; finally, in the third subsection, we provide a lower bound on the convergence time of the h-plurality dynamics for $h > 3$.

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 $\Omega(k \log n)$ steps 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 time step 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 $\{X_t\}$ be the sequence of random variables where X_t is the number of a specific color at time t. If $X_t = n/k + a$ for some $a \leqslant b$ then $X_{t+1} \leqslant n/k + (1 + 3/k)b$ w.h.p.; more precisely, for any $a \leqslant b$ it holds that*

$$\mathbf{P}\left(X_{t+1} \geqslant \frac{n}{k} + \left(1 + \frac{3}{k}\right) b \;\middle|\; X_t = \frac{n}{k} + a \right) \leqslant \frac{1}{n^2}$$

Proof. For a color h and time step t, let $C_{h,t}$ be the random variable indicating the number of nodes with color h, let $C_t = (C_{1,t}, \dots, C_{k,t})$ be the random variable indicating the *coloring* at time t. For any coloring $\bar{c} = (c_1, \dots, c_k)$ with $\sum_{h=1}^k c_h = n$ and any color $h \in [k]$, the expected value of the number of nodes colored h at time $t+1$ given $C_t = \bar{c}$ is (see Lemma 2.1)

$$\mathbf{E}\left[C_{h,t+1} \mid C_t = \bar{c}\right] = c_{h,t}\left(1 + \frac{c_{h,t}}{n} - \frac{1}{n^2}\sum_{j=1}^k c_j^2\right)$$

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, if X_t is the random variable counting the number of nodes of one specific color, then we can give an upper bound on the expectation of X_{t+1} that depends only on X_t and not on the whole coloring at time t, namely

$$\mathbf{E}\left[X_{t+1} \mid X_t\right] \leqslant X_t\left(1+\frac{X_t}{n}-\frac{1}{k}\right)$$

If we condition on the number of nodes of that specific color being of the form $n/k+a$ for some $a \leqslant b$ we get

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

where in the last two inequalities we used that $a \leqslant b$ and $b \leqslant n/k$.[3] Since X_t can be written as a sum of n independent Bernoulli random variables, from Chernoff bound (see Lemma A.1) we thus get that for every $a \leqslant b$ it holds that

$$\mathbf{P}\left(X_{t+1} \geqslant \frac{n}{k}+\left(1+\frac{3}{k}\right) b \mid X_t=n/k+a\right) \leqslant e^{-2(b/k)^2/n} \leqslant \frac{1}{n^2}$$

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

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

THEOREM 4.2. *Let C_t be the random variable indicating the coloring at time t according to the 3-majority protocol and let $\tau=\inf\{t \in \mathbb{N} : C_t$ is monochromatic$\}$ be the random variable indicating the first time step such that C_t is monochromatic. If the initial number of colors is $k \leqslant (n/\log n)^{1/4}$ and the initial coloring is $C_0=(c_1,\ldots,c_k)$ with $\max\{c_h : h=1,\ldots,k\} \leqslant \frac{n}{k}+\left(\frac{n}{k}\right)^{1-\varepsilon}$ then $\tau=\Omega(k \log n)$ w.h.p.*

A full-detailed proof is given in the full-version [4], we here provide its main argument.

Idea of the proof. For a color $h \in [k]$ let us denote the difference $C_{h,t}-n/k$ 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 time step (w.h.p.). Hence, if a color starts with a positive unbalance smaller than $(n/k)^{1-\varepsilon}$, then it will take $\Omega(k \log n)$ time steps to reach an unbalance of n/k w.h.p. By union bounding on all the colors, we can get the stated lower bound. □

It may be worth noticing that what we actually prove in Theorem 4.2 is that $\Omega(k \log n)$ time steps 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.

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

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 time step 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*

$$f : [k]^h \to [k]$$

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

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

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 coloring $\bar{c}$, when running $\mathcal{P}$, with probability at least $1-\varepsilon$ there is a time step 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 $\left(\Theta(\sqrt{\min\{2k,(n/\log n)^{1/3}\} n \log n}), \Theta(1/n)\right)$-solver in $\mathcal{D}_3$. We can now state the main result of this section.

Theorem 4.8 (properties for good solvers). *(a) If a protocol $\mathcal{P}$ 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

$$\begin{aligned} \Delta_r &= \delta(r,r,b) + \delta(r,b,r) + \delta(b,r,r) \\ \Delta_b &= \delta(b,b,r) + \delta(b,r,b) + \delta(r,b,b) \end{aligned}$$

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

$$\begin{aligned} p(r) &= \left(\frac{c_r}{n}\right)^3 + \left(\frac{c_r}{n}\right)^2 \frac{c_b}{n} \cdot \Delta_r + \left(\frac{c_b}{n}\right)^2 \frac{c_r}{n}(3-\Delta_b) \\ &= \frac{c_r}{n^3}\left(c_r^2 + c_b\left(c_r\Delta_r - c_b\Delta_b\right) + 3c_b^2\right) \qquad (7) \end{aligned}$$

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 coloring 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 (7) we get

$$p(r) \leqslant \frac{c_r}{n^3}\left(c_r^2 + c_b(c_r - c_b)\Delta_b - c_r c_b + 3c_b^2\right) \qquad (8)$$

If the majority of nodes is red then $c_r - c_b$ is positive, and since Δ_b can be at most 3 from (8) we get

$$p(r) \leqslant \frac{c_r}{n^3}\left(c_r^2 + 2c_r c_b\right) \qquad (9)$$

Finally, if we put $c_r = n/2 + s$ and $c_b = n/2 - s$, for some positive s, in (9), we get that

$$p(r) \leqslant \frac{c_r}{n^3}\left(\frac{3}{4}n^2 + (n-s)s\right) \leqslant \frac{c_r}{n} \qquad (10)$$

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 (7) we get that

$$p(r) \leqslant \frac{c_r}{n^3}\left(c_r^2 + 2c_b(c_r - c_b) + 3c_b^2\right) = \frac{c_r}{n} \qquad (11)$$

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 time t, from (10) and (11) we get that $\mathbf{E}\left[X_{t+1} \mid X_t\right] \leqslant X_t$, 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 $\mathbf{P}(\tau < \infty) = 1$ and all X_t's have values bounded between 0 and n, from the martingale stopping theorem[4] we get that $\mathbf{E}[X_\tau] \leqslant \mathbf{E}[X_0]$. 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 $\mathbf{E}[X_\tau] = (1-\varepsilon)n$. 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.*

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.

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

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 $p(r) = \mathbf{P}\left(X_v = r \mid C = \bar{c}\right)$ 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

[4]See e.g. Chapter 17 in [13] for a summary of martingales and related results

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

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 time step).

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

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 $\Theta(k \log n)$ time steps 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 time step, 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 $(1 + \Theta(h^2/k))$ is exponentially small.

Lemma 4.11. *Let $j \in [k]$ be a color and let X_t be the random variable counting the number of j-colored nodes at time t. If $k/h = \mathcal{O}(n^{(1-\varepsilon)/4})$, then for every $(n/k) \leqslant a \leqslant 2(n/k)$ it holds that*

$$\mathbf{P}\left(X_{t+1} \geqslant \left(1 + \frac{h^2}{k}\right) a \mid X_t = a\right) \leqslant e^{-\Theta(n^\varepsilon)}$$

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 time step and let Y be the indicator random variable of the event that node u chooses color j at time step $t+1$. 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 time step)

$$\mathbf{P}\left(Y_u = 1\right) \leqslant \mathbf{P}\left(Y_u = 1 \mid N_j = 1\right)\mathbf{P}\left(N_j = 1\right) + \mathbf{P}\left(N_j \geqslant 2\right) \quad (12)$$

Now observe that

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

Hence, in (12) we have that

$$\mathbf{P}\left(Y = 1\right) \leqslant \frac{a}{n} + \frac{h^2}{2}\cdot\frac{a^2}{n^2}$$

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

$$\mathbf{E}\left[X_{t+1} \mid X_t = a\right] \leqslant a + \frac{h^2}{2n}a^2 = a\left(1 + \frac{h^2}{2n}a\right) \leqslant a\left(1 + \frac{h^2}{k}\right)$$

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

$$\mathbf{P}\left(X_{t+1} \geqslant a\left(1 + 2\frac{h^2}{k}\right) \mid X_t = a\right) \leqslant \exp\left(-\frac{2(ah^2/k)^2}{n}\right)$$
$$\leqslant \exp\left(-\Omega(n^\varepsilon)\right)$$

where in the last inequality we used $a \geqslant n/k$ 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 coloring at time t according to the h-plurality protocol and let $\tau = \inf\{t \in \mathbb{N} : \mathbf{C}_t \text{ is monochromatic}\}$. If the initial coloring is $\mathbf{C}_0 = (c_1, \ldots, c_k)$ with $\max\{c_j : j = 1, \ldots, k\} \leqslant \frac{3}{2}\cdot\frac{n}{k}$ then $\tau = \Omega(k/h^2)$ w.h.p.*

Proof. Since in the initial coloring the plurality color has $a \leqslant (3/2)(n/k)$ nodes, 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 $X_{t+1} \leqslant \left(1 + 2h^2/k\right) X_t$ which, in turn, gives

$$X_t \leqslant \left(1 + 2h^2/k\right)^t X_0 \leqslant \left(1 + 2h^2/k\right)^t \frac{3}{2}\cdot\frac{n}{k}$$

We thus have that

$$(3/2)\left(1 + \frac{2h^2}{k}\right)^t \leqslant 2 \quad \text{for} \quad t \leqslant \frac{k}{h^2}\log(4/3)$$

□

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 cannot derive any stronger bound on the required bias, however, in the full-version of the paper [4], we show that there are 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 $s(t)$. On the contrary, 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 mainly 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.

6. REFERENCES

[1] M. A. Abdullah and M. Draief. Majority consensus on random graphs of a given degree sequence. *arXiv:1209.5025*, 2012.

[2] D. Angluin, J. Aspnes, and D. Eisenstat. A simple population protocol for fast robust approximate majority. In *Proc. of 21st DISC*, volume 4731 of *LNCS*, pages 20–32. Springer, 2007.

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

[4] L. Becchetti, A. Clementi, E. Natale, F. Pasquale, R. Silvestri, and L. Trevisan. Simple dynamics for plurality consensus. *arXiv:1310.2858*, 2013.

[5] F. Bénézit, P. Thiran, and M. Vetterli. Interval consensus: from quantized gossip to voting. In *Proc. of ICASSP*, pages 3661–3664. IEEE, 2009.

[6] A. Clementi, M. Di Ianni, G. Gambosi, E. Natale, and R. Silvestri. Distributed community detection in dynamic graphs. In *SIROCCO*. Springer, 2013.

[7] B. Doerr, L. A. Goldberg, L. Minder, T. Sauerwald, and C. Scheideler. Stabilizing consensus with the power of two choices. In *Proc. of 23rd SPAA*, 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. P. Dubhashi and A. Panconesi. *Concentration of measure for the analysis of randomized algorithms.* Cambridge University Press, 2009.

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

[11] M. Kearns and J. Tan. Biased voting and the democratic primary problem. In *Proc. of 4th WINE*, volume 5385 of *LNCS*, pages 639–652. Springer, 2008.

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

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

[14] 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.

[15] E. Mossel and G. Schoenebeck. Reaching consensus on social networks. In *Proc. of ICS*, pages 214–229. Tsinghua University Press, 2010.

[16] D. Peleg. Local majorities, coalitions and monopolies in graphs: a review. *Theor. Comput. Sci.*, 282(2), 2002.

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

APPENDIX

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 $x_1, \dots x_k \in \mathbb{R}$ be k real numbers, then*

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

LEMMA A.3 (BERNSTEIN INEQUALITY [9]). *Let the random variables $X_1, ..., X_n$ be independent with $X_i - \mathbb{E}[X_i] \leqslant b$ for each $i \in [n]$. Let $X = \sum_i X_i$ and let $\sigma^2 = \sum_i \sigma_i^2$ be the variance of X. Then, for any $\lambda > 0$,*

$$\mathbf{P}(X > \mathbb{E}[X] + \lambda) \leqslant \exp\left(-\frac{\lambda^2}{2\sigma^2(1 + b\lambda/3\sigma^2)}\right)$$

We now provide a useful result on finite Markov chains. The proof is given in the full-version [4].

LEMMA A.4. *Let $\{X_t\}_t$ be a finite-state Markov chain with state space S. If $A_1, \dots, A_T$ are such that $S \supseteq A_1 \supseteq A_2 \supseteq \cdots \supseteq A_T$ and, for any $i = 1, \dots, T$, $\mathbf{P}(X_t \in A_i \mid X_{t-1} \in A_i) \geqslant 1 - \epsilon$ and for $i < T$, $\mathbf{P}(X_t \in A_{i+1} \mid X_{t-1} \in A_i) \geqslant 1 - \nu$ where $0 \leqslant \epsilon \leqslant \nu < 1$, then, for any integer $\ell \geqslant 1$ it holds that*

$$\mathbf{P}(X_{\ell T} \in A_T \mid X_0 \in A_1) \geqslant 1 - T(\ell\epsilon + \nu^\ell)$$