

Plurality Consensus in the Gossip Model*

L. Becchetti[†] A. Clementi[‡] E. Natale[†] F. Pasquale[†] R. Silvestri[†]

Abstract

We study *Plurality Consensus* in the *Gossip Model* over a network of n anonymous agents. Each agent supports an initial opinion or *color*. We assume that at the onset, the number of agents supporting the *plurality* color exceeds that of the agents supporting any other color by a sufficiently-large *bias*, though the initial plurality itself might be very far from absolute majority. The goal is to provide a protocol that, with high probability, brings the system into the configuration in which all agents support the (initial) plurality color.

We consider the *Undecided-State Dynamics*, a well-known protocol which uses just one more state (the undecided one) than those necessary to store colors.

We show that the speed of convergence of this protocol depends on the initial color configuration as a whole, not just on the gap between the plurality and the second largest color community. This dependence is best captured by a novel notion we introduce, namely, the *monochromatic distance* $\text{md}(\bar{\mathbf{c}})$ which measures the distance of the initial color configuration $\bar{\mathbf{c}}$ from the closest monochromatic one. In the complete graph, we prove that, for a wide range of the input parameters, this dynamics converges within $O(\text{md}(\bar{\mathbf{c}}) \log n)$ rounds. We prove that this upper bound is almost tight in the strong sense: Starting from *any* color configuration $\bar{\mathbf{c}}$, the convergence time is $\Omega(\text{md}(\bar{\mathbf{c}}))$.

Finally, we adapt the Undecided-State Dynamics to obtain a fast, random walk-based protocol for plurality consensus on *regular expanders*. This protocol converges in $O(\text{md}(\bar{\mathbf{c}}) \text{polylog}(n))$ rounds using only $\text{polylog}(n)$ local memory. A key-ingredient to achieve the above bounds is a new analysis of the maximum node congestion that results from performing n parallel random walks on regular expanders.

All our bounds hold with high probability.

- **Keywords.** Gossip Algorithms, Plurality Consensus, Markov Chains, Random Walks.

1 Introduction

Reaching *Plurality Consensus* is a fundamental task in distributed computing. Each agent of a distributed system initially supports a color, i.e. a number $i \in [k] = \{1, 2, \dots, k\}$ (with $2 \leq k \leq n$). In the initial color configuration $\bar{\mathbf{c}} = \langle \bar{c}_1, \dots, \bar{c}_k \rangle$ (where \bar{c}_i denotes the number of agents supporting color $i \in [k]$), there is an initial *plurality* \bar{c}_1 of agents supporting the *plurality color* (wlog, we assume that color communities are ordered, so that $\bar{c}_i \geq \bar{c}_{i+1}$ for any $i \leq k - 1$). Initially, every agent only knows its own color; the goal is a distributed algorithm that, *with high probability* (in short, *w.h.p.*)¹, brings the system into the *target* configuration, i.e., the monochromatic configuration in which all agents support the initial plurality color. In the remainder, the subset of agents supporting color i is called the *i -color community*.

This problem is also known as *majority consensus* or *proportionate agreement* [3, 1, 30], but we prefer the term *plurality* in this paper, in order to emphasize that the initial plurality \bar{c}_1 might be far from the (absolute) majority: for instance, it could be some root of n . We study plurality consensus in the *Gossip model* [9, 15, 22] in which each of n agents of a communication network can, in every round, contact one (possibly random) neighbor to exchange information. Agents can be anonymous, i.e., they don't need to possess unique labels. A major open question for the plurality consensus problem is whether a plurality protocol exists that converges in polylogarithmic time and uses only polylogarithmic local memory [3, 1, 30].

There is a strong interest for simple plurality protocols (called *dynamics*) in which agents possess just a few more states than those necessary to store the k possible colors [3, 5, 13, 16, 8, 30]. In this paper, we consider the Undecided-State Dynamics², that has been introduced in [3] and analyzed in [3, 30] only in the binary case (i.e. $k = 2$). The analysis of the multivalued case (i.e. $k > 2$) has been proposed in [3, 1, 13, 16, 24] as an open problem. The interest for this dynamics touches areas

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

[†]Sapienza Università di Roma

[‡]Università Tor Vergata di Roma

¹As usual, we say that an event \mathcal{E}_n holds w.h.p. if $\mathbf{P}(\mathcal{E}_n) \geq 1 - n^{-\Theta(1)}$.

²The Protocol has been initially “designed” for the case $k = 2$ and, thus, it has been named the *Third-State Dynamics*.

beyond the borders of computer science. It appears to play a major role in important biological processes modelled as so-called chemical reaction networks [8, 17].

As discussed further in the introduction, in previous work, the performance of this dynamics on the complete graph has been evaluated w.r.t. the following parameters: the number n of nodes, the number k of colors, and the initial *bias* towards the plurality color, with the latter characterized in terms of a parameter that only depends on the relative magnitude³ of \bar{c}_1 and \bar{c}_2 .

However, when $k > 2$, any such measure of the initial bias is not sensitive enough to accurately capture the convergence time of a plurality protocol: a *global* measure is needed, i.e., one that reflects the whole initial color configuration. To better appreciate this issue, consider the two configurations $\bar{\mathbf{c}}$ and $\bar{\mathbf{c}}'$ in Fig. 1. Whether the absolute difference or the relative ratio is used to measure the initial bias, the color configuration $\bar{\mathbf{c}}'$ appears to be not “worse” than $\bar{\mathbf{c}}$. Still, computer simulations and intuitive arguments suggest that, under any “natural” plurality protocol, the almost-uniform color distribution $\bar{\mathbf{c}}'$ can result in much larger convergence times than the highly-concentrated color configuration $\bar{\mathbf{c}}$.

To the best of our knowledge, the impact of the whole initial color configuration on the speed of convergence of plurality protocols has never been analyzed before.




Figure 1: Two different color configurations having the same bias $s = s(c_1, c_2)$

Our Contributions. We first introduce a suitable distance $d(\cdot, \cdot)$ (see Section 2 and Appendix B for a formal definition) on the set \mathcal{S} of all color configurations. It naturally induces a function $\text{md}(\cdot)$, called the *monochromatic distance*, which equals the *distance* between any configuration \mathbf{c} and the target configuration:

$$\text{md}(\mathbf{c}) = \sum_{i=1}^k \left(\frac{c_i}{c_1} \right)^2$$

We use md to characterize the bias of the initial configuration. In particular, note that $\text{md}(\bar{\mathbf{c}})$ measures

the extent to which $\bar{\mathbf{c}}$ is “uniform”: Indeed, the higher the extent of the bias towards a small subset of the colors (including the plurality one), the smaller the value of $\text{md}(\bar{\mathbf{c}})$. As an example, in Fig. 1, $\text{md}(\bar{\mathbf{c}})$ can be substantially smaller than $\text{md}(\bar{\mathbf{c}}')$. At the extremes, when there are only $O(1)$ color communities of size $\Theta(\bar{c}_1)$, we have $\text{md}(\bar{\mathbf{c}}) = \Theta(1)$ while, when $\Theta(k)$ color communities have size $\Theta(n/k)$, we have $\text{md}(\bar{\mathbf{c}}) = \Theta(k)$.

The simple strategy of the Undecided-State Dynamics [3, 30] is to “add” one extra state to somewhat account for the “previous” opinion supported by an agent (see Section 2 and Table 1 for a definition of this dynamics). The convergence time of this dynamics has been analyzed on different distributed models, but only in the binary case [1, 3, 4, 6, 18, 30]. In this restricted setting the complex dependence of the dynamics’ evolution on the overall shape of the initial color configuration is not exhibited.

We analyse the Undecided-State Dynamics using a technique that strongly departs from past work and that allows us to address the plurality consensus problem in the general setting. Our analysis achieves almost-tight bounds on convergence time. Formally, let $k = k(n)$ be any function such that $k = O((n/\log n)^{1/3})$, and consider any initial configuration $\bar{\mathbf{c}} \in \mathcal{S}$ such that $\bar{c}_1 \geq (1 + \alpha)\bar{c}_2$ where $\alpha > 0$ is *any arbitrarily-small* constant (this is a weak-bias condition that ensures the convergence of the process towards the plurality color). Then, the Undecided-State Dynamics converges in $O(\text{md}(\bar{\mathbf{c}}) \log n)$ rounds w.h.p.

This result is almost-tight in a strong sense. In particular, we are able to prove that, for $k = O((n/\log n)^{1/6})$ and for any initial k -colors configurations $\bar{\mathbf{c}}$, the convergence time of the Undecided-State Dynamics is linear in the monochromatic distance $\text{md}(\bar{\mathbf{c}})$ w.h.p.

The best previous results [5, 22] about plurality protocols will be compared to ours later in this introduction. We only emphasize that, when k is some root of n , our refined analysis implies that this dynamics is exponentially faster than the best protocol that uses polylogarithmic bounded memory [5] on a large class of initial color configurations. Moreover, we observe that the Undecided-State Dynamics uses exponentially-smaller message and memory size w.r.t. the fastest (i.e. polylogarithmic-time) gossip protocol in [22].

Our analysis is rather general and it can be extended to other interesting topologies. As a case supporting this claim, we show how to adapt the Undecided-State Dynamics for the class of d -regular expanders [21], for any degree $d \geq 1$. Efficient dynamics for this class of graphs have only been analyzed for the binary case [13, 25].

³Typically, this relative magnitude is defined in terms of the absolute difference or the ratio.

In this variant of the Undecided-State Dynamics, the task of selecting random neighbors is simulated by performing n independent random-walks of suitable length. Thanks to the well-known rapidly-mixing properties of d -regular expanders [21, 23], we can prove that the new protocol converges in time $O(\text{md}(\bar{c})\text{polylog}(n))$, w.h.p.

The major technical hurdle here is proving that this variant of the protocol still requires $\text{polylog}(n)$ local memory. To this aim, we prove that the *node congestion* is at most $\text{polylog}(n)$. The analysis of the process that results from running parallel random walks over a graph has been the subject of extensive research in the past [2, 19, 20, 28, 14]. However, to the best of our knowledge, none has addressed the issues we consider here. In particular, the analysis of node congestion is far from trivial and of independent interest, since efficient protocols for several important tasks in the *Gossip* model (such as *node-sampling* [14], *network-discovery* problems [20], and *averaging* problems [7]) rely on the use of parallel random walks.

Motivations and comparison to previous works.

Plurality consensus (a.k.a. majority consensus or proportionate agreement) is a fundamental problem arising in several areas such as *distributed computing* [3, 16, 29], *communication networks* [30], *social networks* [11, 27, 25] and *biology* [8].

Applications include fault-tolerance in parallel computing and in distributed database management where data redundancy or replication and majority-rules are used to manage the presence of unknown faulty processors [16, 29]. Another application comes from the task of distributed item ranking, in particular when every node initially ranks some item and the goal is to agree on the rank of the item based on the initial plurality opinion [30]. Further areas of interest of the multi-valued case include distributed cooperative decision-making and control in environmental monitoring, surveillance and security [31]. Finally, converging to the plurality color among a (large) set of initial node colors has been recently used as a basic building block for *community detection* in dynamic social networks [11]. We remark that, in all such applications, the data domain can span a relatively-large range of values, hence the importance of this problem for large values of k .

Interestingly enough, only the binary case is essentially settled, even for complete graphs. In the synchronous model, a simple gossip protocol for computing the median can be used to solve the majority consensus problem in the binary case, with constant memory and message size [16]. The proposed protocol converges in $O(\log n)$ time rounds if the initial difference

bias $s = \bar{c}_1 - \bar{c}_2$ is $\Omega(\sqrt{n \log n})$.

More recently, in [13], the authors provide a rigorous analysis of a simple 2-voting dynamics for the binary case on any (possibly random) regular graph: in the latter case, they provide optimal bounds on the convergence time as a function of the second-largest eigenvalue of the graph.

For the multivalued case, in [5] the authors analyze a gossip protocol, called *3-Majority Dynamics*, where at every round, each agent applies a simple majority rule over the colors of three randomly-sampled neighbors. When the initial difference bias is $s = \Omega(\sqrt{kn \log n})$, the 3-Majority Dynamics converges in $\Theta(\min\{k, n^{1/3}\} \log n)$ rounds using $\Theta(\log k)$ memory and message size.

Convergence times of the 3-Majority Dynamics become polylogarithmic only if $\bar{c}_1 \geq n/\text{polylog}(n)$, thus they are not polylogarithmic whenever $k = \omega(\text{polylog}(n))$ and $\bar{c}_1 = o(n/\text{polylog}(n))$. This is the parameter range where our analysis of the Undecided-State Dynamics leads to an exponential speed up w.r.t. the convergence time of the 3-Majority Dynamics. For example, consider an initial “oligarchic” scenario where $k = n^{1/4}$ and a subset $\mathcal{L} \subseteq [k]$ exists such that $|\mathcal{L}| = \text{polylog}(n)$, for any $i \in \mathcal{L}$, $\bar{c}_i \sim n/\sqrt{k}$, and, for any $i \in [k] \setminus \mathcal{L}$, $\bar{c}_i \sim n/k$. Clearly, $1, 2 \in \mathcal{L}$ and the resulting monochromatic distance is $\text{md}(\bar{c}) = \text{polylog}(n)$. Assuming $\bar{c}_1 \geq (1 + \alpha)\bar{c}_2$ for some $\alpha > 0$ our upper bound implies that, starting from any such configuration, the Undecided-State Dynamics converges in polylogarithmic time, whereas the 3-Majority Dynamics converges in $\Theta(k \log n)$ time [5].

In [22], the authors provide a gossip protocol to compute aggregate functions, which can be used to solve plurality consensus in $\text{polylog}(n)$ time starting from any positive bias, but it requires exponentially larger memory and message size (namely $\Theta(k \log n)$). The Undecided-State Dynamics has been introduced and analyzed in [3] for the binary case in the population protocol model (where only one edge is active during a round). They prove that this dynamics has “parallel” convergence time $O(\log n)$ whenever the bias is $\Omega(\sqrt{n \log n})$. In [4, 6, 18, 30, 24], the same dynamics for the binary case has been analyzed in different distributed models. Last but not least, interest for this dynamics was stimulated by recent findings in biology: notably, as shown in [8], the structure and dynamics of the “approximate majority” protocol (as it is called there and in [3]) is to a great extent similar to a mechanism that is collectively implemented in the network that regulates the mitotic entry of the cell cycle in eukaryotes.

We mention that similar majority-consensus prob-

lems have been studied (for example in [1, 26]) in the *LOCAL (communication) model* [19, 28] where, however, node congestion and memory size are linear in the node degree of the network.

2 Preliminaries

We consider a complete graph of n anonymous nodes (*agents*), each of them is initially colored with one out of k possible colors, where $k = k(n) \in [n]$. We assume an initial *plurality* of agents colored with the *plurality color* $j \in [k]$. Wlog, we assume $j = 1$. A synchronous protocol for the *plurality problem* is a finite set of local rules (applied by every agent) that eventually bring the system into the absorbing *target configuration*, in which all agents share the initial plurality color.

The Undecided-State Dynamics. We analyze the synchronous version of the dynamics introduced in [3] and [30] in the (uniform) *GOSIP* model: in every round, each agent pulls the color of a randomly-selected neighbor. If this color differs from its own, the agent enters the *undecided* state, an extra state that an agent can support. When an agent is in the undecided state and pulls a color, it gets that color. Finally, an agent that pulls either the undecided color or its own color remains in its current state (see also Table 1).

$u \backslash v$	undecided	color i	color j
undecided	undecided	i	j
i	i	i	undecided
j	j	undecided	j

Table 1: The update rule of the Undecided-State Dynamics where $i, j \in [k]$ and $i \neq j$.

So, differently from other protocols (e.g., the majority dynamics considered in [5]), after the first round agents can also enter an undecided state, to which no color is associated. At each round t , the global state of the system is completely characterized by the corresponding color configuration, namely by the vector $\mathbf{c}^{(t)} = \langle c_1^{(t)}, c_2^{(t)}, \dots, c_k^{(t)}, q^{(t)} \rangle$, where $c_i^{(t)}$ (respectively $q^{(t)}$) denotes the number of nodes that are colored i (respectively are in the undecided state) at the end of the t -th round. In the initial state, we always have $q^{(0)} = 0$.

As also remarked further, the dynamic process that results from running this protocol on the complete graph induces a finite-state Markov chain defined over the space of all color configurations.

Basic notation. Consider any initial color configuration $\bar{\mathbf{c}} = \langle \bar{c}_1, \bar{c}_2, \dots, \bar{c}_k, 0 \rangle$. Assume wlog that $\bar{c}_i \geq \bar{c}_{i+1}$ for any $i \leq k-1$. Then, at any time $t \geq 0$, the execution of the protocol (uniquely) de-

termines the probability distribution of the (vectorial) random variable indicating the state at time t : $\mathbf{C}^{(t)} = \langle C_1^{(t)}, C_2^{(t)}, \dots, C_k^{(t)}, Q^{(t)} \rangle$. Notice that we omit the dependence of the random state on the initial color configuration in our notation. Since we are considering complete graphs, this random process is clearly a finite-state Markov chain. Lower-case letters will be used to denote functions of the observed color configuration at any specified time. Upper-case letters instead will denote *random variables* (r.v.s). In particular, $Q^{(t)}$ and $C_i^{(t)}$ denote the number of nodes that are undecided and that have color i , respectively, at time t .

Finally, when we condition the system to be in a fixed state \mathbf{c} at some round, the random community sizes in the next round will be denoted by C'_i and Q' .

Global bias. We define a *distance*⁴ between color configurations as follows:

$$d(\mathbf{c}, \mathbf{c}') = \sum_i \left(\frac{c_i}{c_1} - \frac{c'_i}{c'_1} \right)^2$$

In particular, consider the set M of the k possible *monochromatic* color configurations. For any \mathbf{c} , let $d(\mathbf{c}, M) = \min_{\mathbf{c}' \in M} \{d(\mathbf{c}, \mathbf{c}')\}$. It is easy to see that $\text{md}(\mathbf{c}) = d(\mathbf{c}, M) + 1$.

3 Analysis of the Undecided-State Dynamics

Generally speaking, when the initial configuration is sufficiently biased, the dynamics' evolution follows a typical pattern, characterized by well-distinct phases.

Understanding such a pattern requires a careful analysis. In this section, we provide an overview of this analysis, quantitatively describing a typical evolution of the process. We start from the expectations of a few key r.v.s

$$(3.1) \quad \mathbf{E} \left[C_i^{(t+1)} \mid \mathbf{c}^{(t)} \right] = c_i^{(t)} \cdot \frac{c_i^{(t)} + 2q^{(t)}}{n}$$

$$(3.2) \quad \mathbf{E} \left[Q^{(t+1)} \mid \mathbf{c}^{(t)} \right] = \frac{(q^{(t)})^2 + (n - q^{(t)})^2 - \sum_i (c_i^{(t)})^2}{n}$$

These equations follow directly from the definition of the Undecided-State Dynamics. From (3.1), we can appreciate the crucial role of the function $\frac{c_i^{(t)} + 2q^{(t)}}{n}$: It represents the expected *growth rate* of every color

⁴Note that $d(\bar{\mathbf{c}}, \bar{\mathbf{c}}')$ is not a distance in the strict sense. See Appendix B for a formal discussion of this notion.

community. The corresponding r.v. $C_i^{(t+1)} + 2Q^{(t+1)}$ is of particular interest when i is the plurality color⁵. In fact, a major novelty of our contribution is the discovery of a clean mathematical connection between the expected growth rate of the plurality and the monochromatic distance of the current configuration. The following expression formalizes this connection and plays a key role in our analysis. For every $t \geq 0$,

$$(3.3) \quad \mathbf{E} \left[C_1^{(t+1)} + 2Q^{(t+1)} \mid \mathbf{c}^{(t)} \right] = \\ = \frac{n^2 + \left(n - 2q^{(t)} - c_1^{(t)} \right)^2 + 2(R(\mathbf{c}^{(t)}) - \text{md}(\mathbf{c}^{(t)})) \left(c_1^{(t)} \right)^2}{n}$$

where

$$R(\mathbf{c}) = \sum_{i=1}^k \frac{c_i}{c_1}$$

Notice that $1 \leq \text{md}(\mathbf{c}), R(\mathbf{c}) \leq k$ and $R(\mathbf{c}) \geq \text{md}(\mathbf{c})$ (see (B.4) in the appendix). The derivation of (3.3) becomes straightforward only after guessing the (non obvious) key role played by md as a measure of global bias. We observe that it is not linear in several parameters and its recursive form depends, through R and md , on the previous color configuration, as a whole. The resulting process evolution is thus rather complex and hard to analyze in a rigorous way (the details of this analysis can be found in Appendix C). However, (3.3) allows us to informally characterize the main drivers of the process evolution. At the extremes, we have two complementary mechanisms that may determine an exponential (or quasi exponential) growth of C_1 and that qualitatively explain the leftmost (first phase) and rightmost (third phase) regions of Fig. 2: Namely, large values of Q or of C_1 itself. In the latter case, growth follows a preferential attachment-like pattern. In the middle, we have a phase of relative “flat” growth that corresponds to Q dropping to a value close to $n/2$ and C_1 not being large enough to self-sustain an exponential growth. During this phase, growth is basically driven by the term $(R(\mathbf{c}) - \text{md}(\mathbf{c}))c_1^2/n$, i.e., it crucially depends on the distance from the closest monochromatic configuration.

A further remark concerning (3.3) is that its proof crucially relies on properties of the plurality, the argument does not carry over to other colors. In the next subsection, we give an overview of our analysis. For the sake of space and simplicity, we omit some major

technical aspects, mostly related to the rigorous characterization of phase-transition timings and the derivation of concentration bounds.

The process in a nutshell The typical behaviour of the Undecided-State Dynamics follows a characteristic pattern that exhibits three distinct phases, as exemplified in Fig. 2. Note that *the quantitative overview we provide below applies to typical evolutions*. We remark that the typical behavior holds w.h.p. under the assumption that $\bar{c}_1 \geq (1 + \alpha) \cdot \bar{c}_2$, where α is an arbitrarily small positive constant. Indeed this assumption guarantees that the initial plurality is preserved w.h.p. along the whole process.

First round: Rise of the undecided. The initial state is extremely unstable⁶, since any node has a high probability of sampling a node of different color in the first round, ending up in the undecided state. Thus, the first round sees dramatic changes in the system: i) In general, a drastic drop in $C_i^{(1)}$'s (with “small”⁷ ones simply disappearing w.h.p.); ii) An explosive surge in $Q^{(1)}$, that possibly come to account for the vast majority; iii) The initial plurality is preserved w.h.p., though it drops in absolute terms. The results below follow immediately from (3.1) and (3.2) with $t = 0$ and recalling that $q^{(0)} = 0$,

$$(3.4) \quad \mathbf{E} \left[C_1^{(1)} \mid \bar{\mathbf{c}} \right] = \frac{n}{R(\bar{\mathbf{c}})^2},$$

$$(3.5) \quad \mathbf{E} \left[Q^{(1)} \mid \bar{\mathbf{c}} \right] = n \left(1 - \frac{1}{\Lambda(\bar{\mathbf{c}})} \right)$$

where

$$\Lambda(\bar{\mathbf{c}}) = \frac{R(\bar{\mathbf{c}})^2}{\text{md}(\bar{\mathbf{c}})}$$

and notice that $1 \leq \Lambda(\bar{\mathbf{c}}) \leq k$ (see (B.5) in the appendix). Furthermore, $C_1^{(1)}$ and $Q^{(1)}$ are concentrated around their expectations (see Lemma C.3 in Appendix C).

First phase: Age of the undecided. The first phase starts right after round 1. In this phase, the C_i 's grow (almost) exponentially fast while Q decreases. The duration of this phase depends on $\Lambda(\bar{\mathbf{c}})$ (and not just the magnitude of the initial bias). Those facts are discussed in the proof of Claim 1 that highlights key properties of the process marking the end of the first phase (for rigorous statements see Lemmas C.5 and C.6 in Appendix C).

⁵We are implicitly assuming that 1 remains the plurality color across the whole process. This holds w.h.p. under the assumptions of Theorem 3.2.

⁶Exceptions include cases that are less interesting, such as the one in which we have a strong absolute majority already at the onset.

⁷Namely, $o(\sqrt{n})$ in size.




Figure 2: Typical evolution of the Undecided-State Dynamics after the first round, for $n = 7 \cdot 10^{10}$ nodes and $k = (\frac{n}{\log n})^{\frac{1}{4}}$ colors, with $c_1^{(0)} = 2\frac{n}{k}$ and $c_i^{(0)} = \frac{n}{k}(1 - \frac{2}{k})$ for every $i \neq 1$.

CLAIM 1. *Within $T = O(\log \Lambda(\bar{\mathbf{c}}))$ rounds the system reaches a configuration such that w.h.p.*

$$\begin{aligned} Q^{(T)} &= \frac{n}{2} \left(1 \pm \Theta \left(\frac{1}{\text{md}(\bar{\mathbf{c}})} \right) \right) \\ C_1^{(T)} &= \Theta \left(\frac{n}{\text{md}(\bar{\mathbf{c}})} \right) \end{aligned}$$

Furthermore, the relative ratios C_1/C_i are approximately preserved.

Sketch of Proof. We first sketch the proof for the bound on Q . Assume that at some time t we are in a configuration $\mathbf{c}^{(t)}$ such that $q^{(t)} = (n/2)(1 + \beta)$ for some $\beta > 0$. Notice that, choosing $\beta = 1 - \Theta(1/\Lambda(\bar{\mathbf{c}}))$, this assumption holds w.h.p. for $t = 1$ from the above overview of the first round. Then, from (3.2), we immediately have:

$$\mathbf{E} \left[Q^{(t+1)} \mid \mathbf{c}^{(t)} \right] = \frac{n}{2}(1 + \beta^2) - \frac{1}{n} \sum_i \left(c_i^{(t)} \right)^2$$

Under reasonable assumptions on k , from the above inequality we have $Q^{(t+1)} \leq (n/2)(1 + \beta^2)$ w.h.p. (see the proof of Lemma C.6 in Appendix C). Unfolding this argument for t rounds after round 1, we obtain $Q^{(t+1)} \leq (n/2)(1 + \beta^{2^t})$ w.h.p. Recalling that $\beta = 1 - \Theta(1/\Lambda(\bar{\mathbf{c}}))$, we obtain $Q^{(T)} \leq (n/2)(1 + \Theta(1/\text{md}(\bar{\mathbf{c}})))$ for $T = \log \Lambda(\bar{\mathbf{c}}) + O(\log \log \text{md}(\bar{\mathbf{c}}))$.

Moreover, whenever $Q^{(t)} \geq \frac{n}{2} \left(1 + \Theta \left(\frac{1}{\text{md}(\bar{\mathbf{c}})} \right) \right)$, we have $Q^{(t+1)} \geq \frac{n}{2} \left(1 - \Theta \left(\frac{1}{\text{md}(\bar{\mathbf{c}})} \right) \right)$ w.h.p., hence $|Q^{(T)} - n/2| \leq \Theta(1/\text{md}(\bar{\mathbf{c}}))$ w.h.p.

As for the claim for C_1 , we next consider the evolution of the term $C_1^{(t)} + 2Q^{(t)}$ which, up to the factor $1/n$, determines the growth rate of $C_1^{(t+1)}$. Assume that $c_1^{(1)} + 2q^{(1)} = (1 + \epsilon)n$. We know from the analysis of the first round and in particular from (3.5), that this assumption holds w.h.p if we choose

$\epsilon \approx 1 - \Theta(1/\Lambda(\bar{\mathbf{c}}))$ (note that we are neglecting the contribution of $C_1^{(1)}$). Consequently, from (3.3) we get $\mathbf{E} \left[C_1^{(t+1)} + 2Q^{(t+1)} \mid \mathbf{c}^{(t)} \right] \approx (1 + \epsilon^2)n$.

Informally, by applying the argument above iteratively we obtain

$$\begin{aligned} C_1^{(2)} + 2Q^{(2)} &\approx (1 + \epsilon^2)n; \\ C_1^{(3)} + 2Q^{(3)} &\approx (1 + \epsilon^4)n; \\ &\dots \\ C_1^{(t)} + 2Q^{(t)} &\approx (1 + \epsilon^{2^{t-1}})n \end{aligned}$$

At this point, from (3.1) we get

$$\begin{aligned} C_1^{(t)} &\approx C_1^{(1)} \prod_{i=0}^{t-1} (1 + \epsilon^{2^i}) \approx C_1^{(1)} \prod_{i=0}^{t-1} \exp(\epsilon^{2^i}) \\ &\approx C_1^{(1)} \exp \left(\sum_{i=0}^{t-1} \epsilon^{2^i} \right) \\ &\approx C_1^{(1)} \exp \left(\sum_{i=0}^{t-1} \left(1 - \frac{1}{\Lambda(\bar{\mathbf{c}})} \right)^{2^i} \right) \end{aligned}$$

Since $T = \log \Lambda(\bar{\mathbf{c}}) + O(\log \log \text{md}(\bar{\mathbf{c}}))$ it holds that

$$C_1^{(T)} \approx C_1^{(1)} \cdot \Theta(\Lambda(\bar{\mathbf{c}})) \approx \Theta \left(\frac{n}{\text{md}(\bar{\mathbf{c}})} \right)$$

The last derivation follows from (3.5), which approximately holds w.h.p. (see also Lemma C.3 in Appendix C). \square

The proof outlined above highlights the following properties of the first phase: i) The growth rate of Plurality keeps “almost” exponential, while it quickly decreases mirroring the decrease of Q ; ii) The duration of the second phase is determined by $\log \Lambda(\bar{\mathbf{c}})$ (this can be as large as $\Theta(\log n)$ and as small as $O(1)$); iii) From (3.5) it is possible to see that the factor $1/\text{md}(\bar{\mathbf{c}})$, appearing in the expression of $C_1^{(T)}$ in the statement of Claim 1, corresponds to the fraction of the not-undecided nodes that belong to the plurality at the end of round 1.

Second phase: Plateau or Age of stability.

The second phase is characterized by a slow increase of C_1 , roughly at a rate $1 + \Theta(1/\text{md}(\bar{\mathbf{c}}))$ and a substantial stability of Q around the value $n/2$. Indeed, if the system is in a color configuration \mathbf{c} such that

$$q = \frac{n}{2} \left(1 \pm \Theta \left(\frac{1}{\text{md}(\bar{\mathbf{c}})} \right) \right) \text{ and } c_1 = \Theta \left(\frac{n}{\text{md}(\bar{\mathbf{c}})} \right)$$

Equations (3.1) and (3.2) imply that

$$\begin{aligned} \mathbf{E}[Q' | \mathbf{c}] &\approx \frac{n}{2} \left(1 - \Theta \left(\frac{1}{\text{md}(\bar{\mathbf{c}})} \right) \right) \\ \mathbf{E}[C_1' | \mathbf{c}] &\approx \left(1 + \Theta \left(\frac{1}{\text{md}(\bar{\mathbf{c}})} \right) \right) c_1 \end{aligned}$$

By choosing the suitable constants we prove that the above relations hold w.h.p. (see Lemma C.7 in Appendix C). This is also the main argument for proving our lower bound.

THEOREM 3.1. *Let $k = O((n/\log n)^{1/6})$. Starting from any color configuration $\bar{\mathbf{c}}$ the convergence time of the Undecided-State Dynamics is $\Omega(\text{md}(\bar{\mathbf{c}}))$ w.h.p.*

However, as discussed above, since C_1 increases at a rate $1 + \Theta(1/\text{md}(\bar{\mathbf{c}}))$, after a plateau of $O(\text{md}(\bar{\mathbf{c}}) \log \text{md}(\bar{\mathbf{c}}))$ rounds the system reaches a configuration $\mathbf{c}^{(t)}$ such that $R(\mathbf{c}^{(t)}) = 1 + o(1)$. This fact marks the end of the second phase, since the next phase yields a much faster growth of C_1 . For a rigorous analysis of this part see Lemma C.8 and Lemma C.9 in Appendix C.

Third phase: From plurality to totality. Observe that, by definition of R , $C_1 = \frac{n-Q}{R}$ and, when the third phase starts, we have $R = 1 + o(1)$: hence, $C_1 \approx n - Q$. Now, from (3.3), the leading term of the growth rate $\frac{C_1+2Q}{n}$ becomes $1 + (\frac{Q}{n})^2$. So, as long as Q is large (say $Q = \Theta(n)$), C_1 has an exponential growth while Q decreases. The above arguments, rigorously described in the proofs of Lemma C.9 and Theorem C.2 in Appendix C, are the main ingredients to bound the time of the last phase. Finally, the whole analysis above yields the following upper bound

THEOREM 3.2. *Let $k = O((n/\log n)^{1/3})$ and let $\bar{\mathbf{c}}$ be any initial configuration such that $c_1 \geq (1 + \alpha) \cdot c_2$ where α is an arbitrarily small positive constant. Then, w.h.p. within time $O(\text{md}(\bar{\mathbf{c}}) \cdot \log n)$ the system reaches the target configuration.*

4 The Undecided-State Dynamics on regular expander graphs

We next show how to adapt the Undecided-State Dynamics to achieve plurality consensus on the class of d -regular expander graphs [21] (with d denoting the degree of the nodes) at a polylogarithmic extra-cost in terms of local memory and time. The simple idea is to simulate the (uniform) random sampling of nodes' colors by using n tokens, each originating at a different node and performing a (short) independent random-walk over the graph. It is well known [23] that in every d -regular expander $G = (V, E)$ a lazy random walk has a uniform stationary distribution.

Moreover, it is *rapidly mixing*, i.e., its mixing time is $\bar{t} = O(\log(1/\epsilon) \log n)$ where ϵ is the desired bound on the total variation distance.

The modified dynamics works in synchronous *phases*, each of them consisting of exactly 2τ rounds (the suitable value for τ will be defined later). During the first τ rounds a *forward process* takes place: Every node sends a token performing a random walk of at least \bar{t} -hops and thus sampling the color of a random node. In the next τ rounds we have a *backward process*: Every token is sent back to its source by “reversing” the path followed in the forward process.

If we were in the *LOCAL* model [28], where each node can communicate with all its neighbors in one round, each phase of the above protocol would last exactly $2\bar{t}$ rounds. In the *GOSSIP* model [9], each node can instead activate only one (bidirectional) link per round. Moreover, since we want *messages of limited size*, we assume that through each direction of an active link only one token can be transmitted.

We further assume that nodes enqueue tokens with a *FIFO* policy, breaking ties arbitrarily. The random walk performed by a token will likely require more than \bar{t} rounds to perform (at least) \bar{t} hops of the random walk, depending on the *congestion*, i.e. the maximum number of tokens in the queue of a node. We thus need to bound the maximal congestion and use this bound, together with \bar{t} , to suitably set the right value for τ , so that every random walk is w.h.p. “mixed” enough. At time 2τ each node gets back its own token, and updates its state according to the Undecided-State Dynamics. After that, a new phase starts, and the process iterates. Further important details and remarks about this modified dynamics:

- During the forward process, every token records the link labels of its random-walk and each node records, for any round, the (local) link label it has used (if any) to send a token at that round. Thanks to this information, every node can easily perform the backward process: At every round each node knows (if any) the neighbor it must contact to receive the right token back⁸. Notice that, since the backward process is perfectly specular to the forward one, the congestion is the same in both phases. Hence, both node memory and token message require $\Theta(\tau \log d)$ bits.

- By setting a suitable value for τ , every token will w.h.p. perform at least \bar{t} hops (some tokens may perform more hops than others). Thanks to the rapidly-mixing property, the color reported to the sender belongs to a random node, i.e., each node has probability $1/n \pm \epsilon$ to

⁸Recall that in the *GOSSIP* model [9], agents can indeed contact one *arbitrary* neighbor per round.

be sampled (our analysis works setting $\epsilon = O(1/n^2)$).

In the next paragraph, we provide the main arguments of our congestion analysis (a formal analysis with all the details can be found in Appendix D).

Highlights on the congestion analysis. Let $u \in [n]$ be a node, for every round $t \in [2\tau]$ of a phase, we consider the r.v. \mathcal{Q}_t defined as the number of tokens in u at round t . Consider the number Y_t of tokens received by node u at round t (for brevity's sake, we will omit index u in any r.v.). Then we can write $Y_t = \sum_{i \in [d]} X_{i,t}$ where $X_{i,t} = 1$ if the i -th neighbor of u sends a token to u and 0 otherwise. Observe that the r.v.s $X_{i,t}$ are not mutually independent. However, the crucial fact is that, for any t and any i , $\mathbf{P}(X_{i,t} = 1) \leq 1/d$, regardless of the state of the system (in particular, independently of the value of the other r.v.s). So, if we consider a family $\{\hat{X}_{i,t} : i \in [d] t \in [2\tau]\}$ of i.i.d. Bernoulli r.v.s with $\mathbf{P}(\hat{X}_{i,t} = 1) = 1/d$, then Y_t is stochastically dominated by $\hat{Y}_t = \sum_{i \in [d]} \hat{X}_{i,t}$. For any node u and any round t , the r.v. \mathcal{Q}_t is thus stochastically dominated by the r.v. $\hat{\mathcal{Q}}_t$ defined recursively as follows.

$$\begin{cases} \hat{\mathcal{Q}}_t &= \hat{\mathcal{Q}}_{t-1} + \hat{Y}_t - \chi_t \\ \hat{\mathcal{Q}}_0 &= 1 \end{cases}$$

where $\chi_t = \begin{cases} 1 & \text{if } \hat{\mathcal{Q}}_{t-1} > 0 \\ 0 & \text{otherwise} \end{cases}$

Since our goal is to provide a concentration upper bound on \mathcal{Q}_t , we can do this by considering the ‘‘simpler’’ process $\hat{\mathcal{Q}}_t$. It turns out that ‘‘unrolling’’ $\hat{\mathcal{Q}}_t$ directly is far from trivial: we thus need the ‘‘right’’ way to write it by using only i.i.d. Bernoulli r.v.s. To this aim, for any $t \in [2\tau]$ and for any $s \in [t]$, we define the r.v.

$$(4.6) \quad Z_{s,t} = \sum_{i=s}^t \hat{Y}_i - (t-s)$$

Informally speaking, $Z_{s,t}$ matches the value of $\hat{\mathcal{Q}}_t$ whenever $s \leq t$ was the last previous round s.t. $\hat{\mathcal{Q}}_s = 0$.

As a key fact (see the claim in the proof of Lemma D.1 in the Appendix D), we show that $\hat{\mathcal{Q}}_t$ can be written as a suitable function of $Z_{s,t}$ and χ_t so that it holds

$$(4.7) \quad \hat{\mathcal{Q}}_t \leq \max_{s \in [t]} \{Z_{s,t}\} \text{ and thus } \max_{t \in [2\tau]} \{\mathcal{Q}_t\} \leq \max_{s \leq t \leq 2\tau} \{Z_{s,t}\}$$

From (4.6), the r.v. $Z_{s,t} + (t-s)$ is a sum of $d \cdot (t-s+1)$ i.i.d. Bernoulli r.v.s, each with expectation $1/d$. From Chernoff's bound, it thus follows that, for constant $c > 0$ and any $1 \leq s \leq t \leq 2\tau$ we have

$$\mathbf{P}\left(Z_{s,t} \leq \max\left\{\sqrt{c(t-s+1)\log n}, 3c\log n\right\}\right) \geq 1 - n^{-c/3}$$

By taking the union bound over all $1 \leq s \leq t \leq 2\tau$, from the above bound and (4.7), we get the desired concentration bound on the maximal node congestion during every phase:

$$\mathbf{P}\left(\max_{1 \leq t \leq 2\tau} \mathcal{Q}_t \leq \max\left\{\sqrt{c\tau \log n}, 3c\log n\right\}\right) \geq 1 - \frac{\tau^2}{n^{c/3}}$$

The above congestion bound allows us to set the right value of τ , thus getting the following final result (its proof is given in Appendix D).

THEOREM 4.1. *Let $G = (V, E)$ be any regular expander graph. For any initial configuration $\bar{\mathbf{c}}$ such that the Undecided-State Dynamics on the clique computes plurality consensus in $O(\text{md}(\bar{\mathbf{c}}) \log n)$ rounds w.h.p., the modified Undecided-State Dynamics computes plurality consensus on G in $O(\text{md}(\bar{\mathbf{c}}) \text{polylog}(n))$ rounds w.h.p.*

Remark. Notice that our analysis on the congestion also works in a scenario where every node generates a new token whenever its queue is empty. For this reason our analysis does not take care of the bound n on the overall number of nodes, and thus it is not tight. However, we believe that any improvement of the analysis taking into account this bound is far from trivial.

5 Open Problems

There are several open research directions related to the plurality problem on the gossip model. One of the most interesting (and challenging) ones concerns the monochromatic distance we have introduced in this paper. We believe that this distance might represent a general lower bound on the convergence time of *any* plurality dynamics which uses only $\log k + \Theta(1)$ bits of local memory.

Another interesting future research is the study of the Undecided-State Dynamics (or other simple dynamics) over other classes of graphs. In our paper, we combined this dynamics with parallel random walks in order to get an efficient protocol for regular expander graphs. We believe that similar protocols can work also in other classes of graphs such as Erdős-Rényi graphs and dynamic graphs [12, 10].

References

- [1] M. A. Abdullah and M. Draief. Majority consensus on random graphs of a given degree sequence. *arXiv:1209.5025*, 2012.
- [2] N. Alon, C. Avin, M. Koucky, G. Kozma, Z. Lotker, and M. R. Tuttle. Many random walks are faster than one. In *Proc. of 20th ACM SPAA*, 2008.

- [3] D. Angluin, J. Aspnes, and D. Eisenstat. A simple population protocol for fast robust approximate majority. *Distributed Computing*, 21(2):87–102, 2008. Ext. abs. in DISC 2007.
- [4] A. Babaee and M. Draief. Distributed multivalued consensus. In *Computer and Information Sciences III*, pages 271–279. Springer, 2013.
- [5] L. Becchetti, A. Clementi, E. Natale, F. Pasquale, R. Silvestri, and L. Trevisan. Simple dynamics for plurality consensus. In *Proc. of 26th ACM SPAA*, 2014.
- [6] F. Bénézit, P. Thiran, and M. Vetterli. Interval consensus: from quantized gossip to voting. In *Proc. of 34th IEEE ICASSP*, 2009.
- [7] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, pages 2508–2530, 2006.
- [8] L. Cardelli and A. Csikász-Nagy. The cell cycle switch computes approximate majority. *Scientific Reports*, Vol. 2, 2012.
- [9] K. Censor-Hillel, B. Haeupler, J. Kelner, and P. Maymounkov. Global computation in a poorly connected world: Fast rumor spreading with no dependence on conductance. In *Proc. of 44th ACM STOC*, 2012.
- [10] A. Clementi, P. Crescenzi, C. Doerr, P. Fraigniaud, M. Isopi, A. Panconesi, F. Pasquale, and R. Silvestri. Rumor spreading in random evolving graphs. In *Proc. of 21st ESA*, 2013.
- [11] A. Clementi, M. Di Ianni, G. Gambosi, E. Natale, and R. Silvestri. Distributed community detection in dynamic graphs. In *Proc. of 20th SIROCCO*, 2013.
- [12] A. Clementi, C. Macci, A. Monti, F. Pasquale, and R. Silvestri. Flooding time of edge-Markovian evolving graphs. *SIAM J. Discrete Math.*, 24(4):1694–1712, 2010.
- [13] C. Cooper, R. Elsässer, and T. Radzik. The power of two choices in distributed voting. In *Proc. of 41st ICALP*, 2014.
- [14] A. Das Sarma, A.R. Molla, and G. Pandurangan. Near-optimal random walk sampling in distributed networks. In *Proc. of 31st IEEE INFOCOM*, 2012.
- [15] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *Proc. of 6th ACM PODC*, 1987.
- [16] B. Doerr, L. A. Goldberg, L. Minder, T. Sauerwald, and C. Scheideler. Stabilizing consensus with the power of two choices. In *Proc. of 23rd ACM SPAA*, 2011.
- [17] D. Doty. Timing in chemical reaction networks. In *Proc. of 25th ACM-SIAM SODA*, 2014.
- [18] M. Draief and M. Vojnovic. Convergence speed of binary interval consensus. *SIAM J. on Control and Optimization*, 50(3):1087–1109, 2012.
- [19] P. Fraigniaud, A. Korman, and D. Peleg. Local distributed decision. In *Proc. of 52nd IEEE FOCS*, 2011.
- [20] B. Haeupler, G. Pandurangan, D. Peleg, R. Rajaraman, and Z. Sun. Discovery through gossip. In *Proc. of 24th ACM SPAA*. ACM, 2012.
- [21] S. Hoory, Linial N., and Wigderson W. Expander graphs and their applications. *Bull. Amer. Math. Soc. (N.S)*, 43:439–561, 2006.
- [22] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *Proc. of 43rd IEEE FOCS*, 2003.
- [23] D. Levin, Y. Peres, and E. L. Wilmer. *Markov Chains and Mixing Times*. AMS, 2008.
- [24] G. B. Mertzios, S. E. Nikolettseas, C. Raptopoulos, and P. G. Spirakis. Determining majority in networks with local interactions and very small local memory. In *Proc. of 41st ICALP*, 2014.
- [25] E. Mossel, J. Neeman, and O. Tamuz. Majority dynamics and aggregation of information in social networks. *Autonomous Agents and Multi-Agent Systems*, pages 1–22, 2012.
- [26] E. Mossel, J. Neeman, and O. Tamuz. Majority dynamics and aggregation of information in social networks. *Autonomous Agents and Multi-Agent Systems*, 28(3), 2014.
- [27] E. Mossel and G. Schoenebeck. Reaching consensus on social networks. In *Proc. of 2nd ITCS*, 2010.
- [28] D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, 2000.
- [29] D. Peleg. Local majorities, coalitions and monopolies in graphs: a review. *Theor. Comput. Sci.*, 282(2):231–257, 2002.
- [30] E. Perron, D. Vasudevan, and M. Vojnovic. Using three states for binary consensus on complete graphs. In *Proc. of 28th IEEE INFOCOM*, 2009.
- [31] Y. Ruan and Y. Mostofi. Binary consensus with soft information processing in cooperative networks. In *Proc. of 47th IEEE CDC*, 2008.

Appendix

A Probabilistic bounds

LEMMA A.1. (CHERNOFF BOUND, MULT. FORM)
 Let $\{X_i\}_{i \in [n]}$ be a family of n independent r.v.s, let $\delta \in (0, 1]$ and let μ_1, μ_2 , and μ_3 be such that $\mu_1 \leq \mathbf{E} \left[\sum_{i \in [n]} X_i \right] \leq \mu_2$ and $\mu_3 \geq 6 \cdot \mathbf{E} \left[\sum_{i \in [n]} X_i \right]$. It holds that

$$(A.1) \quad \mathbf{P} \left(\sum_{i \in [n]} X_i \leq (1 - \delta) \cdot \mu_1 \right) \leq e^{-\frac{\delta^2}{2} \cdot \mathbf{E} \left[\sum_{i \in [n]} X_i \right]}$$

$$(A.2) \quad \mathbf{P} \left(\sum_{i \in [n]} X_i \geq (1 + \delta) \cdot \mu_2 \right) \leq e^{-\frac{\delta^2}{3} \cdot \mathbf{E} \left[\sum_{i \in [n]} X_i \right]}$$

$$(A.3) \quad \mathbf{P} \left(\sum_{i \in [n]} X_i \geq \mu_3 \right) \leq 2^{-\mu_3}$$

LEMMA A.2. (CHERNOFF BOUND, ADDITIVE FORM)
 Let X_1, \dots, X_n be a sequence of independent $\{0, 1\}$ r.v.s, let $X = \sum_{i=1}^n X_i$ be their sum, and let $\mu = \mathbf{E}[X]$. Then for $0 < \lambda < 1$ it holds that

$$\begin{aligned} \mathbf{P}(X \geq \mu + n\lambda) &\leq e^{-2n\lambda^2} \\ \mathbf{P}(X \leq \mu - n\lambda) &\leq e^{-2n\lambda^2}. \end{aligned}$$

LEMMA A.3. Let a and b be two constants such that $a > b > 0$, let B be an event and let $\{A_i\}_{i \in I}$ be a family of events such that $|I| = O(n^b)$ and $\mathbf{P}(A_i | B) \geq 1 - n^{-a}$. Then, the event $\bigcap_{i \in I} A_i | B$ holds with probability at least $1 - \frac{|I|}{n^a}$.

LEMMA A.4. If $f(n) = \omega(1)$ and $g(n) = o(f(n))$ then

$$\left(1 \pm \frac{1}{f(n)}\right)^{g(n)} = 1 \pm O\left(\frac{g(n)}{f(n)}\right)$$

B Global bias

Our analysis will highlight a fundamental dependence of convergence properties of the Undecided-State Dynamics on a particular measure of the initial global bias. To mathematically characterize this we next introduce the following notion of distance between *equivalent* color configurations.

Given any color configuration $\mathbf{c} \langle c_1, c_2, \dots, c_k, q \rangle$, consider the following ratio

$$R(\mathbf{c}) = \sum_{i=1}^k \frac{c_i}{c_1}$$

This allows us to define an equivalence relation \equiv in the space \mathcal{S}

$$\mathbf{c} \equiv \mathbf{c}' \quad \text{iff} \quad R(\mathbf{c}) = R(\mathbf{c}')$$

and the following function over pairs of equivalence classes (with an abuse of notation, for any color configuration \mathbf{c} , we will denote its equivalence class as \mathbf{c} as well)

$$d(\mathbf{c}, \mathbf{c}') = \sum_i \left(\frac{c_i}{c_1} - \frac{c'_i}{c'_1} \right)^2$$

It is easy to verify that the function $d(\cdot, \cdot)$ is a distance over the quotient space of \mathcal{S} . Let us now consider the equivalence class \mathcal{M} of the (k) possible *monochromatic* color configurations and recall the definition of *monochromatic distance* (given in the introduction),

$$\text{md}(\mathbf{c}) = \sum_{i=1}^k \left(\frac{c_i}{c_1} \right)^2$$

Then, we immediatly have

$$\text{md}(\mathbf{c}) = d(\mathbf{c}, \mathcal{M}) + 1$$

The simple considerations above entail that md defines a notion of distance from the monochromatic configuration that corresponds to the initial plurality. Consistently, it is straightforward to see that md is maximized by “uniform” configurations, i.e., configurations \mathbf{c} such that $c_1 \approx n/k$. For every \mathbf{c} , it holds that

$$(B.4) \quad 1 \leq R(\mathbf{c}), \text{md}(\mathbf{c}) \leq k$$

Finally, let us define the following ratio

$$\Lambda(\mathbf{c}) := \frac{R(\mathbf{c})^2}{\text{md}(\mathbf{c})}$$

From the definitions of $R(\mathbf{c})$ and $\text{md}(\mathbf{c})$ and from a simple application of the Cauchy-Schwartz inequality to $R(\mathbf{c})$, we get

$$(B.5) \quad \Lambda(\mathbf{c}) \leq k$$

for every configuration \mathbf{c} .

C Details of the analysis

In this section we give detailed proofs for the process analysis sketched in Section 3.

C.1 General bounds. We here provide some crucial properties that hold along the entire process.

If $\mathbf{c} = \langle c_1, \dots, c_k, q \rangle$ is the current color configuration (i.e. state) of the Markov chain, then we can easily derive the “expectation” of the next color configuration

$$(C.6) \quad \mu_i = \mathbf{E}[C'_i | \bar{\mathbf{c}}] = c_i \cdot \frac{c_i + 2q}{n} \quad (i \in [k])$$

$$(C.7) \quad \begin{aligned} \mu_q &= \mathbf{E}[Q' | \bar{\mathbf{c}}] = \frac{q^2 + \sum_{i \neq j} c_i \cdot c_j}{n} \\ &= \frac{q^2 + (n - q)^2 - \sum_i c_i^2}{n} \end{aligned}$$

From (C.6), we can see the crucial role of the quantity $\frac{c_i + 2q}{n}$: it in fact represents the expected *growth rate* of every color community. The following lemma in fact formalizes such a connection by means of $R(\mathbf{c})$ and it plays a key role in our analysis of the entire process evolution. As will see in Lemma C.4, $R(\mathbf{c})$ and $\text{md}(\mathbf{c})$ are in fact strongly related.

LEMMA C.1. (PLURALITY DRIFT) *Assume that, at some round, the system is in a color configuration \mathbf{c}*

such that $c_1 \geq (1 + \alpha)c_i$ for any $i \neq 1$ and for some constant $\alpha > 0$. Then, at the next round, it holds that

$$\mathbf{E} \left[\frac{C'_1 + 2Q'}{n} \mid \mathbf{c} \right] \geq 1 + \Gamma(\mathbf{c})$$

where

$$\Gamma(\mathbf{c}) = \left(1 - \frac{c_1 + 2q}{n}\right)^2 + 2(1 - \gamma)(R(\mathbf{c}) - 1) \left(\frac{c_1}{n}\right)^2$$

with $\gamma = (1 + \alpha)^{-1}$.

Proof. Let $\beta = (1 - \gamma)$. By using the hypothesis

$$\frac{c_i}{c_1} \leq \frac{1}{(1 + \alpha)}$$

we get

$$\text{md}(\mathbf{c}) = \sum_i \frac{c_i^2}{c_1^2} \leq 1 + \frac{1}{(1 + \alpha)} \sum_{i \neq 1} \frac{c_i}{c_1} = \gamma R(\mathbf{c}) + \beta$$

Moreover, we can write q as $q = n - R(\mathbf{c})c_1$. Thanks to the above equations and (C.6) and (C.7), by simple manipulations, we get

$$\begin{aligned} \mathbf{E} \left[\frac{C'_1 + 2Q'}{n} \mid \mathbf{c} \right] &= c_1 \cdot \frac{c_1 + 2q}{n^2} + 2 \frac{q^2 + (n - q)^2 - \sum_i (c_i)^2}{n^2} \\ &= c_1 \cdot \frac{c_1 + 2q}{n^2} + 2 \frac{q^2 + (R(\mathbf{c})^2 - \text{md}(\mathbf{c})) \cdot (c_1)^2}{n^2} \\ &\geq c_1 \cdot \frac{c_1 + 2q}{n^2} + 2 \frac{q^2 + (R(\mathbf{c})^2 - \gamma R(\mathbf{c}) - \beta) \cdot (c_1)^2}{n^2} \\ &= 1 + \left(1 - \frac{c_1 + 2q}{n}\right)^2 + 2(1 - \gamma)(R(\mathbf{c}) - 1) \frac{c_1^2}{n^2} \end{aligned}$$

□

Another useful property that is often used in our analysis is the fact that some crucial r.v.s are essentially monotone along the entire process. In the next lemma, we prove this monotonicity for the r.v.s $R(\mathbf{C}')$ and the ratios C'_i/C'_1 (for $i \neq 1$).

LEMMA C.2. (MONOTONICITY) *Assume that, at some round, the system is in a color configuration \mathbf{c} such that, for some constant $\alpha > 0$ and a large enough constant $\lambda > 0$ it holds*

$$c_1 \geq (1 + \alpha)c_i \text{ for any } i \neq 1 \text{ and } \mu_1 \geq \lambda \log n$$

Then, at the next round, w.h.p. it holds that:

$$(C.8) \quad R(\mathbf{C}') < R(\mathbf{c}) \cdot \left(1 + O\left(\sqrt{\frac{\log n}{\mu_1}}\right)\right)$$

$$(C.9) \quad C'_1 \geq (1 + \alpha) \cdot C'_i \cdot \left(1 - O\left(\sqrt{\frac{\log n}{\mu_1}}\right)\right)$$

Proof. As for Claim (C.8), since $R(\mathbf{C}') = \frac{\sum_i C'_i}{C'_1}$, it suffices to bound, respectively, C'_1 and $\sum_i C'_i$. By applying the Chernoff bounds (A.1) and (A.2) and by using the hypothesis $\mu \geq \mu_1 \geq \lambda \log n$ we get

$$(C.10) \quad \mathbf{P} \left(C'_1 \leq \mu_1 \cdot \left(1 - \sqrt{\frac{2a \cdot \log n}{\mu_1}}\right) \mid \mathbf{c} \right) \leq \frac{1}{n^a}$$

$$(C.11) \quad \mathbf{P} \left(C'_1 \geq \mu_1 \cdot \left(1 + \sqrt{\frac{3a \log n}{\mu_1}}\right) \mid \mathbf{c} \right) \leq \frac{1}{n^a}$$

$$(C.12) \quad \mathbf{P} \left(\sum_i C'_i \geq \mu \cdot \left(1 + \sqrt{\frac{3a \log n}{\mu}}\right) \mid \mathbf{c} \right) \leq \frac{1}{n^a}$$

for any constant $a \in (0, \frac{\lambda}{3})$.

Let A be the event in (C.10), let B be the event in (C.12) and let A^c and B^c be their complimentary events, respectively. Observe that, from Lemma A.3, it holds $\mathbf{P}(A^c \cap B^c) \geq 1 - \frac{2}{n^a}$. Moreover, by using that

$$\frac{1 + \sqrt{\frac{3a \log n}{\mu}}}{1 - \sqrt{\frac{2a \log n}{\mu_1}}} \leq \frac{1 + \sqrt{\frac{3a \log n}{\lambda \log n}}}{1 - \sqrt{\frac{2a \log n}{\lambda \log n}}} \leq 1 + \sqrt{\frac{ba \log n}{\lambda \log n}}$$

where $b = \left(\frac{\sqrt{3} - \sqrt{2}}{1 - 3\sqrt{2}a}\right)^2$. Using the latter two facts, we have that

$$\begin{aligned} \mathbf{P} \left(R(\mathbf{C}') = \frac{\sum_i C'_i}{C'_1} < \frac{\sum_i c_i}{c_1} \cdot \left(1 + \sqrt{\frac{ba \log n}{\mu}}\right) \mid \mathbf{c} \right) &\geq \\ &\geq \mathbf{P} \left(\frac{\sum_i C'_i}{C'_1} < \frac{\sum_i c_i \cdot (c_i + q)}{c_1 \cdot (c_1 + q)} \cdot \left(1 + \sqrt{\frac{ba \log n}{\mu}}\right) \mid \mathbf{c} \right) = \\ &= \mathbf{P} \left(\frac{\sum_i C'_i}{C'_1} < \frac{\mu}{\mu_1} \cdot \left(1 + \sqrt{\frac{ba \log n}{\mu}}\right) \mid \mathbf{c} \right) \geq \\ &\geq \mathbf{P} \left(\frac{\sum_i C'_i}{C'_1} < \frac{\mu \cdot \left(1 + \sqrt{\frac{3a \log n}{\mu}}\right)}{\mu_1 \cdot \left(1 - \sqrt{\frac{2a \log n}{\mu_1}}\right)} \mid \mathbf{c} \right) \\ &\geq \mathbf{P}(A^c \cap B^c) \geq 1 - \frac{2}{n^a} \end{aligned}$$

As for Claim (C.9), the hypothesis $c_1 \geq (1 + \alpha)c_i$ clearly implies $\mu_1 \geq (1 + \alpha) \cdot \mu_i$. Thus, by (C.10) we get

$$(C.13) \quad \mathbf{P} \left(C'_1 \leq (1 + \alpha) \cdot \mu_i \cdot \left(1 - \sqrt{\frac{2a \log n}{\mu_1}}\right) \mid \mathbf{c} \right) \leq \mathbf{P} \left(C'_1 \leq \mu_1 \cdot \left(1 - \sqrt{\frac{2a \log n}{\mu_1}}\right) \mid \mathbf{c} \right) \leq \frac{1}{n^a}$$

We now consider two cases. If $\mu_i < \mu_1/(6(1+\alpha))$ then, by Chernoff bound (A.3) (choosing $\delta = \mu_1/(1+\alpha)$), with probability $1 - n^{-\frac{\lambda}{1+\alpha}}$ it holds that $C'_i \leq \mu_1/(1+\alpha)$. Together with (C.10), this implies that w.h.p.

$$C'_1 > \mu_1 \cdot \left(1 - \sqrt{\frac{2a \log n}{\mu_1}}\right) > (1+\alpha) C'_i \cdot \left(1 - \sqrt{\frac{2a \log n}{\mu_1}}\right)$$

On the other hand, if $\mu_i \geq \mu_1/(6(1+\alpha))$ then, from the Chernoff bound (A.1) we get that

$$(C.14) \quad \mathbf{P} \left(C'_i \geq \mu_i \cdot \left(1 + \sqrt{\frac{3a \log n}{\mu_i}}\right) \mid \mathbf{c} \right) \\ \leq \mathbf{P} \left(C'_i \geq \mu_i \cdot \left(1 + \sqrt{\frac{3a \log n}{\mu_i/6(1+\alpha)}}\right) \mid \mathbf{c} \right) \leq \frac{1}{n^a}$$

for any $a \in \left(0, \frac{\lambda}{18(1+\alpha)}\right)$. Thus, by using (C.13), (C.14) and Lemma A.4 we get that w.h.p.

$$C'_1 \geq (1+\alpha) \cdot C'_i \cdot \left(1 - O\left(\sqrt{\frac{\log n}{\mu_1}}\right)\right)$$

□

C.2 First Round: Rise of the undecided. After the first round, a strong decrease of the color communities happens, while the undecided community get to a large majority of the agent.

The next lemmas provide some formal statements about this behaviour which represent the key start-up of the process (and its analysis).

We will implicitly assume that the process starts in a fixed initial color configuration

$$\bar{\mathbf{c}} = \langle \bar{c}_1, \bar{c}_2, \dots, \bar{c}_k \rangle$$

So, in the next lemmas, events and related probabilities are conditioned on some fixed $\bar{\mathbf{c}}$.

We observe that when k is large, i.e. when $k = \omega(n^b)$ for some $b \in (\frac{1}{2}, 1]$, if the process starts from “almost-uniform” color configurations then, after the first round, even the plurality may disappear (w.h.p.): indeed, if we consider any $\bar{\mathbf{c}}$ such that $\bar{c}_1 = O\left(\frac{n}{k}\right)$, then a simple application of the Markov inequality implies that $C'_1 = 0$ w.h.p. We will thus focus on ranges of k such that $k < \sqrt{n/\log n}$.

LEMMA C.3. *Let $k = o\left(\sqrt{n/\log n}\right)$. Given any initial color configuration $\bar{\mathbf{c}}$, after the first round w.h.p. it*

holds:

$$\frac{1}{2} \frac{n}{R(\bar{\mathbf{c}})^2} \leq C'_1 \leq 2 \frac{n}{R(\bar{\mathbf{c}})^2} \\ n \left(1 - \frac{2}{\Lambda(\bar{\mathbf{c}})}\right) \leq Q' \leq n \left(1 - \frac{1}{2\Lambda(\bar{\mathbf{c}})}\right)$$

Proof. From (C.6) and recalling that in the initial configuration $q = 0$, we get

$$\mu_1 = \frac{(\bar{c}_1)^2}{n} = \frac{n}{R(\bar{\mathbf{c}})^2}$$

Similarly, from (C.7) we get

$$\mu_q = \frac{n^2 - \sum_i (\bar{c}_i)^2}{n} = \frac{n^2 - \text{md} \cdot (\bar{c}_1)^2}{n} = n \left(1 - \frac{1}{\Lambda(\bar{\mathbf{c}})}\right)$$

where the second equality follows from the definition of md , while the third one from the definition of $R(\bar{\mathbf{c}})$ and from simple manipulations. Since we assumed $k \leq o\left(\sqrt{n/\log n}\right)$ then we have that

$$\mu_q = \frac{n}{R(\bar{\mathbf{c}})^2} \geq \frac{n}{k^2} = \omega(\log n)$$

The above inequality allows us to apply the Chernoff bound and prove the first claim (i.e. that on C'_1).

Similarly, from (B.5), it holds

$$\frac{n}{\Lambda(\bar{\mathbf{c}})} \geq \frac{n}{k}$$

This allows us to apply the additive version of the Chernoff’s bound and prove the second claim (i.e. that on Q'). □

The next lemma relates $R(\mathbf{c})$ to $\text{md}(\bar{\mathbf{c}})$ after the first round.

LEMMA C.4. *Let $k = o\left(\sqrt{n/\log n}\right)$. Given any initial color configuration $\bar{\mathbf{c}}$, after the first round w.h.p. it holds*

$$R(\mathbf{C}^{(1)}) \leq \text{md}(\bar{\mathbf{c}}) \cdot (1 + o(1))$$

Proof. By definition of plurality color, it holds that $c_1 > n/k$. Therefore, by the hypothesis on k and (C.6), we get $\mu_1 = \omega(\log n)$ and then, by using the Chernoff bounds of Lemma A.1, we can get concentration bounds on both the numerator and the denominator of $R(\mathbf{C}^{(1)})$ (as we did in the proof of Lemma C.2). Formally, we have that w.h.p.

$$R(\mathbf{C}^{(1)}) = \frac{\sum_i C_i^{(1)}}{C_1^{(1)}} \leq \frac{\mu}{\mu_1} \cdot (1 + o(1))$$

Observe that, since in the initial color configuration $q = 0$, it holds

$$\frac{\mu}{\mu_1} = \frac{\sum_i (\bar{c}_i)^2}{(\bar{c}_1)^2}$$

It follows that w.h.p.

$$\begin{aligned} R(\mathbf{C}^{(1)}) &\leq \frac{\mu}{\mu_1} \cdot (1 + o(1)) = \frac{\sum_i (\bar{c}_i)^2}{(\bar{c}_1)^2} \cdot (1 + o(1)) \\ &= \text{md} \cdot (1 + o(1)) \end{aligned}$$

□

C.3 First phase: Age of the undecided. In this phase, the undecided community rapidly decreases to a value close to $n/2$ while the plurality reaches a size close to $n/(2\text{md})$. When this happens, the ratios C_i/C_1 and $R(\mathbf{c})$ will essentially keep their initial values and the Q will decrease to a value very close to $n/2$. The length of this phase is at most logarithmic.

The next lemma formalizes the aspects of this phase that will be used to get the upper bound on the convergence time of the process.

LEMMA C.5. *Let $k = o\left(\sqrt{n/\log^2 n}\right)$ and let ϵ be any constant in $(0, \frac{1}{2})$. Let $\bar{\mathbf{c}}$ be any initial configuration such that, for any $j \neq 1$ and for some arbitrarily small constant $\alpha > 0$, $c_1 \geq (1 + \alpha) \cdot c_j$. Then w.h.p. at some round $\tilde{t} = O(\log n)$ the process reaches a configuration $\mathbf{C}^{(\tilde{t})}$ such that:*

$$\begin{aligned} \text{(C.15)} \quad &\left\{ C_1^{(\tilde{t})} \geq \left(\frac{1}{16} - \frac{\epsilon}{8}\right) \frac{n}{R(\mathbf{C}^{(\tilde{t})})} \right. \\ \text{(C.16)} \quad &\left. R(\mathbf{C}^{(\tilde{t})}) \leq \text{md} \cdot (1 + o(1)) \right. \\ \text{(C.17)} \quad &\left. C_1^{(\tilde{t})} \geq \left(1 + \frac{\alpha}{2}\right) \cdot C_i^{(\tilde{t})} \text{ for any color } i \neq 1 \right. \\ \text{(C.18)} \quad &\left. \frac{C_1^{(\tilde{t})} + 2Q^{(\tilde{t})}}{n} > 1 + \frac{\epsilon^2}{4} \right. \end{aligned}$$

Proof. We prove one claim at a time.

Proof of (C.15). Let $\tilde{\epsilon}$ be any positive constant in $(\epsilon/2, \epsilon)$.

Two cases may arise. If $\bar{c}_1 > \left(\frac{1}{4} - \frac{\tilde{\epsilon}}{2}\right) \cdot n$, by applying the Chernoff bound (A.1) on the expected value of $C_1^{(1)}$ and using (B.4), it is easy to see that w.h.p.

$$C_1^{(1)} \geq \left(\frac{1}{16} - \frac{\epsilon}{8}\right) n \geq \left(\frac{1}{16} - \frac{\epsilon}{8}\right) \frac{n}{R(\mathbf{C}^{(1)})}$$

Assume now $\bar{c}_1 \leq \left(\frac{1}{4} - \frac{\tilde{\epsilon}}{2}\right) \cdot n$. From Lemma C.3 at round $t = 1$ we have w.h.p.

$$Q^{(1)} \geq n \left(1 - \frac{2}{\Lambda(\bar{\mathbf{c}})}\right) \geq n \left(1 - \frac{2c_1}{n}\right) \geq \frac{n}{2} + \tilde{\epsilon} \cdot n$$

where we used that $\Lambda(\bar{\mathbf{c}}) \geq R(\bar{\mathbf{c}}) = n/\bar{c}_1$.

In the generic configuration \mathbf{c} , as long as $q \geq \frac{n}{2} + \tilde{\epsilon} \cdot n$, from (C.6) we have

$$\mu_1 \geq c_1 \cdot \left(\frac{1}{2} + \tilde{\epsilon}\right)$$

thus, by applying the Chernoff bound (A.1), we see that w.h.p. C_1 grows exponentially fast.

It follows that we can consider the first round such that $\tilde{t} = O(\log n)$ and $Q^{(\tilde{t})} < \frac{n}{2} + \tilde{\epsilon} \cdot n$. This implies that

$$n - Q^{(\tilde{t})} \geq \frac{n}{2} - \tilde{\epsilon} \cdot n$$

hence

$$C_1^{(\tilde{t})} = \frac{n - Q^{(\tilde{t})}}{R(\mathbf{C}^{(\tilde{t})})} \geq \frac{\frac{n}{2} - \tilde{\epsilon} \cdot n}{R(\mathbf{C}^{(\tilde{t})})}$$

This proves (C.15).

Proof of (C.16). Observe that, since $\bar{c}_1 \geq \frac{n}{k}$, then from (C.6) and the Chernoff bound (A.1) it holds w.h.p. that $C_1^{(1)} = \omega(\log^2 n)$. As we have already shown in the proof of Claim (C.15), after the first round C_1 grows exponentially until round \tilde{t} . It follows that we can repeatedly apply Lemma C.2 and, together with Lemma C.4, we get w.h.p. holds w.h.p. that

$$R(\mathbf{C}^{(\tilde{t})}) \leq \text{md} \cdot \left(1 + o\left(\frac{1}{\log n}\right)\right)^{\log n} \leq \text{md} \cdot (1 + o(1))$$

This proves (C.16).

Proof of (C.17). Similarly to the previous Claim proof, the repeated application of Lemma C.2 until round \tilde{t} and Lemma A.4 implies that w.h.p.

$$\begin{aligned} C_1^{(\tilde{t})} &\geq (1 + \alpha) \cdot C_i^{(\tilde{t})} \cdot \left(1 - o\left(\frac{1}{\log n}\right)\right)^{\log n} \\ &= (1 + \alpha) \cdot C_i^{(\tilde{t})} \cdot (1 - o(1)) \geq \left(1 + \frac{\alpha}{2}\right) \cdot C_i^{(\tilde{t})} \end{aligned}$$

This proves (C.17).

Proof of (C.18). Since, by the definition of \tilde{t} , it holds $q^{(\tilde{t}-1)} \geq \frac{n}{2} + \tilde{\epsilon}$, then by Lemma C.1 we get that

$$\mathbf{E} \left[C_1^{(\tilde{t})} + 2Q^{(\tilde{t})} \mid \mathbf{c}^{(\tilde{t}-1)} \right] \geq (1 + \tilde{\epsilon}^2) \cdot n$$

Observe that $\mathbf{E} \left[C_1^{(\tilde{t})} + 2Q^{(\tilde{t})} \mid \mathbf{c}^{(\tilde{t}-1)} \right]$ can be written as the expected value of the sum of the following independent r.v.s: given a color configuration $\mathbf{c}^{(\tilde{t}-1)}$, for each node i

$$X_i = \begin{cases} 1 & \text{if node } i \text{ is 1-colored at the next round,} \\ 2 & \text{if node } i \text{ is undecided at the next round.} \end{cases}$$

Then (C.18) is an easy application of the Chernoff bound (A.1). □

From the state conditions achieved after the first round (see Lemma C.3), the next lemma shows that, within $O(\log n)$ rounds, the process w.h.p. reaches a configuration where Q gets very close to $n/2$ and C_1 is still relatively small. In the next section, we will prove (see Theorem C.1) that this fact forces the process to “wait” for a time period $\Omega(\text{md}(\bar{\mathbf{c}}))$ before the plurality (re-)starts to grow fastly. This is the key ingredient of the lower bound in Theorem C.1.

LEMMA C.6. *Let $k \leq \varepsilon \cdot (n/\log n)^{1/6}$ be the initial number of colors, where $\varepsilon > 0$ is a sufficiently small positive constant. Let $\bar{\mathbf{c}}$ be the initial color configuration and let $\mathbf{c}^{(1)}$ be the color configuration after the first round. If it holds that:*

$$\begin{aligned} \frac{1}{2} \frac{n}{R(\bar{\mathbf{c}})^2} &\leq c_1^{(1)} \leq 2 \frac{n}{R(\bar{\mathbf{c}})^2} \\ n \left(1 - \frac{2}{\Lambda(\bar{\mathbf{c}})}\right) &\leq q^{(1)} \leq n \left(1 - \frac{1}{2\Lambda(\bar{\mathbf{c}})}\right) \end{aligned}$$

within the next $O(\log n)$ rounds there will be a round \bar{t} such that

$$C_1^{(\bar{t})} \leq \gamma \frac{n}{\text{md}(\bar{\mathbf{c}})} \quad \text{and} \quad \left|Q^{(\bar{t})} - \frac{n}{2}\right| \leq 2 \frac{\gamma^2}{\text{md}(\bar{\mathbf{c}})}$$

w.h.p., where $\gamma > 0$ is a sufficiently large constant.

Proof. First, we prove that if at an arbitrary round t the number of undecided nodes is $q = (1 + \delta)(n/2)$ with $1/\text{md}(\bar{\mathbf{c}}) \leq \delta \leq 1 - (2\Lambda(\bar{\mathbf{c}}))^{-1}$, then at the next round it holds that $Q' \leq (1 + \delta^2)(n/2)$ w.h.p. Indeed, if we replace $q = (1 + \delta)(n/2)$ in (C.7), we get that the expected value of Q' at the next round is

$$\begin{aligned} \mu_q &= \frac{1}{n} \left(\left((1 + \delta) \frac{n}{2} \right)^2 + \left((1 + \delta) \frac{n}{2} \right)^2 - \sum_{j=1}^k (c_j)^2 \right) \\ &= (1 + \delta^2) \frac{n}{2} - \frac{1}{n} \sum_{j=1}^k (c_j)^2 \end{aligned}$$

Now observe that

$$\begin{aligned} \frac{1}{n} \sum_{j=1}^k (c_j)^2 &\geq \frac{1}{n} k \left(\frac{n - q}{k} \right)^2 = \frac{n}{4k} (1 - \delta)^2 \\ &\geq \frac{n}{4k} \cdot \left(\frac{1}{2\Lambda(\bar{\mathbf{c}})} \right)^2 \geq \frac{n}{16k^3} \end{aligned}$$

where in the last inequality we used (B.5), that is $\Lambda(\bar{\mathbf{c}}) \leq k$.

Therefore, since Q' is a sum of independent Bernoulli r.v., from the Chernoff bound (Lemma A.2 with $\lambda = 1/16k^3$) it follows that

$$\begin{aligned} \text{(C.19)} \quad \mathbf{P} \left(Q' \geq (1 + \delta^2) \frac{n}{2} \mid \mathbf{c} \right) &\leq \exp \left(-\frac{n}{128k^6} \right) \\ &\leq n^{-1/(128\varepsilon^6)} \end{aligned}$$

where in the last inequality we used the hypothesis on k .

Now we show that the number Q of undecided nodes, while decreasing quickly, cannot jump over the whole interval

$$\left[\frac{n}{2} - 2\gamma^2 \frac{n}{\text{md}(\bar{\mathbf{c}})}, \frac{n}{2} + 2\gamma^2 \frac{n}{\text{md}(\bar{\mathbf{c}})} \right]$$

Observe that function $f(q) = q^2 + (n - q)^2$ has a minimum for $q = n/2$, so for any $q \geq n/2 + 2\gamma^2 n/\text{md}(\bar{\mathbf{c}})$ it holds that $f(q) \geq f(n/2 + 2\gamma^2 n/\text{md}(\bar{\mathbf{c}}))$. Hence if at some round t we have that $q \geq (n/2)(1 + 4\gamma^2/\text{md}(\bar{\mathbf{c}}))$ and $c_1 \leq \gamma n/\text{md}(\bar{\mathbf{c}})$, in (C.7) we get

$$\begin{aligned} \mu_q &\geq \frac{1}{n} \left(\left(\frac{n}{2} + 2\gamma^2 \frac{n}{\text{md}(\bar{\mathbf{c}})} \right)^2 + \left(\frac{n}{2} + 2\gamma^2 \frac{n}{\text{md}(\bar{\mathbf{c}})} \right)^2 - \sum_{j=1}^k c_j^2 \right) \\ &= \frac{n}{2} + 4\gamma^4 \frac{n}{\text{md}(\bar{\mathbf{c}})^2} - \frac{1}{n} \sum_{j=1}^k (c_j)^2 \\ &\geq \frac{n}{2} - \frac{1}{n} \sum_{j=1}^k (c_j)^2 = \frac{n}{2} - \frac{(c_1)^2 \text{md}(\bar{\mathbf{c}})}{n} \geq \frac{n}{2} - \gamma^2 \frac{n}{\text{md}(\bar{\mathbf{c}})} \end{aligned}$$

where in the last inequality we used that $c_1 \leq \gamma n/\text{md}(\bar{\mathbf{c}})$. Since Q' is a sum of n independent Bernoulli r.v., from Chernoff bound it follows that

$$\begin{aligned} \text{(C.20)} \quad \mathbf{P} \left(Q' \leq n/2 - 2\gamma^2 n/\text{md}(\bar{\mathbf{c}}) \mid \mathbf{c} \right) &\leq \\ &\leq \exp \left(-2\gamma^2 \frac{n}{\text{md}(\bar{\mathbf{c}})^2} \right) \leq \exp \left(-2\gamma^2 \frac{n}{k^2} \right) \\ &\leq \exp \left(-\Omega \left(n^{2/3} \right) \right) \end{aligned}$$

From (C.19), we get that w.h.p.

$$\text{(C.21)} \quad Q^{(t)} \leq \left(1 + \delta^{2^t}\right) \frac{n}{2}$$

Hence, within

$$\log(\Lambda(\bar{\mathbf{c}})) + O(\log \log \text{md}(\bar{\mathbf{c}}))$$

rounds, the number Q of undecided nodes will be below $(n/2)(1 + 4\gamma^2/\text{md}(\bar{\mathbf{c}}))$ w.h.p. Moreover, from (C.20) it follows that in one of such rounds we will have that

$$\left| Q - \frac{n}{2} \right| \leq 2\gamma^2 n/\text{md}(\bar{\mathbf{c}})$$

w.h.p. It remains to show that, during this time, the plurality C_1 does not increase from less $2n/R(\bar{\mathbf{c}})^2$ to more than $\gamma n/\text{md}(\bar{\mathbf{c}})$.

To simplify notation, let us define

$$\begin{aligned} l &= \log(\Lambda(\bar{\mathbf{c}})) \\ L &= \log(\Lambda(\bar{\mathbf{c}})) + O(\log \log \text{md}(\bar{\mathbf{c}})) \end{aligned}$$

From (C.6) and (C.21) it follows that, as long as $c_1 \leq \gamma n / \text{md}(\bar{\mathbf{c}})$, the increasing rate of C_1 at round t is w.h.p. at most

$$1 + \delta^{2^t} + \frac{\gamma}{\text{md}(\bar{\mathbf{c}})}$$

For the first l rounds, we can bound the above increasing rate with 2. Thus, after l rounds we get that the plurality is $C_1 \leq 2n / \text{md}(\bar{\mathbf{c}})$ w.h.p. As for the next $O(\log \log \text{md}(\bar{\mathbf{c}}))$ rounds, we have that the plurality is w.h.p. at most

$$\begin{aligned} & 2 \frac{n}{\text{md}(\bar{\mathbf{c}})} \cdot \prod_{t=l}^L \left(1 + \delta^{2^t} + \frac{\gamma}{\text{md}(\bar{\mathbf{c}})} \right) \leq \\ & \leq 2 \frac{n}{\text{md}(\bar{\mathbf{c}})} \cdot \exp \left(\sum_{t=l}^L \left(\delta^{2^t} + \frac{\gamma}{\text{md}(\bar{\mathbf{c}})} \right) \right) \\ & \leq 2 \frac{n}{\text{md}(\bar{\mathbf{c}})} \cdot \exp \left(O(1) + \frac{\log \log \text{md}(\bar{\mathbf{c}})}{\text{md}(\bar{\mathbf{c}})} \right) \\ & \leq \gamma \frac{n}{\text{md}(\bar{\mathbf{c}})} \end{aligned}$$

where in the last inequality we need to choose γ sufficiently large. \square

Remark. The two lemmas above refer to some rounds $\tilde{t}, \bar{t} = O(\log n)$ in which the process lies in a state satisfying certain properties. We observe that our analysis does never combine the two lemmas and thus it does not require that $\tilde{t} = \bar{t}$, indeed the first lemma is used to get the upper bound while the second one to get the lower bound on the convergence time. However, it is possible to prove that there is in fact a time interval (at the end of Phase 2) where both claims of the lemmas hold w.h.p.

C.4 Second phase: *Plateau or Age of stability.*

This phase is characterized by a slow increase of c_1 , roughly at a rate $1 + \Theta(1 / \text{md}(\bar{\mathbf{c}}))$. This fact is formalized in the next lemma and it will be used to derive the lower bound on the convergence time of the process in Theorem C.1.

LEMMA C.7. *Let $\bar{\mathbf{c}}$ be the initial color configuration, let $k \leq \varepsilon \cdot (n / \log n)^{1/4}$ be the initial number of colors, where $\varepsilon > 0$ is a sufficiently small positive constant. If there is a round \bar{t} such that*

$$\left| q^{(\bar{t})} - \frac{n}{2} \right| \leq 2\gamma^2 \frac{n}{\text{md}(\bar{\mathbf{c}})} \quad \text{and} \quad c_1^{(\bar{t})} \leq \gamma(n / \text{md}(\bar{\mathbf{c}}))$$

(where γ is an arbitrary positive constant), then the plurality C_1 remains smaller than $2\gamma(n / \text{md}(\bar{\mathbf{c}}))$ for the next $\Omega(\text{md}(\bar{\mathbf{c}}))$ rounds w.h.p.

Proof. Let us define $\delta = q - n/2$ and let Δ' be the random variable $Q' - n/2$ in the next round. From (C.6) we get

$$(C.22) \quad \mathbf{E}[\Delta' | \mathbf{c}] = \frac{1}{n} \left(2\delta^2 - \sum_{j=1}^k (c_j)^2 \right)$$

$$(C.23) \quad \mu_i = \left(1 + \frac{2\delta + c_i}{n} \right) c_i$$

We now show that, if $\delta \in (-2\gamma^2 n / \text{md}(\bar{\mathbf{c}}), 2\gamma^2 n / \text{md}(\bar{\mathbf{c}}))$ and $c_1 \leq 2\gamma n / \text{md}(\bar{\mathbf{c}})$, then the increasing rate of C_1 is smaller than $(1 + \Theta(1 / \text{md}(\bar{\mathbf{c}})))$ w.h.p. More precisely, we prove that w.h.p.

$$\begin{cases} |\delta| \leq 2\gamma^2 \frac{n}{\text{md}(\bar{\mathbf{c}})} \\ c_1 \leq 2\gamma \frac{n}{\text{md}(\bar{\mathbf{c}})} \end{cases} \implies \begin{cases} |\Delta'| \leq 2\gamma^2 \frac{n}{\text{md}(\bar{\mathbf{c}})} \\ C'_1 \leq \left(1 + \frac{2\gamma(\gamma+1)+1}{\text{md}(\bar{\mathbf{c}})} \right) c_1 \end{cases}$$

As for the increasing rate of the plurality, from (C.23) it follows that

$$\begin{aligned} \mu_1 &= \left(1 + \frac{2\delta + c_1}{n} \right) c_1 \\ &\leq \left(1 + \frac{2\gamma^2 n / \text{md}(\bar{\mathbf{c}}) + 2\gamma n / \text{md}(\bar{\mathbf{c}})}{n} \right) c_1 \\ &= \left(1 + \frac{2\gamma(\gamma+1)}{\text{md}(\bar{\mathbf{c}})} \right) c_1 \end{aligned}$$

Since C'_1 can be written as a sum of $q + c_1 \leq n$ independent Bernoulli random variables, from Chernoff bound (Lemma A.2 with $\lambda = c_1 / (n \text{md}(\bar{\mathbf{c}}))$) it follows that

$$(C.24) \quad \mathbf{P} \left(C_1 \geq \left(1 + \frac{2\gamma(1+\gamma)+1}{\text{md}(\bar{\mathbf{c}})} \right) c_1 \mid \mathbf{c} \right) \leq \\ \leq \exp \left(- \frac{2(c_1 / \text{md}(\bar{\mathbf{c}}))^2}{n} \right) \\ \leq \exp \left(- \frac{2n}{9k^4} \right) \leq n^{-2/(9\varepsilon^4)}$$

where in the second inequality we used the fact that $c_1 \geq n - q/k \geq n/(3k)$ and $\text{md}(\bar{\mathbf{c}}) \leq k$, and in the last inequality we used the hypothesis $k \leq \varepsilon \cdot (n / \log n)^{1/4}$.

As for $\mathbf{E}[\Delta' | \mathbf{c}]$, according to (C.22), we have the upper bound

$$(C.25) \quad \mathbf{E}[\Delta' | \mathbf{c}] \leq 2 \frac{\delta^2}{n} \leq 8\gamma^4 \frac{n}{(\text{md}(\bar{\mathbf{c}}))^2} \leq \gamma^2 \frac{n}{\text{md}(\bar{\mathbf{c}})^2}$$

where in the first inequality we discarded the non-negative term $\sum_{j=1}^k (c_j)^2$, in the second inequality we

have used $|\delta| \leq 2\gamma^2 n / \text{md}(\bar{\mathbf{c}})$, and in the third one we simply assumed that $\text{md}(\bar{\mathbf{c}})$ is a sufficiently large constant, namely $\text{md}(\bar{\mathbf{c}}) \geq 8\gamma^2$.

On the other hand, we have the lower bound

$$(C.26) \quad \mathbf{E}[\Delta' | \mathbf{c}] = \frac{1}{n} \left(2\delta^2 - \sum_{j=1}^k (c_j)^2 \right) \geq \\ \geq -\frac{1}{n} \sum_{j=1}^k (c_j)^2 \geq -\frac{k}{n} \left(\frac{n-q}{k} \right)^2 \geq \\ \geq -\frac{4}{9} \cdot \frac{n}{k} \geq -\frac{4}{9} \cdot \frac{n}{\text{md}(\bar{\mathbf{c}})}$$

From the first to the second line we used the fact that all c_j 's are smaller than $n-q$. Then we used the fact that q is close to $n/2$, so $n-q$ is smaller than, say, $(2/3)n$. Finally we used the fact that $k \geq \text{md}(\bar{\mathbf{c}})$.

Hence, from (C.25) and (C.26) we get

$$-\frac{4}{9} \frac{n}{\text{md}(\bar{\mathbf{c}})} \leq \mathbf{E}[\Delta' | \mathbf{c}] \leq \gamma^2 \frac{n}{\text{md}(\bar{\mathbf{c}})}$$

Since $\Delta' = Q' - n/2$ can be written as a sum of n independent random variables taking values $\pm 1/2$, from the appropriate version of Chernoff bound it thus follows that

$$(C.27) \quad \mathbf{P} \left(\Delta' \notin \left(-2\gamma^2 \frac{n}{\text{md}(\bar{\mathbf{c}})}, 2\gamma^2 \frac{n}{\text{md}(\bar{\mathbf{c}})} \right) \mid \mathbf{c} \right) \leq \\ \leq \exp \left(-\Omega \left(\frac{n}{\text{md}(\bar{\mathbf{c}})^2} \right) \right) \leq \exp \left(-\Omega \left(n^{1/2} \right) \right)$$

where in the last inequality we used again the fact that $\text{md}(\bar{\mathbf{c}}) \leq k \leq \varepsilon (n/\log n)^{1/4}$.

In order to formally complete the proof, let us now define event $\mathcal{E}_t = \mathcal{A}_t \wedge \mathcal{B}_t$, where \mathcal{A}_t and \mathcal{B}_t are the events

$$\mathcal{A}_t = \text{“} |\Delta^{(t)}| \leq 2\gamma^2 \frac{n}{\text{md}(\bar{\mathbf{c}})} \text{”}$$

$$\mathcal{B}_t = \text{“} C_1^{(t)} \leq \left(1 + \frac{2\gamma(1+\gamma)+1}{\text{md}(\bar{\mathbf{c}})} \right)^t \cdot \gamma \frac{n}{\text{md}(\bar{\mathbf{c}})} \text{”}$$

Observe that

$$\left(1 + \frac{2\gamma(1+\gamma)+1}{\text{md}(\bar{\mathbf{c}})} \right)^t \leq 2 \text{ for } t \leq \frac{1}{4\gamma(1+\gamma)} \cdot \text{md}(\bar{\mathbf{c}})$$

Hence, if we set $T = \left\lfloor \frac{1}{4\gamma(1+\gamma)} \text{md}(\bar{\mathbf{c}}) \right\rfloor$, from (C.24) and (C.27) it follows that, for every $j \in [\bar{t}, \bar{t}+T]$, we get

$$\mathbf{P} \left(\mathcal{E}_j \mid \cap_{i=1}^{j-1} \mathcal{E}_i \right) \geq (1 - n^{-c})$$

for a positive constant c that we can make arbitrarily large. Thus, starting from the given color configuration $\mathbf{c}^{(\bar{t})}$, the probability that after T rounds the plurality $C_1^{(\bar{t}+T)}$ is at most $2\gamma n / \text{md}(\bar{\mathbf{c}})$ is

$$\mathbf{P} \left(C_1^{(\bar{t}+T)} \leq 2\gamma \frac{n}{\text{md}(\bar{\mathbf{c}})} \mid \mathbf{c}^{(\bar{t})} \right) \geq \\ \geq \mathbf{P} \left(\bigcap_{j=\bar{t}}^{\bar{t}+T} \mathcal{E}_j \right) = \prod_{j=\bar{t}}^{\bar{t}+T} \mathbf{P} \left(\mathcal{E}_j \mid \bigcap_{i=\bar{t}}^{j-1} \mathcal{E}_i \right) \\ \geq (1 - n^{-c})^T \geq 1 - Tn^{-c} \geq 1 - n^{-\Omega(1)}$$

□

THEOREM C.1. *Let $\bar{\mathbf{c}}$ be the initial color configuration. If the initial number of colors is $k \leq \varepsilon \cdot (n/\log n)^{1/6}$, where $\varepsilon > 0$ is a sufficiently small positive constant, then the convergence time of the Undecided-State Dynamics is $\Omega(\text{md}(\bar{\mathbf{c}}))$ w.h.p.*

Proof. From Lemma C.3 and Lemma C.6 it follows that there is a round \bar{t} , within the first $O(\log n)$ rounds, such that the process lies in a color configuration $\mathbf{c}^{(\bar{t})}$ where the number of undecided nodes is $|Q^{(\bar{t})} - n/2| \leq 2\gamma^2 / \text{md}(\bar{\mathbf{c}})$ and the plurality is $C_1^{(\bar{t})} \leq \gamma(n/\text{md}(\bar{\mathbf{c}}))$ w.h.p., where γ is a sufficiently large constant. From Lemma C.7, it then follows that the plurality C_1 remains smaller than $2\gamma(n/\text{md}(\bar{\mathbf{c}}))$ for the next $\Omega(\text{md}(\bar{\mathbf{c}}))$ rounds. □

There is, however, a “positive” drift for the plurality working in this “long” phase as well: this minimal drift (see the next lemma) allows the process to reach a state (representing the end of this phase) by which the plurality can re-start to grow fast (this phase-convergence time is formalized in Lemma C.9).

LEMMA C.8. (MINIMAL DRIFT) *Let $k = o\left(\sqrt{\frac{n}{\log n}}\right)$ and let $\epsilon \in (0, \frac{1}{2})$ be an arbitrarily small positive constant. Given a color configuration \mathbf{c} such that*

$$\begin{cases} c_1 \geq \beta \cdot \frac{n}{R(\mathbf{c})} & \text{for some constant } \beta > 0 \\ c_1 \geq (1 + \alpha) c_i & \text{for some constant } \alpha > 0 \\ & \text{and any } i \neq 1 \end{cases}$$

w.h.p. it holds either

$$R(\mathbf{C}') \leq 1 + \frac{\epsilon}{3} \text{ and } Q' \leq \epsilon n$$

or

$$\frac{C_1' + 2Q'}{n} \geq 1 + \Omega\left(\frac{1}{R(\mathbf{c})}\right)$$

Proof. First, let us derive a lower bound on $C'_1 + 2Q'$ from (C.28) we get that holds w.h.p.

By Lemma C.1

$$\mathbf{E}[C'_1 + 2Q' | \mathbf{c}] = n \cdot (1 + \Gamma(\mathbf{c}))$$

where

$$\Gamma(\mathbf{c}) = \left(1 - \frac{c_1 + 2q}{n}\right)^2 + 2(1 - \gamma)(R(\mathbf{c}) - 1) \left(\frac{c_1}{n}\right)^2$$

with $\gamma = (1 + \alpha)^{-1}$. As in the proof of Lemma C.5, observe that $\mathbf{E}[C'_1 + 2Q' | \mathbf{c}]$ can be written as the expected value of the sum of the following independent r.v.s: given \bar{c} , for each node i

$$X_i = \begin{cases} 1 & \text{if node } i \text{ is 1-colored at round } t + 1, \\ 2 & \text{if node } i \text{ is undecided at round } t + 1. \end{cases}$$

Thus, we can apply the Chernoff bound (A.1) to them and get that w.h.p.

$$(C.28) \quad C'_1 + 2Q' \geq n \cdot (1 + \Gamma(\mathbf{c})) \left(1 - O\left(\sqrt{\frac{\log n}{n}}\right)\right)$$

Let us analyze (C.28) when $R(\mathbf{c}) > 1 + \frac{\epsilon}{4}$ or $Q' > \frac{3}{4}\epsilon n$. If $R(\mathbf{c}) > 1 + \frac{\epsilon}{4}$ we have that

$$(C.29) \quad \begin{aligned} \Gamma(\mathbf{c}) &\geq 2(1 - \gamma)(R(\mathbf{c}) - 1) \left(\frac{c_1}{n}\right)^2 \\ &\geq 2(1 - \gamma) \left(1 - \frac{1}{R(\mathbf{c})}\right) R(\mathbf{c}) \cdot \left(\frac{\beta}{R(\mathbf{c})}\right)^2 \\ &> \frac{\alpha\epsilon\beta^2}{2(1 + \alpha)(1 + \epsilon/4)} \frac{1}{R(\mathbf{c})} \end{aligned}$$

On the other hand, if $R(\mathbf{c}) \leq 1 + \frac{\epsilon}{4}$ then

$$c_1 = \frac{n - q}{R(\mathbf{c})} \geq \frac{n - q}{1 + \epsilon/4} \geq (n - q)(1 - \epsilon/4) \geq n - q - \frac{\epsilon}{4}n$$

hence, if it also holds that $q > \frac{3}{4}\epsilon n$, the latter inequality implies that

$$1 - \frac{c_1 + 2q}{n} \leq \frac{\epsilon}{4} - \frac{q}{n} \leq -\frac{\epsilon}{2}$$

that is

$$(C.30) \quad \Gamma(\mathbf{c}) \geq \left(1 - \frac{c_1 + 2q}{n}\right)^2 \geq \frac{\epsilon^2}{4}$$

Therefore, if $R(\mathbf{c}) > 1 + \frac{\epsilon}{4}$ or $q > \frac{3}{4}\epsilon n$, then using (C.29), (C.30) and the given upper bound on the value of $R(\mathbf{c})$,

$$\begin{aligned} \frac{C'_1 + 2Q'}{n} &\geq (1 + \Gamma(\mathbf{c})) \left(1 - O\left(\sqrt{\frac{\log n}{n}}\right)\right) \\ &\geq \left(1 + \frac{\sigma}{R(\mathbf{c})}\right) \left(1 - O\left(\sqrt{\frac{\log n}{n}}\right)\right) \\ &\geq \left(1 + \frac{\sigma}{2R(\mathbf{c})}\right) \end{aligned}$$

where

$$\sigma = \min \left\{ \frac{\epsilon^2}{4} R(\mathbf{c}), \frac{\alpha\epsilon\beta^2}{2(1 + \alpha)(1 + \epsilon/4)} \right\}$$

It remains to show that if $R(\mathbf{c}) \leq 1 + \frac{\epsilon}{4}$ and $q \leq \frac{3}{4}\epsilon n$ then w.h.p. $R(\mathbf{C}') \leq 1 + \frac{\epsilon}{3}$ and $Q' \leq \epsilon n$.

In order to do so, observe that

$$\sum_{i \neq 1} c_i = (R(\mathbf{c}) - 1)c_1 \leq \frac{\epsilon}{4}n$$

It follows that

$$\begin{aligned} \mu_q &= \frac{q^2 + \sum_{i \neq j} c_i c_j}{n} \\ &\leq \frac{q^2 + 2c_1 \sum_{j \neq 1} c_j + \sum_{i \neq 1} c_i \sum_{j \neq 1} c_j}{n} \\ &\leq \left(\frac{3}{4}\epsilon\right)^2 n + \frac{\epsilon}{2}c_1 + \frac{\epsilon^2}{16}n \end{aligned}$$

Thanks to the Chernoff bound (A.2) and since $\epsilon < \frac{1}{2}$, the previous inequality implies that w.h.p. $Q' \leq \epsilon n$. As for $R(\mathbf{C}')$, by applying Lemma C.2 and using the Chernoff bound (A.2), we get that w.h.p. $R(\mathbf{C}') \leq 1 + \frac{\epsilon}{3}$, concluding the proof. \square

LEMMA C.9. *Let $k = O((n/\log n)^{1/4})$ and let $\epsilon > 0$ be an arbitrarily small constant. If the process is in a color configuration $\mathbf{c}^{(\bar{t})}$ that satisfies the following conditions:*

$$(C.31) \quad \left\{ \begin{aligned} \frac{c_1^{(\bar{t})} + 2q^{(\bar{t})}}{n} &= 1 + \Omega\left(\frac{1}{R(\mathbf{c}^{(\bar{t})})}\right) \end{aligned} \right.$$

$$(C.32) \quad \left\{ \begin{aligned} c_1^{(\bar{t})} &\geq \frac{1}{17} \frac{n}{R(\mathbf{c}^{(\bar{t})})} \end{aligned} \right.$$

$$(C.33) \quad \left\{ \begin{aligned} R(\mathbf{c}^{(\bar{t})}) &= O(\text{md}(\bar{\mathbf{c}})) \end{aligned} \right.$$

$$(C.34) \quad \left\{ \begin{aligned} c_1^{(\bar{t})} &\geq (1 + \alpha) \cdot c_i^{(\bar{t})} \text{ for some constant } \alpha > 0 \\ &\text{and for any color } i \neq 1 \end{aligned} \right.$$

then, after $T = O(\text{md}(\bar{\mathbf{c}}) \cdot \log n)$ rounds, the process is

w.h.p. in a color configuration $\mathbf{C}^{(\tilde{t}+T)}$ such that □

$$\begin{cases} C_1^{(\tilde{t}+T)} \geq \frac{1}{17} \frac{n}{R(\mathbf{C}^{(\tilde{t}+T)})} \\ R(\mathbf{C}^{(\tilde{t}+T)}) \leq 1 + \frac{\epsilon}{3} \\ Q^{(\tilde{t}+T)} \leq \epsilon n \\ C_1^{(\tilde{t}+T)} \geq (1 + \alpha) \cdot C_i^{(\tilde{t}+T)} (1 - o(1)) \end{cases} \text{ for any color } i \neq 1$$

Proof. First, we show that, if we start in a color configuration \mathbf{c} satisfying properties (C.31), (C.32), (C.33) and (C.34), then w.h.p. \mathbf{C}' still satisfies the conditions (C.32), (C.33) and (C.34).

Using the Chernoff bound (A.1) and conditions (C.32) and (C.31), we get that w.h.p.

$$\begin{aligned} C_1' &\geq \frac{c_1^{(\tilde{t})} + 2q^{(\tilde{t})}}{n} c_1 \left(1 - O\left(\sqrt{\frac{\log n}{\mu_1}}\right) \right) \\ &= \left(1 + \Omega\left(\frac{1}{R(\mathbf{c})}\right) \right) c_1 \geq \frac{1}{17} \frac{n}{R(\mathbf{c})} \end{aligned}$$

In the first equality, we used that (C.31) and (C.32) together imply that $\mu_1 \geq c_1 \geq \frac{1}{17} \frac{n}{R(\mathbf{c})} \gg \frac{1}{R(\mathbf{c})}$ w.h.p., thus proving that \mathbf{C}' also satisfies Condition (C.32) w.h.p. Moreover, Condition (C.32) allows us to apply Lemma C.2 to get that w.h.p.

$$\begin{aligned} C_1' &\geq (1 + \alpha) \cdot C_i' \cdot \left(1 - O\left((\log n / \mu_1)^{1/2}\right) \right) \\ R(\mathbf{C}') &< R(\mathbf{c}) \cdot \left(1 + O\left((\log n / \mu_1)^{1/2}\right) \right) \end{aligned}$$

proving that w.h.p. \mathbf{C}' satisfies the hypotheses (C.33) and (C.34).

Now, by Lemma C.8 and (C.33), it follows that w.h.p. either $R(\mathbf{C}') \leq 1 + \frac{\epsilon}{3}$ and $Q' \leq \epsilon n$ (in which case, we have done), or it holds w.h.p. that

$$\frac{C_1' + 2Q'}{n} = 1 + \Omega\left(\frac{1}{R(\mathbf{c})}\right) = 1 + \Omega\left(\frac{1}{\text{md}(\bar{\mathbf{c}})}\right)$$

In the latter case, \mathbf{C}' satisfies also Condition (C.31) and the above argument can be iterated again. In particular, (C.31) implies that after $T = \Omega(\text{md}(\bar{\mathbf{c}}) \log n)$ further rounds w.h.p. we have

$$\begin{aligned} C_1^{(\tilde{t}+T)} &= \left(1 + \Omega\left(\frac{1}{\text{md}(\bar{\mathbf{c}})}\right) \right) C_1^{(\tilde{t}+T-1)} = \dots = \\ &= \left(1 + \Omega\left(\frac{1}{\text{md}(\bar{\mathbf{c}})}\right) \right)^T C_1^{(\tilde{t})} = n - o(n) \end{aligned}$$

and thus

$$R(\mathbf{C}^{(\tilde{t}+T)}) - 1 = \frac{\sum_{i \neq 1} C_i^{(\tilde{t}+T)}}{C_1^{(\tilde{t}+T)}} \leq \frac{\epsilon}{3} \text{ and } Q^{(\tilde{t}+T)} \leq \epsilon n$$

C.5 Third phase: From plurality to totality.

The next theorem connects the results achieved in the previous sections into a consistent picture, establishing an upper bound on the overall convergence time of the process. Its proof also highlights the main features of the final phase, during which plurality turns into totality of the agents at an exponential rate.

THEOREM C.2. *Let $k = O((n/\log n)^{1/3})$ and let \bar{c} be any initial configuration such that for any $i \neq 1$ $c_1 \geq (1 + \alpha) \cdot c_i$ holds, where α is an arbitrarily small positive constant. Then, w.h.p. after at most $T = O(\text{md} \cdot \log n)$ time steps all agents support the initial plurality color.*

Proof. Let $\epsilon > 0$ be an arbitrarily small positive constant. Thanks to Lemma C.5, we can assume that at some time $\tilde{t} = O(\log n)$ the process w.h.p. reaches a configuration $\mathbf{C}^{(\tilde{t})}$ where

$$\begin{cases} \frac{C_1^{(\tilde{t})} + 2Q^{(\tilde{t})}}{n} = 1 + \Omega\left(\frac{1}{R(\mathbf{c}^{(\tilde{t})})}\right) \\ C_1^{(\tilde{t})} \geq \frac{1}{17} \frac{n}{R(\mathbf{c}^{(\tilde{t})})} \\ R(\mathbf{c}^{(\tilde{t})}) = O(\text{md}) \\ C_1^{(\tilde{t})} \geq (1 + \alpha) \cdot c_i^{(\tilde{t})} (1 - o(1)) \text{ for any color } i \neq 1 \end{cases}$$

Assuming $\mathbf{c}^{(\tilde{t})}$, Lemma C.9 determines the kick-off condition for a new phase in which both the undecided and the non-plurality color communities decrease exponentially fast. In particular, it implies that w.h.p., within $O(\text{md} \log n)$ further rounds, the process reaches a configuration $\mathbf{C}^{(t_{\text{end}})}$ such that the following properties hold:

$$(C.35) \quad C_1^{(t_{\text{end}})} \geq \frac{1}{17} \frac{n}{R(\mathbf{c}^{(t_{\text{end}})})}$$

$$(C.36) \quad \begin{cases} C_1^{(t_{\text{end}})} \geq (1 + \alpha) \cdot C_i^{(t_{\text{end}})} (1 - o(1)) \\ \text{for any color } i \neq 1 \end{cases}$$

$$(C.37) \quad R(\mathbf{c}^{(t_{\text{end}})}) \leq 1 + \frac{\epsilon}{3}$$

$$(C.38) \quad Q^{t_{\text{end}}} \leq \epsilon n$$

Now, we show that starting from any configuration satisfying the conditions above, any community (including the undecided) other than the plurality decreases exponentially fast until disappearance. To this aim, let $\psi = \sum_{i \neq 1} c_i + q$ and, as usual, let Ψ' be the r.v. associated to the value of ψ at the next time step. We prove

that the following holds in any round following t_{end} : i) w.h.p., both Q and $\sum_{i \neq 1} C_i$ are bounded by quantities that decrease by a constant factor, so that at any time following t_{end} , Ψ is (upper) bounded by a quantity that decreases exponentially fast, thus $C_1 = n - \Psi$ is (lower) bounded by an increasing quantity; ii) properties (C.36), still holds. In the rest of this proof we assume $\epsilon < 1/3$, which is consistent with the assumptions of Lemma C.9.

To begin with, note that Property (C.37) implies $\sum_{i \neq 1} c_i \leq \frac{\epsilon}{3}n$, so that

$$\sum_{i \neq j} c_i \cdot c_j \leq 2c_1 \sum_{j \neq 1} c_i + \sum_{i \neq 1} c_i \sum_{j \neq 1} c_j \leq \left(\frac{2}{3}\epsilon + \frac{\epsilon^2}{9} \right) n^2$$

Therefore, properties (C.37) and (C.38) together imply

$$(C.39) \quad \mu_q = \frac{(q)^2 + \sum_{i \neq j} c_i \cdot c_j}{n} \leq \left(\epsilon^2 + \frac{2}{3}\epsilon + \frac{\epsilon^2}{9} \right) n < \frac{3}{4}\epsilon n$$

$$(C.40) \quad \mathbf{E} \left[\sum_{i \neq 1} C'_i \mid \mathbf{c} \right] = \sum_{i \neq 1} \left(c_i \frac{c_i + 2q}{n} \right) \leq \frac{1}{3} \left(\frac{1}{3} + 2 \right) \epsilon^2 n = \frac{7}{9} \epsilon^2 n < \frac{7}{27} \epsilon n$$

where we use the assumption that $\epsilon < 1/3$. At this point, we can use the Chernoff bound (A.2) to show that (C.39) and (C.40) hold w.h.p. (up to a multiplicative factor $1 + o(1)$). This proves that w.h.p., both Q and $\sum_{i \neq 1} C_i$ (and hence Ψ) decrease by a constant factor in a round⁹. It remains to observe that, when q and/or $\sum_{i \neq 1} c_i$ become $O(\log n)$, an application of the Chernoff bound (A.3) shows that w.h.p., they remain below this value in the subsequent rounds. This completes the proof of i). Moreover, since $C'_1 = n - \Psi'$, i) implies that C'_1 is lower bounded by an increasing quantity w.h.p. Additionally, property (C.35) and i) just proved, together with property (C.36), imply the assumptions of Lemma C.2, allowing us to show that w.h.p. property (C.36) still holds at the end of next round as well.

As a consequence, we have that in at most $\tau = O(\log n)$ rounds w.h.p. we reach a color configuration $\tilde{C}^{(t_{end} + \tau)}$ such that $Q^{(t_{end} + \tau)} + \sum_{i \neq 1} C_i^{(t_{end} + \tau)} = O(\log n)$.

Finally, we can apply Markov's inequality on the value of $\sum_{i \neq 1} C_i^{(t_{end} + \tau)}$ to show that at the next round w.h.p. all color communities except for the plurality one disappear. \square

⁹In fact, a more careful analysis, unnecessary to prove our result, could use (C.40) to show that $\sum_{i \neq 1} C_i$ decreases superexponentially fast.

D Node congestion analysis

The parallel random walks described in Section 4 yield variable token queues in the nodes. For each node $u \in [n]$, and for every round $t \in [2\tau]$ of the phase, we consider the r.v. $\mathcal{Q}_{t,u}$ defined as the number of tokens in u at round t of any phase of the modified dynamics. In the next lemma we prove a useful bound on the maximal congestion in a phase of length 2τ .

LEMMA D.1. *Consider a phase of length $2\tau \geq 1$ of the above protocol on a d -regular graph $G = (V, E)$. Let $u \in V$ be any node and let t be any round of the phase. Then, for any constant $c > 0$, it holds that*

$$\mathbf{P} \left(\max_{1 \leq t \leq 2\tau} \mathcal{Q}_{t,u} \leq \max \left\{ \sqrt{2c\tau \log n}, 3c \log n \right\} \right) \geq 1 - \frac{(2\tau)^2}{n^{c/3}}$$

Proof. Consider the number Y_t of tokens received by a fixed node u at round t (for brevity's sake, we will omit index u in any r.v.). Then we can write

$$Y_t = \sum_{i \in [d]} X_{i,t}$$

where $X_{i,t} = 1$ if the i -th neighbor of u sends a token to u and 0 o.w.. Observe (again) that the r.v.s $X_{i,t}$ are not mutually independent. However, the crucial fact is that, for any t and any i , it holds $\mathbf{P}(X_{i,t} = 1) \leq 1/d$, regardless the state of the system (in particular, independently of the value of the other r.v.s).

So, if we consider a family $\{\hat{X}_{i,t} : i \in [d] \ t \in [2\tau]\}$ of i.i.d. Bernoulli r.v.s with $\mathbf{P}(\hat{X}_{i,t} = 1) = 1/d$, then Y_t is stochastically smaller than

$$\hat{Y}_t = \sum_{i=1}^d \hat{X}_{i,t}$$

For any node u and any round t , the r.v. \mathcal{Q}_t is thus stochastically smaller than the r.v. $\hat{\mathcal{Q}}_t$ defined recursively as follows.

$$\begin{cases} \hat{\mathcal{Q}}_t &= \hat{\mathcal{Q}}_{t-1} + \hat{Y}_t - \chi_t \\ \hat{\mathcal{Q}}_0 &= 1 \end{cases}$$

$$\text{where } \chi_t = \begin{cases} 1 & \text{if } \hat{\mathcal{Q}}_{t-1} > 0 \\ 0 & \text{otherwise} \end{cases}$$

Since our goal is to provide a concentration upper bound on \mathcal{Q}_t , we can do it by considering the ‘‘simpler’’ process $\hat{\mathcal{Q}}_t$. By the way, unrolling $\hat{\mathcal{Q}}_t$ directly is far from trivial: We need the ‘‘right’’ way to write it by using only i.i.d. Bernoulli r.v.s. Let's see how.

For any $t \in [2\tau]$ and for any $s \in [t]$, define the r.v.

$$(D.41) \quad Z_{s,t} = \sum_{i=s}^t \hat{Y}_i - (t-s)$$

Informally speaking, $Z_{s,t}$ matches the value of \hat{Q}_t whenever $s \leq t$ was the last previous round s.t. $\hat{Q}_s = 0$.

As a key-fact we show that \hat{Q}_t can be bounded by the maximum of $Z_{s,t}$ for $s \leq t$.

CLAIM 2. For any $t \in [2\tau]$ it holds that

$$\hat{Q}_t \leq \max\{Z_{s,t} : s = 1, \dots, t\}$$

and thus

$$(D.42) \quad \max\{Q_t : 1 \leq t \leq 2\tau\} \leq \max\{Z_{s,t} : 1 \leq s \leq t \leq 2\tau\}$$

Proof. (of the Claim). For any $s \in [t]$, let

$$\chi_{s,t} = \prod_{r=s}^t \chi_r$$

be the r.v. taking value 1 if $\hat{Q}_{r-1} > 0$ for all $s \leq r \leq t$ and 0 otherwise. It is easy to prove by induction that \hat{Q}_t can be written as

$$(D.43) \quad \hat{Q}_t = \sum_{s=2}^t (1 - \chi_{s-1}) \chi_{s,t} Z_{s-1,t} + \chi_{1,t} Z_{1,t} + (1 - \chi_t) Z_{t,t}$$

Since

$$\sum_{s=2}^t (1 - \chi_{s-1}) \chi_{s,t} + \chi_{1,t} = 1$$

the sum in (D.43) is not larger than the maximum of the $Z_{s,t}$, hence

$$\hat{Q}_t \leq \max\{Z_{s,t} : s = 1, \dots, t\}$$

and

$$\max\{Q_t : 1 \leq t \leq 2\tau\} \leq \max\{Z_{s,t} : 1 \leq s \leq t \leq 2\tau\}$$

□(of the Claim).

Let us consider (D.41): The r.v. $Z_{s,t} + (t - s)$ is a sum of $d \cdot (t - s + 1)$ i.i.d. Bernoulli r.v.s each one with expectation $1/d$. From the Chernoff bounds (A.2) and (A.3), for any $1 \leq s \leq t$, it holds that

$$\mathbf{P}\left(Z_{s,t} \leq \max\left\{\sqrt{c(t-s+1)\log n}, 6c\log n\right\}\right) \geq 1 - n^{-c/3}$$

By taking the union bound over all $1 \leq s \leq t \leq 2\tau$, from the above bound and (D.42) we can get the desired concentration bound on the maximal node congestion during every phase:

$$\mathbf{P}\left(\max_{1 \leq t \leq 2\tau} Q_t \leq \max\left\{\sqrt{2c\tau \log n}, 6c\log n\right\}\right) \geq 1 - \frac{(2\tau)^2}{n^{c/3}}$$

Let $t_{\text{mix}}^G(\epsilon)$ be the first round such that the total variation distance between the simple random walk starting at an arbitrary node and the uniform distribution is smaller than ϵ , i.e.,

$$t_{\text{mix}}^G(\epsilon) = \inf\{t \in \mathbb{N} : \|P^t(u, \cdot) - \pi\| \leq \epsilon \text{ for all } u \in V\}$$

Notice that for any $\epsilon > 0$ it holds that (see e.g. (4.36) in [23])

$$(D.44) \quad t_{\text{mix}}^G(\epsilon) \leq \log(1/\epsilon) t_{\text{mix}}^G(1/(2e))$$

As a consequence of the above Lemma, we can now set the right value of τ , thus getting the following result.

THEOREM D.1. Let $G = ([n], E)$ be a d -regular graph with $t_{\text{mix}}^G(1/4) = \text{polylog}(n)$. Each round of the Undecided-State Dynamics on the clique can be simulated on G in the GOSSIP model in $\text{polylog}(n)$ rounds by exchanging messages of $\text{polylog}(n)$ size, w.h.p.

Proof. Let $2\tau = \alpha \bar{t}^2 \log n$ be the length of the phase, where $\bar{t} = t_{\text{mix}}^G(1/n^2)$ and α is a suitable constant that we fix later. From Lemma D.1, we have that the maximum number of tokens in every node at any round of the phase is w.h.p at most

$$\sqrt{2c\tau \log n} = \sqrt{\alpha c} \cdot \bar{t} \log n$$

Since tokens are enqueued with a FIFO policy, each single hop of the random walk performed by a token can be delayed for at most the above number of rounds. Hence, in order to perform \bar{t} hops of the random walk, a token takes at most $\sqrt{\alpha c} \cdot \bar{t}^2 \log n$ rounds w.h.p.

By choosing $\alpha \geq 4c$ we have that this number is smaller than τ , this allows us to set τ so that the forward process and the backward one can both complete safely.

By union bounding over all tokens we thus have that during the phase all tokens perform at least \bar{t} hops of a random walk and report back to the sender the color of the node they reached after \bar{t} hops w.h.p.

Finally, notice that from (D.44) it follows that $\bar{t} = \text{polylog}(n)$. The phase length and the size of the exchanged messages are thus $\text{polylog}(n)$ as well. □

Since a lazy random walk on regular expanders (see e.g. [21]) has $\text{polylog}(n)$ mixing time, from the above theorem and our result on the Undecided-State Dynamics on the clique we easily get the following final result.

COROLLARY 1. From any initial configuration \bar{c} such that the Undecided-State Dynamics on the clique completes plurality consensus in $O(\text{md}(\bar{c}) \log n)$ rounds w.h.p., the modified Undecided-State Dynamics completes plurality consensus on any d -regular expander graph within $O(\text{md}(\bar{c}) \cdot \text{polylog}(n))$ rounds w.h.p.

□