

A Tight Analysis of the Parallel Undecided-State Dynamics with Two Colors

Andrea Clementi
Università di Roma “Tor Vergata”
 Roma, Italy
 clementi@mat.uniroma2.it

Mohsen Ghaffari
ETH Zurich
 Zurich, Switzerland
 ghaffari@mit.edu

Luciano Gualà
Università di Roma “Tor Vergata”
 Roma, Italy
 guala@mat.uniroma2.it

Emanuele Natale
Max-Planck-Institut für Informatik
 Saarbrücken, Germany
 enatale@mpi-inf.mpg.de

Francesco Pasquale
Università di Roma “Tor Vergata”
 Roma, Italy
 pasquale@mat.uniroma2.it

Giacomo Scornavacca
Università dell’Aquila
 L’Aquila, Italy
 scornavacca@graduate.univaq.it

Abstract—The *Undecided-State Dynamics* is a well-known protocol that achieves *Consensus* in distributed systems formed by a set of n anonymous nodes interacting via a communication network. We consider this dynamics in the parallel *PULL* communication model on the complete graph for the binary case, i.e., when every node can either support one of two possible colors (say, **Alpha** and **Beta**) or stay in the *undecided* state. Previous work in this setting only considers initial color configurations with no undecided nodes and a large *bias* (i.e., $\Theta(n)$) towards the majority color.

An interesting open question here is whether this dynamics reaches consensus *quickly*, i.e. within a polylogarithmic number of rounds.

In this paper we present an *unconditional* analysis of the Undecided-State Dynamics which answers to the above question in the affirmative. Our analysis shows that, starting from *any* initial configuration, the Undecided-State Dynamics reaches a monochromatic configuration within $\mathcal{O}(\log n)$ rounds, with high probability (w.h.p.). This bound is tight since a lower bound $\Omega(\log n)$ for this process is known. Moreover, we prove that if the initial configuration has bias $\Omega(\sqrt{n \log n})$, then the dynamics converges toward the initial majority color within $\mathcal{O}(\log n)$ round, w.h.p.

I. INTRODUCTION

Strong research interest has been recently focussed on the study of simple, local mechanisms for *Consensus* problems in distributed systems [3], [2], [17], [18], [23], [24]. In one of the basic versions of the consensus problem, the system consists of a finite set of n anonymous entities (nodes) that run elementary operations and interact by exchanging messages. Every node initially supports a value (i.e. a *color*) chosen from a finite alphabet Σ and a *Consensus Protocol* is a local procedure that, starting from any color configuration, let the system converge in finite time to a monochromatic configuration where every node supports the same color. The consensus is *valid* if the *winning* color is a *valid* one: It is one among those initially supported by at least one node. Moreover, the consensus configurations must result *equilibria* of the protocol process: Once the system reaches a consensus configuration, it will stay there forever unless some external event takes place.

We study the consensus problem in the *PULL* model [12], [16], [21] in which, at every round, each active node of a communication network contacts one neighbor uniformly at random to pull information. A well-studied natural consensus protocol in this model is the *Undecided-State Dynamics*¹ (for short, the *U-Dynamics*) in which the state of a node can be either a color or the *undecided* state. When a node is activated, it pulls the state of a random neighbors and updates its state according to the following updating rule (see Table I): If a colored node pulls a different color from its current one, then it becomes undecided, while in all other cases it keeps its color; moreover, if the node is in the undecided state then it will take the state of the pulled neighbor.

The U-Dynamics has been studied in both *sequential* and *parallel* models: Informally, in the former, at every round, only one random node is activated, while in the latter, at every round, all nodes are activated synchronously.

As for the sequential model², [3] provides an unconditional analysis showing (among other results) that the U-Dynamics solves the *binary* consensus problem (i.e. when $|\Sigma| = 2$) in the complete graph within $\mathcal{O}(n \log n)$ activations (and, thus in $\mathcal{O}(\log n)$ “parallel” time), *with high probability*³.

As for the parallel *PULL* model, even though it is easy to verify that the U-Dynamics achieves consensus in the complete graph (with high probability), the convergence time of this dynamics is still an interesting open issue, even in the binary case. We remark that the stochastic process yielded by the parallel dynamics significantly departs from the process yielded by the sequential one. To get just one immediate evidence of this difference, observe that, in the former model, the system can converge to the (non-valid) configuration where

¹In some previous papers [24] on the binary case ($|\Sigma| = 2$), this protocol has been also called the *Third-State Dynamics*. We here prefer the term “undecided” since it also holds for the non-binary case and, moreover, the term well captures the role of this additional state.

²[3] in fact considers the *Population-Protocol* model which is, in our specific context, equivalent to the sequential *PULL* model.

³As usual, we say that an event \mathcal{E}_n holds *w.h.p.* if $\mathbf{P}(\mathcal{E}_n) \geq 1 - n^{-\Theta(1)}$.

all nodes are undecided even if starting from a “fully-colored” configuration (where all nodes are not undecided). On the other hand, it is easy to see that this evolution cannot happen in the sequential setting. A deeper, crucial difference lies in the random number of nodes that may change color at every round: In the sequential model, this is at most one⁴, while in the parallel one, *all* nodes may change state in one shot and indeed, for most phases of the process, the expected number of changes is linear in n . It thus turns out that the probabilistic arguments used in the analysis of [3] appear not useful in the parallel setting. In [5], the author analyzes the U-Dynamics in the parallel \mathcal{PULL} model on the complete graph when the alphabet Σ has size k , where $k = o(n^{1/3})$. The analysis in [5] considers this dynamics as a protocol for *Plurality Consensus* [2], [3], [22], a variant of Consensus, where the goal is to reach consensus on the color that was initially supported by the *plurality* of the nodes: Their analysis requires that the initial configuration must have a relatively-large *bias* $s = c_1 - c_2$ between the size c_1 of the (unique) initial plurality and the size c_2 of the second-largest color. More in details, in [5] it is assumed that $c_1 \geq \alpha c_2$, for some absolute constant $\alpha > 1$ and, thus, this condition for the binary case would result into requiring a very-large initial bias, i.e., $s = \Theta(n)$. This analysis clearly does not show that the U-Dynamics efficiently solves the binary consensus problem, mainly because it does not manage *balanced* initial configurations.

Our results: We prove that, starting from any color configuration⁵ on the complete graph, the U-Dynamics reaches a monochromatic configuration (thus consensus) within $\mathcal{O}(\log n)$ rounds, with high probability. This bound is tight since, for some (in fact, a large number of) initial configurations, the process requires $\Omega(\log n)$ rounds to converge.

Not assuming a large initial bias of the majority color significantly complicates the analysis. Indeed, the major technical issues arise from the analysis of *balanced* initial configurations where the system “needs” to *break symmetry* without having a strong expected drift towards any color. Essentially, previous analysis of this phase consider either *sequential* processes of interacting particles that can be modeled as *birth-and-death* chains [3] or parallel processes whose local rule is fully symmetric w.r.t. the states/colors of the nodes (such as majority rules)[6], [17]. The U-Dynamics process falls neither in the former nor in the latter scenario: it works in parallel rounds and the role of the undecided nodes makes the local rule not symmetric. Informally speaking, in Section IV, we show an “efficient” way to reduce all “critical” almost-balanced starting of the process to a specific regime along which the system keeps a number q of undecided nodes which is some suitable constant fraction of n until the bias s has reached an $\Omega(\sqrt{n \log n})$ magnitude: In other words, during this regime, with very high probability the system never jumps to almost-balanced configurations having either too many or too few undecided nodes. This fact is crucial essentially because of

two reasons: along this regime, (i) the *variance* of the bias s is large (i.e. $\Theta(n)$) and (ii) whenever the bias s gets $\Omega(\sqrt{n})$, its drift turns out to be *exponential* with non-negligible, increasing probability (w.r.t. s itself). Then, using a suitable coupling to a “simplified” pruning process, we can apply (a suitable version) of a general lemma [17] (see Claim 9.2 in [17]) that provides a logarithmic bound on the hitting time of some Markov chains that have Properties (i) and (ii) above.

The symmetry-breaking phase terminates when the U-Process reaches some configuration having a bias $s = \Omega(\sqrt{n \log n})$. Then (see Section V) we prove that, starting from *any* configuration having that bias, the process reaches consensus within $\mathcal{O}(\log n)$ rounds, with high probability. Even though our analysis of this “majority” part of the process is based on standard concentration arguments, it must cope with some *non-monotone* behaviour of the key random variables (such as the bias and the number of undecided nodes at the next round): Again, this is due to the non-symmetric role played by the undecided nodes. A good intuition about this “non-monotone” process can be gained by looking at the mutually-related formulas giving the expectation of such key random variables (see Equations (1)-(3)). Our refined analysis shows that, during this majority phase, the winning color never changes and, thus, the U-Dynamics also ensures *Plurality Consensus* in logarithmic time whenever the initial bias is $s = \Omega(\sqrt{n \log n})$.

Interestingly enough, we also show that configurations with $s = \mathcal{O}(\sqrt{n})$ exist so that the system may converge toward the minority color with non-negligible probability.

Further motivation and related work: On the U-Dynamics. The interest in the U-Dynamics arises in fields beyond the borders of Computer Science and it seems to have a key-role in important biological processes modelled as so-called chemical reaction networks [11], [18]. For such reasons, the convergence time of this dynamics has been analyzed on different communication models [1], [3], [4], [8], [14], [17], [19], [22], [24].

As previously mentioned, the U-Dynamics has been analysed in the parallel \mathcal{PULL} model in [5] and their results concern the evolution of the process for the multi-color case when there is a significant initial bias (as a protocol for plurality consensus). As for the sequential model, the U-Dynamics has been introduced and analyzed in [3] in the complete graph. They prove that this dynamics, with high probability, converges to a valid consensus within $\mathcal{O}(n \log n)$ activations and, moreover, it converges to the majority whenever the initial bias is $\omega(\sqrt{n \log n})$.

Still concerning the sequential model, [22] recently analyzes, besides other protocols, the U-Dynamics in arbitrary graphs when the initial configuration is sampled uniformly at random between the two colors. In this (average-case) setting, they prove that the system converges to the initial majority color with higher probability than the initial minority one. They also give results for special classes of graphs where the minority can win with large probability if the initial configuration is chosen in a suitable way. Their proof for this last result relies

⁴This number becomes 2 if the sequential communication model activate a random edge per round, rather than one single node [3].

⁵Our analysis also considers initial configurations with undecided nodes.

on an exponentially-small upper bound on the probability that a certain minority can win in the complete graph (see [22] for more details).

In [4], [8], [19], [24], the same dynamics for the binary case has been analyzed in further sequential communication models.

On some other consensus dynamics. Recently, further simple consensus protocols have been deeply analyzed in several papers, thus witnessing the high interest of the scientific community on such processes [3], [7], [10], [11], [14], [15], [17], [24].

The parallel 3-MAJORITY is a protocol where at every round, each node picks the colors of three random neighbors and updates its color according to the majority rule (taking the first one or a random one to break ties). All theoretical results for 3-MAJORITY consider the complete graph. The authors of [7] assume that the bias is $\Omega(\min\{\sqrt{2k}, (n/\log n)^{1/6}\} \cdot \sqrt{n \log n})$. Under this assumption, they prove that consensus is reached with high probability in $\mathcal{O}(\min\{k, (n/\log n)^{1/3}\} \cdot \log n)$ rounds, and that this is tight if $k \leq (n/\log n)^{1/4}$. The first result without bias [6] restricts the number of initial colors to $k = \mathcal{O}(n^{1/3})$. Under this assumption, they prove that 3-MAJORITY reaches consensus with high probability in $\mathcal{O}((k^2(\log n)^{1/2} + k \log n) \cdot (k + \log n))$ rounds. Very recently, such result has been generalized to the whole range of k in [9].

In [17] the authors provide an analysis of the 3-median rule, in which every node updates its value to the median of its random sample. They show that this dynamics converges to an almost-agreement configuration (which is even a good approximation of the global median) within $\mathcal{O}(\log k \cdot \log \log n + \log n)$ rounds, w.h.p. It turns out that, in the binary case, the median rule is equivalent to the 2-CHOICES dynamics, a variant of 3-MAJORITY, thus their result implies that this is a stabilizing consensus protocol with $\mathcal{O}(\log n)$ convergence time. As mentioned earlier, our analysis borrows a hitting-time bound on general Markov chains from [17].

In [14], [15], the authors consider the 2-CHOICES dynamics for plurality consensus in the binary case (i.e. $k = 2$). For random d -regular graphs, [14] proves that all nodes agree on the majority color in $\mathcal{O}(\log n)$ rounds, provided that the bias is $\omega(n \cdot \sqrt{1/d + d/n})$. The same holds for arbitrary d -regular graphs if the bias is $\Omega(\lambda_2 \cdot n)$, where λ_2 is the second largest eigenvalue of the transition matrix. In [15], these results are extended to general expander graphs.

II. PRELIMINARIES

We analyze the parallel version of the dynamics called U-Dynamics in the (uniform) \mathcal{PULL} model on a complete graph: Starting from an initial configuration where every node supports a color, i.e. a value from a set Σ of possible colors, at every round, each node u pulls the color of a randomly-selected neighbor v . If the color of node v differs from its own color, then node u enters in an *undecided* state (an extra state with no color). When an node is in the undecided state and pulls a color, it gets that color. Finally, an node that pulls either an

undecided node or an node with its own color remains in its current state.

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

TABLE I
THE UPDATE RULE OF THE U-DYNAMICS WHERE $i, j \in [k]$ AND $i \neq j$.

In this paper we consider the case in which there are two possible colors (say color Alpha and color Beta). Let us name \mathcal{C} the space of all possible configurations and observe that, since we are on the complete graph, a configuration $\mathbf{x} \in \mathcal{C}$ is completely determined by the number of nodes with color Alpha and the number of nodes with color Beta , say $a(\mathbf{x})$ and $b(\mathbf{x})$, respectively.

It is convenient to give names also to two other quantities that will appear often in the analysis: the number $q(\mathbf{x}) = n - a(\mathbf{x}) - b(\mathbf{x})$ of undecided nodes and the difference $s(\mathbf{x}) = a(\mathbf{x}) - b(\mathbf{x})$ between the numbers of Alpha -colored and Beta -colored nodes. We will call $s(\mathbf{x})$ the *bias* of configuration \mathbf{x} . Notice that any two of the quantities $a(\mathbf{x})$, $b(\mathbf{x})$, $q(\mathbf{x})$, and $s(\mathbf{x})$ uniquely determine the configuration. When it will be clear from the context, we will omit \mathbf{x} and write a, b, q , and s instead of $a(\mathbf{x}), b(\mathbf{x}), q(\mathbf{x})$, and $s(\mathbf{x})$.

Observe that the U-Dynamics defines a finite-state Markov chain $\{\mathbf{X}_t\}_{t \geq 0}$ with state space \mathcal{C} and three absorbing states, namely, $q = n$, $a = n$, and $b = n$. We call *U-Process* the random process obtained by applying the U-Dynamics starting at a given state. Once we fix the configuration \mathbf{x} at round t of the process, i.e. $\mathbf{X}_t = \mathbf{x}$, we use the capital letters A, B, Q , and S to refer to the random variables $a(\mathbf{X}_{t+1}), b(\mathbf{X}_{t+1}), q(\mathbf{X}_{t+1}), s(\mathbf{X}_{t+1})$.

From the definition of U-Dynamics it is easy to compute the following expected values (see also Section 3 in [5])

$$\mathbf{E}[A | \mathbf{X}_t = \mathbf{x}] = a \left(\frac{a + 2q}{n} \right) \quad (1)$$

$$\mathbf{E}[Q | \mathbf{X}_t = \mathbf{x}] = \frac{q^2 + 2ab}{n} \quad (2)$$

$$\mathbf{E}[S | \mathbf{X}_t = \mathbf{x}] = \frac{a(a + 2q)}{n} - \frac{b(b + 2q)}{n} = s \left(1 + \frac{q}{n} \right) \quad (3)$$

A. The expected evolution of the U-Dynamics

Equations (1)-(3) can be used to have a preliminary intuitive idea on the expected evolution of the U-Dynamics. From (3) it follows that the bias s increases exponentially, in expectation, as long as the number q of undecided nodes is a constant fraction of n (say, $q \geq \delta n$, for some positive constant δ). By rewriting (2) in terms of q and s we have that

$$\begin{aligned} \mathbf{E}[Q | \mathbf{X}_t = \mathbf{x}] &= \frac{q^2 + 2ab}{n} = \frac{2q^2 + (n - q)^2 - s^2}{2n} \\ &\geq \frac{n}{3} - \frac{s^2}{2n} \end{aligned} \quad (4)$$

where in the inequality we used the fact that the minimum of $2q^2 + (n - q)^2$ is achieved at $q = n/3$ and its value is $2n^2/3$. From (4) it thus follows that, as long as the magnitude of the bias is smaller than a constant fraction of n (say $|s| < 2n/3$), the expected number of undecided nodes will be larger than a constant fraction of n at the next round (say, $\mathbf{E}[Q | \mathbf{X}_t = \mathbf{x}] \geq n/9$).

When the magnitude of the bias $|s|$ reaches $2n/3$, it is easy to see that the expected number of nodes with the *minority* color decreases exponentially. Indeed, suppose wlog that B is the minority color and let us rewrite (1) for B and in terms of b and s , we get

$$\mathbf{E}[B | \mathbf{X}_t = \mathbf{x}] = b \left(\frac{b + 2q}{n} \right) = b \left(1 - \frac{2s + 3b - n}{n} \right). \quad (5)$$

Hence, when $s > 2n/3$ we have that $\mathbf{E}[B | \mathbf{X}_t = \mathbf{x}] \leq (1 - 2/3)b$.

The above sketch of the analysis *in expectation* would suggest that the process should end up in a monochromatic configuration within $\mathcal{O}(\log n)$ rounds. Indeed, in Theorem 2 we prove that this is what happens with high probability (w.h.p., from now on) when the process starts from a configuration that already has some bias, namely $s = \Omega(\sqrt{n \log n})$.

When the process starts from a configuration with a smaller bias, the analysis *in expectation* loses its predictive power. As an extremal example, observe that when $a = b = n/3$ the system is “in equilibrium” according to (1)-(3). However, the equilibrium is “unstable” and the symmetry is broken by the *variance* of the process (as long as $s = o(\sqrt{n})$) and by the increasing drift towards majority (as soon as $s > \sqrt{n}$). As mentioned in the Introduction, the analysis of this *symmetry-breaking* phase is the key technical contribution of the paper and it will be described in Section IV. This analysis will show that, starting from any initial configuration, the system reaches a configuration where the magnitude of the bias is $\Omega(\sqrt{n \log n})$ within $\mathcal{O}(\log n)$ rounds, w.h.p.

III. MAIN RESULTS AND THE DIGRAPH OF THE U-PROCESS' PHASES

As informally discussed in the introduction, we prove the two following results characterizing the evolution of the U-Dynamics on the synchronous *PULL* model in the complete graph.

Theorem 1 (Consensus). *Let the U-Process start from any configuration in \mathcal{C} . Then the process converges to a (valid) monochromatic configuration within $\mathcal{O}(\log n)$ rounds, w.h.p. Furthermore, if the initial configuration has at least one colored node (i.e. $q \leq n - 1$), then the process converges to a configuration such that $|s| = n$, w.h.p.*

Theorem 2 (Plurality consensus). *Let γ be any positive constant and assume that the U-Process starts from any biased configuration such that $|s| \geq \gamma\sqrt{n \log n}$ and assume w.l.o.g. the majority color is Alpha. Then the process converges to the monochromatic configuration with $a = n$ within $\mathcal{O}(\log n)$ rounds, w.h.p. Furthermore, the result is almost tight in a*

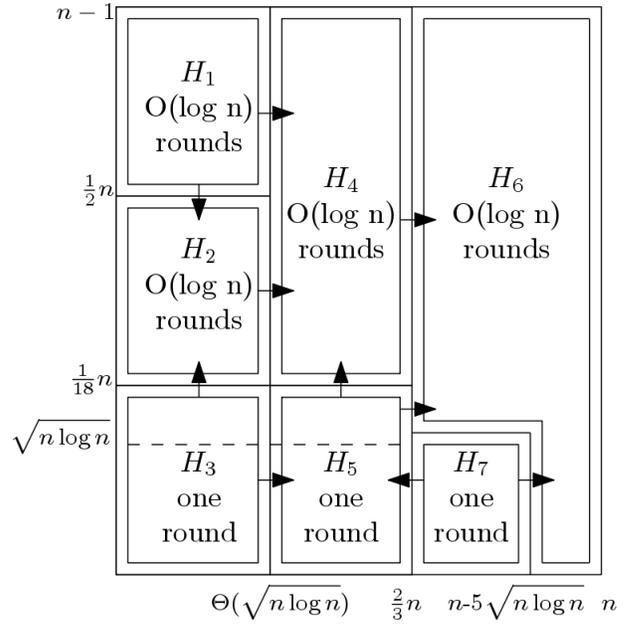


Fig. 1. $\{H_1, \dots, H_7\}$ is the considered partitioning of the configuration space \mathcal{C} . On the x axis we represent the bias s , on the y axis the number of undecided nodes q . Missing arrows are transitions that have negligible probabilities.

twofold sense: (i) An initial configuration exists, with $|s| = \Omega(\sqrt{n \log n})$, such that the process requires $\Omega(\log n)$ rounds to converge w.h.p., and (ii) there is an initial configuration with $|s| = \Theta(\sqrt{n})$ such that the process converges to the minority color with constant probability.

Outline of the two proofs. The two theorems above are consequences of our refined analysis of the evolution of the U-Process. The analysis is organized into a set of possible process phases, each of them is defined by specific ranges of parameters q and s . A high-level description of this structure is shown in Fig. 1 where every rectangular region represents a subset of configurations with specific ranges of s and q and it is associated to a specific phase. In details, let γ be any positive constant, then the regions are defined as follows: H_1 is the set of configurations such that $s \leq \gamma\sqrt{n \log n}$ and $q \geq \frac{1}{2}n$; H_2 is the set of configurations such that $s \leq \gamma\sqrt{n \log n}$ and $\frac{1}{18}n \leq q \leq \frac{1}{2}n$; H_3 is the set of configurations such that $s \leq \gamma\sqrt{n \log n}$ and $q \leq \frac{1}{18}n$. H_4 is the set of configurations such that $\gamma\sqrt{n \log n} \leq s \leq \frac{n}{6}$ and $q \geq \frac{1}{6}n$; H_5 is the set of configurations such that $\gamma\sqrt{n \log n} \leq s \leq \frac{2}{3}n$ and $q \leq \frac{1}{6}n$; H_7 is the set of configurations such that $\frac{2}{3}n \leq s \leq n - 5\sqrt{n \log n}$ and $q \leq \sqrt{n \log n}$. H_6 is the set of configurations such that $s \geq \frac{2}{3}n$ minus H_7 .

For each region, Fig. 1 specifies our upper bound on the exit time from the corresponding phase, while black arrows represent all possible phase transitions which may happen with non-negligible probability.

As a first, important remark, we point out that the scheme of Fig. 1 can be seen as a directed acyclic graph G with

a single sink H_6 , which is reachable from any other region. We also remark that, starting from certain configurations, the monochromatic state may be reached via different paths in G . This departs from previous analysis of consensus processes [5], [7], [17] in which the phase transition graph is essentially a path.

We now outline the proofs of the two main results of this paper.

Outline of the Proof of Theorem 2. Consider an initial configuration \mathbf{x} such that $s(\mathbf{x}) \geq \gamma\sqrt{n \log n}$, for some positive constant γ , and assume w.l.o.g. that the majority color in \mathbf{x} is Alpha. In Section V, we first show (see Lemma 6) that if the process lies in H_4 the bias grows exponentially fast and thus the process enters in H_6 within $\mathcal{O}(\log n)$ rounds. Then we prove Lemma 7 stating that, starting from any configuration in H_6 , the process ends in the monochromatic configuration where $a = n$ in $\mathcal{O}(\log n)$ rounds. Next, we show that, starting from any configuration in H_5 , the process falls into H_4 or in H_6 in one round (Lemma 8) and that, starting from any configuration in H_7 the process falls into H_4, H_5 or H_6 in one round (Lemma 9). Concerning the tightness of the result stated in the second part of the theorem, we have that the lower bound on the convergence time is an immediate consequence of Claim (ii) of Lemma 6. While, Claim (ii), concerning the lower bound on the initial bias, will be proved in Claim 13 which is provided in Appendix P.

Outline of the Proof of Theorem 1. We first observe that the configuration where all nodes are undecided (i.e. $q = n$) is an absorbing state of the U-Process and thus, for this initial configuration, Theorem 1 trivially holds. In Section IV, we will show that, starting from any *balanced* configuration, i.e. with $|s| = o(\sqrt{n \log n})$, the U-Process “breaks symmetry” reaching a configuration \mathbf{y} with $|s(\mathbf{y})| = \Omega(\sqrt{n \log n})$ within $\mathcal{O}(\log n)$ rounds, w.h.p. Then, the thesis easily follows by applying Theorem 2 with initial configuration \mathbf{y} . As for the symmetry-breaking phase, in Lemma 3 we prove that, if the process starts from a configuration in H_1 or H_3 (see Figure 1), then after $\mathcal{O}(\log n)$ rounds either the bias between the two colors becomes $\Omega(\sqrt{n \log n})$ or the system reaches some configuration in H_2 , w.h.p. In Lemma 5 we then prove that, if the process is in a configuration in H_2 , then the bias between the two colors will become $\Omega(\sqrt{n \log n})$ within $\mathcal{O}(\log n)$ rounds, w.h.p.

IV. SYMMETRY BREAKING

In this section we show that, starting from any (almost-) balanced configuration, i.e. those with $|s| = o(\sqrt{n \log n})$, the U-Process “breaks symmetry” reaching a configuration with $|s| = \Omega(\sqrt{n \log n})$ within $\mathcal{O}(\log n)$ rounds, w.h.p. This part of our analysis is organized as follows.

In Lemma 3 we prove that, if the process starts in a configuration in H_1 or H_3 (see Figure 1), i.e., when the number of undecided nodes is either smaller than $n/18$ or larger than $n/2$, then, after $\mathcal{O}(\log n)$ rounds, either the bias between the two colors already gets magnitude $\Omega(\sqrt{n \log n})$ or the system reaches some configuration in H_2 (i.e., a configuration where

the number of undecided nodes is between $n/18$ and $n/2$). In Lemma 5 we then prove that, if the process is in a configuration in H_2 , then the bias between the two colors will get magnitude $\Omega(\sqrt{n \log n})$ within $\mathcal{O}(\log n)$ rounds, w.h.p.

Lemma 3 is a simple consequence of the following three claims. Claims 1 and 2 follow from Chernoff bound applied to (4) and (1), respectively (their proofs can be found in Appendix C).

Claim 1. *Let $\mathbf{x} \in \mathcal{C}$ be any configuration with $s(\mathbf{x}) \leq (2/3)n$. Then, at the next round, the number of undecided nodes of the U-Process is $Q \geq n/18$, w.h.p.*

Claim 2. *Let $\mathbf{x} \in \mathcal{C}$ be any configuration with $q(\mathbf{x}) \geq n/2$ and $a(\mathbf{x}) \geq \log n$. Then, at the next round, the number of Alpha-colored nodes of the U-Process is $A \geq (1 + 3/4)a(\mathbf{x})$, w.h.p.*

The next claim is a consequence of fact that, when the number of colored nodes is very small, the U-Process behaves essentially like a *pull* process.

Claim 3. *Starting from any configuration $\mathbf{x} \in \mathcal{C}$ with $1 \leq a(\mathbf{x}) + b(\mathbf{x}) < 2 \log n$, the U-Process reaches a configuration \mathbf{X}' with $a(\mathbf{X}') + b(\mathbf{X}') \geq 2 \log n$ within $\mathcal{O}(\log n)$ rounds, w.h.p.*

Sketch of Proof. Let us consider the random variable counting the number of *colored nodes* $A + B$ and its evolution during the process. As long as $1 \leq a(\mathbf{x}) + b(\mathbf{x}) < 2 \log n$, the probability that in one round an Alpha-colored node picks a Beta-colored node (or vice versa) is less than $\frac{(2 \log n)^2}{n}$. Applying the union bound for $\mathcal{O}(\log n)$ rounds, we get that the probability that this “bad” event happens in one of such rounds is negligible.

Now, assuming no such bad events happen, the colored nodes will remain colored. Moreover, we know that an undecided node becomes colored if it picks a colored node. So, discarding the difference between colors, the process over the undecided nodes turns out to be a standard *rumor-spreading* process via *PULL* messages (a colored node is in fact an informed node). The claim then follows by observing that this spreading process is known to inform at least $2 \log n$ nodes within $\mathcal{O}(\log n)$ rounds, w.h.p. (see for instance [20]). ■

Lemma 3 (Phases H_1 and H_3 : Starters). *Starting from any configuration $\mathbf{x} \in H_1$, the U-Process reaches a configuration $\mathbf{X}' \in (H_2 \cup H_4)$ within $\mathcal{O}(\log n)$ rounds, w.h.p.*

Starting from any configuration $\mathbf{x} \in H_3$, the U-Process reaches a configuration $\mathbf{X}' \in (H_1 \cup H_2 \cup H_4)$ in one round, w.h.p.

If the system lies in a configuration of H_2 , we need more complex probabilistic arguments to prove that the bias between the two colors reaches $\Omega(\sqrt{n \log n})$ within $\mathcal{O}(\log n)$ rounds w.h.p.

We will make use of a useful bound on some hitting time of a Markov chain having suitable drift properties. This result can be obtained via a simple adaptation of Claim 2.9 in [17] while a self-contained proof is given in Appendix G.

Lemma 4. Let $\{X_t\}_t$ be a Markov Chain with finite state space Ω and let $f : \Omega \mapsto [0, n]$ be a function that maps states to integer values. Let $m = \mathcal{O}(\sqrt{n \log n})$ be a target value and let h, c_1, c_2, ε be four positive constants with $h \geq 3, (3c_2 - 1)h > 2$, and $\varepsilon h > 2$. If the following properties hold

1) For any $x \in \Omega$ such that $f(x) < h\sqrt{n}$

$$\mathbf{P}(f(X_{t+1}) < h\sqrt{n} | X_t = x) < c_1 < 1,$$

2) For any $x \in \Omega$ such that $h\sqrt{n} \leq f(x) < m$

$$\mathbf{P}(f(X_{t+1}) < (1 + \varepsilon)f(X_t) | X_t = x) < e^{-c_2 f(x)^2/n},$$

then the process reaches a state x with $f(x) \geq m$ within $\mathcal{O}(\log n)$ rounds, w.h.p.

The basic idea would be to apply the above lemma to the U-Process with $f(X_t) = |s(X_t)|$ in order to get an upper bound on the number of rounds needed to reach a configuration such that the bias is $\Omega(\sqrt{n \log n})$. To this aim we first show that for any configuration in H_2 Properties 1 and 2 in the above lemma are satisfied.

Claim 4. If $\frac{n}{18} \leq q \leq \frac{n}{2}$ then four positive constants h, c_1, c_2, ε exist with $h \geq 3, (3c_2 - 1)h > 2$, and $\varepsilon h > 2$ such that:

1) If $s < h\sqrt{n}$ then $\mathbf{P}(S < h\sqrt{n}) < c_1$;

2) If $s \geq h\sqrt{n}$ then $\mathbf{P}(S \geq (1 + \varepsilon)s) \geq 1 - e^{-c_2 s^2/n}$.

Sketch of Proof. From the additive form of the Chernoff bound (see Appendix B) it follows that

$$\mathbf{P}\left(A < \mathbf{E}[A] - \frac{1}{72}s\right) < e^{-2s^2/72^2n},$$

$$\mathbf{P}\left(B > \mathbf{E}[B] + \frac{1}{72}s\right) < e^{-2s^2/72^2n}.$$

Thus:

$$\begin{aligned} S &= A - B > \mathbf{E}[A] - \frac{1}{72}s - \mathbf{E}[B] - \frac{1}{72}s \\ &= \mathbf{E}[A - B] - \frac{1}{36}s = \mathbf{E}[S] - \frac{1}{18}s \\ &= \left(1 + \frac{q}{n}\right)s - \frac{1}{36}s = \left(1 + \frac{1}{18} - \frac{1}{36}\right)s \\ &= \left(1 + \frac{1}{36}\right)s. \end{aligned}$$

Hence, the second item is obtained setting $\varepsilon = \frac{1}{36}$ and $c_2 = \frac{2}{72^2}$. As for the first item, we observe that the bias s is a difference of two binomial distributions. So we can choose a large enough constant h to ensure the parameter conditions in Lemma 4, i.e., $h \geq 3, (3c_2 - 1)h > 2, \varepsilon h > 2$, while keeping c_1 a constant less than one. The formal proof is a simple adaptation of Claim 13 in Appendix P. ■

We remark that the above claim ensures Properties 1 and 2 of Lemma 4 whenever $\frac{1}{18}n \leq q \leq \frac{1}{2}n$. Unfortunately, Lemma 4 requires such properties to hold for *any* (almost-)balanced configuration: if $q = n - o(n)$ Property 1 does not hold, while Property 2 is not satisfied if $q = o(n)$. In order to manage this issue, in Subsection IV-A, we define a *pruned* process, a

variant of U-Process where it is possible to apply Lemma 4. Then, in Subsection IV-B we show a coupling between the U-Process and the pruned one.

A. The pruned process

The helpful, key point is that, starting from any configuration $\mathbf{x} \in H_2$, the probability that the process goes in one of those “bad” configurations with $q < \frac{1}{18}n$ or $q \geq \frac{1}{2}n$ is negligible (see Claim 5). Thus, intuitively speaking, all the configurations *actually visited* by the process before exiting H_2 do satisfy Lemma 4. In order to make this intuitive argument rigorous, in what follows, we first define a suitably *pruned* process by removing from H_2 all the *unwanted* transitions.

Let $\bar{s} \in [n]$ and $\mathbf{z}(\bar{s})$ the configuration such that $s(\mathbf{z}(\bar{s})) = \bar{s}$ and $q(\mathbf{z}(\bar{s})) = \frac{1}{2}n$. Let $p_{\mathbf{x}, \mathbf{y}}$ the probability of a transition from the configuration \mathbf{x} to the configuration \mathbf{y} in the U-Process. We define a new stochastic process: the U-Pruned-Process. The U-Pruned-Process behaves exactly like the original process but every transition from a configuration $\mathbf{x} \in H_2$ to a configuration \mathbf{y} such that $q(\mathbf{y}) < \frac{1}{18}n$ or $q(\mathbf{y}) > \frac{1}{2}n$ now have probability $p'_{\mathbf{x}, \mathbf{y}} = 0$. Moreover, for any $\bar{s} \in [n]$, starting from any configuration $\mathbf{x} \in H_2$ the probability of reaching the configuration $\mathbf{z}(\bar{s})$ is:

$$p'_{\mathbf{x}, \mathbf{z}(\bar{s})} = p_{\mathbf{x}, \mathbf{z}(\bar{s})} + \sum_{\mathbf{y}: (q(\mathbf{y}) < \frac{1}{18}n \vee q(\mathbf{y}) > \frac{1}{2}n) \wedge s(\mathbf{y}) = \bar{s}} p_{\mathbf{x}, \mathbf{y}}.$$

Finally, all the other transition probabilities remain the same.

Observe that, since the U-Pruned-Process is defined in such a way it has exactly the same marginal probability of the original process with respect to the random variable s , then Claim 4 holds for the U-Pruned-Process as well. Thus, we can choose constants h, c_1, c_2, ε such that the two properties of Lemma 4 are satisfied. Then we get the following:

Corollary 1. Starting from any configuration $\mathbf{x} \in H_2$, the U-Pruned-Process reaches a configuration $\mathbf{x}' \in H_4$ within $\mathcal{O}(\log n)$ rounds, w.h.p.

B. Back to the original process.

The definition of the U-Pruned-Process suggests an obvious coupling between the original process and the pruned one: if the two process are in different state they act independently, if they are in the same configuration \mathbf{x} they move together unless the U-Process go in a configuration \mathbf{y} such that $q(\mathbf{y}) < \frac{1}{18}n$ or $q(\mathbf{y}) > \frac{1}{2}n$. In that case the U-Pruned-Process goes in $\mathbf{z}(s(\mathbf{x}))$.

Using this simple coupling, we first show that, if the two processes are in the same configuration, the probability that they get separated is negligible. Then, we show that the H_2 exit time of the pruned procedure stochastically dominate the H_2 exit time of the original process.

Claim 5. For every configuration $\mathbf{x} \in H_2$, the probability that the number of undecided nodes in the next round of the U-Process is not between $n/18$ and $n/2$ is

$$\mathbf{P}\left(q(X_{t+1}) \notin \left[\frac{n}{18}, \frac{n}{2}\right] \mid X_t = \mathbf{x}\right) \leq e^{-\Theta(n)}.$$

Lemma 5 (Phase H_2). *Starting from any configuration $\mathbf{x} \in H_2$, the U-Process reaches a configuration $\mathbf{X}' \in H_4$ within $\mathcal{O}(\log n)$ rounds, w.h.p.*

Proof: Let $\{\mathbf{X}_t\}$ and $\{\mathbf{Y}_t\}$ be the original process and the pruned one, respectively. Let $\mathbf{x} \in H_2$, note that if $\mathbf{X}_t = \mathbf{Y}_t = \mathbf{x}$ then

$$\mathbf{Y}_{t+1} = \begin{cases} \mathbf{X}_{t+1} & \text{if } \mathbf{X}_{t+1} \in H_2 \\ \mathbf{z}(s(\mathbf{X}_{t+1})) & \text{otherwise} \end{cases}$$

Let $\tau = \inf\{t \in \mathbb{N} : |s(\mathbf{X}_t)| \geq \sqrt{n \log n}\}$ and let $\tau^* = \inf\{t \in \mathbb{N} : |s(\mathbf{Y}_t)| \geq \sqrt{n \log n}\}$. For any $\mathbf{x} \in H_2$ and any round t we define ρ_x^t the event $\{\mathbf{X}_t\}$ and $\{\mathbf{Y}_t\}$ separated at round $t+1$, i.e. $\rho_x^t = (\mathbf{X}_t = \mathbf{Y}_t = \mathbf{x}) \wedge (\mathbf{X}_{t+1} \neq \mathbf{Y}_{t+1})$. Observe that, if the two coupled processes start in the same configuration $\mathbf{x}_0 \in H_2$ and $\tau > c \log n$, then either $\tau^* > c \log n$ as well, or a round $t \leq c \log n$ exists such that, for some $\mathbf{x} \in H_2$ the event ρ_x^t occurred. Hence,

$$\begin{aligned} & \mathbf{P}_{\mathbf{x}_0, \mathbf{x}_0}(\tau > c \log n) \leq \\ & \leq \mathbf{P}_{\mathbf{x}_0, \mathbf{x}_0} \left(\{\tau^* > c \log n\} \cup \left\{ \begin{array}{l} \exists t \leq c \log n \\ \exists \mathbf{x} \in H_2 \end{array} : \rho_x^t \right\} \right) \\ & \leq \mathbf{P}_{\mathbf{x}_0, \mathbf{x}_0}(\tau^* > c \log n) + \mathbf{P}_{\mathbf{x}_0, \mathbf{x}_0} \left(\begin{array}{l} \exists t \leq c \log n \\ \exists \mathbf{x} \in H_2 \end{array} : \rho_x^t \right). \quad (6) \end{aligned}$$

As for the first term in (6), from the analysis of the pruned process (Corollary 1) we have that it is upper bounded by $1/n$. As for the second term, we have that

$$\begin{aligned} \mathbf{P}_{\mathbf{x}_0, \mathbf{x}_0} \left(\begin{array}{l} \exists t \leq c \log n \\ \exists \mathbf{x} \in H_2 \end{array} : \rho_x^t \right) & \leq \sum_{t=1}^{c \log n} \mathbf{P}_{\mathbf{x}_0, \mathbf{x}_0}(\exists \mathbf{x} \in H_2 : \rho_x^t) \\ & = \sum_{t=1}^{c \log n} \sum_{\mathbf{x} \in H_2} \mathbf{P}_{\mathbf{x}_0, \mathbf{x}_0}(\rho_x^t) \\ & \leq \sum_{t=1}^{c \log n} \frac{n^2}{e^{-\Theta(n)}} \\ & \leq \frac{1}{n}, \end{aligned} \quad (7)$$

where in 7 we used Claim 5 and the fact that $|H_2|$ is at most all the possible combinations of the parameters q and s . ■

V. CONVERGENCE TO THE MAJORITY

In this section we provide the arguments needed to prove our second main result, namely Theorem 2, which essentially states that starting from any sufficiently biased configuration, the U-Process converges to the monochromatic configuration where all nodes have the majority color. Remind that the outline of the proof is given in Section III. Here, we formalize the arguments of the provided high-level description. Due to space limitation the proofs of the technical claims are moved to Appendix.

Phase H_4 : The age of the undecideds: We first show that under some parameter ranges including H_4 (and hence when the number of the undecideds are large enough) the growth of the bias is exponential.

Claim 6. *Let γ be any positive constant and $\mathbf{x} \in \mathcal{C}$ be any configuration such that $s \geq \gamma \sqrt{n \log n}$ and $q \geq \frac{1}{18}n$. Then, it holds that $s(1 + \frac{1}{36}) < S < 2s$, w.h.p.*

Lemma 6 (Phase H_4). *Let $\mathbf{x} \in H_4$ be a configuration with $a > b$. Then, (i) starting from \mathbf{x} , the U-Process reaches a configuration $\mathbf{X}' \in H_6$ with $a > b$ within $\mathcal{O}(\log n)$ rounds, w.h.p. Moreover, (ii) an initial configuration $\mathbf{y} \in H_4$ exists such that the U-Process stays in H_4 for $\Omega(\log n)$ rounds, w.h.p.*

Proof: We iteratively apply Claim 6 and Claim 1 and after $t = \Theta(\log n)$ rounds we have that either there is a round $t' < t$ such that $s(\mathbf{X}_{t'}) > \frac{2}{3}n$, or $s(\mathbf{X}_t) > (1 + 1/36)^t s(\mathbf{x}) \geq (1 + 1/36)^t$. In both cases, the process has reached a configuration \mathbf{X}' such that $s(\mathbf{X}') \geq \frac{2}{3}n$ and $q(\mathbf{X}') \geq \frac{n}{18}$: So \mathbf{X}' belongs to H_6 . Since each step of the iteration holds w.h.p. and the number of steps is $\mathcal{O}(\log n)$, we easily obtain that the result holds w.h.p. by a simple application of the Union Bound.

Concerning the second part of the lemma, consider an initial configuration \mathbf{y} such that $s(\mathbf{y}) = n^{2/3}$. By iteratively applying (the upper bound of) Claim 6 and Claim 1 for $t = \frac{1}{4} \log n$ rounds, we have that $s(\mathbf{X}_t) < 2^t s(\mathbf{y}) = 2^t n^{2/3} = n^{1/4} n^{2/3} = o(n)$. ■

Phase H_6 : The victory of the majority: This is the phase in which a large bias let the nodes converge to the majority color within a logarithmic number of rounds. We first prove that the number of nodes that support the minority color decreases exponentially fast (Claim 7) and that the bias is preserved round by round (Claim 8 and Claim 9). Then, when $b \leq 2\sqrt{n \log n}$, the undecided nodes start to decrease exponentially fast as well (Claim 10). At the very end, when there are only few nodes (i.e., $\mathcal{O}(\sqrt{n \log n})$) that do not still support the majority color, the minority color disappears in few steps and thus the U-Process converges to majority within $\mathcal{O}(\log n)$ rounds (Claim 11).

Claim 7. *Let $\mathbf{x} \in \mathcal{C}$ be any configuration such that $s \geq \frac{2}{3}n$ and $b \geq \log n$ then it holds that $B \leq b(1 - \frac{1}{9})$, w.h.p.*

In order to iteratively apply the above claim we now show that, if there are enough undecided nodes, the bias is preserved round by round until the number of Beta-colored nodes decreases below $2\sqrt{n \log n}$.

Claim 8. *Let $\mathbf{x} \in \mathcal{C}$ be any configuration such that $s \geq \frac{2}{3}n$ and $q \geq \sqrt{n \log n}$. Then it holds that $S \geq \frac{2}{3}n$, w.h.p.*

Claim 9. *Let $\mathbf{x} \in \mathcal{C}$ be any configuration such that $s \geq \frac{2}{3}n$ and $b \geq 2\sqrt{n \log n}$. Then it holds that $Q > \sqrt{n \log n}$, w.h.p.*

The three above claims imply that, after $\mathcal{O}(\log n)$ rounds, the process reaches a configuration such that $s \geq \frac{2}{3}n$, $q \geq \sqrt{n \log n}$ and $b \leq 2\sqrt{n \log n}$. The next claim shows that starting from any such configuration the number of undecided nodes decrease exponentially fast. Next, we show that if the process reaches a configuration such that $q \leq 12\sqrt{n \log n}$ and $b \leq 2\sqrt{n \log n}$ then within few rounds the U-Process converges to the configuration where all nodes support Alpha.

Claim 10. Let $\mathbf{x} \in \mathcal{C}$ be any configuration such that $12\sqrt{n \log n} \leq q \leq \frac{1}{3}n$ and $b \leq 2\sqrt{n \log n}$ it holds that $Q \leq q(1 - \frac{1}{9})$, w.h.p.

Claim 11. Let γ be any positive constant and let $\mathbf{x} \in \mathcal{C}$ be any configuration such that $q \leq \gamma\sqrt{n \log n}$ and $b \leq 2\sqrt{n \log n}$ then the U-Process reaches a configuration \mathbf{X}' with $a(\mathbf{X}') = n$ within $\mathcal{O}(\log n)$ rounds, w.h.p.

We are now ready to show the following

Lemma 7 (Phase H_6). Starting from any configuration $\mathbf{x} \in H_6$ with $a > b$, the U-Process ends in the monochromatic configuration where $a = n$ within $\mathcal{O}(\log n)$ rounds, w.h.p.

Proof: Let us first assume that $s(\mathbf{x}) \geq n - 5\sqrt{n \log n}$ and $q(\mathbf{x}) \leq \sqrt{n \log n}$. This implies that $b(\mathbf{x}) \leq 2\sqrt{n \log n}$ and thanks to Claim 11 we get that the process end in the configuration such that $a = n$ within $\mathcal{O}(\log n)$ rounds. Otherwise $s(\mathbf{x}) \geq \frac{2}{3}n$ and $q \geq \sqrt{n \log n}$. Then, starting from \mathbf{x} , we iteratively apply Claim 7 together with Claim 8 and Claim 9, and we get that the process reaches a configuration \mathbf{X}' such that $s(\mathbf{X}') \geq \frac{2}{3}n$, $q(\mathbf{X}') \geq \sqrt{n \log n}$ and $b(\mathbf{X}') \leq 2\sqrt{n \log n}$ in $\mathcal{O}(\log n)$ rounds. Then we iteratively apply Claim 10 together with Claim 7 (if $b < \log n$ we cannot apply Claim 7 in order to show that B does not overtake $2\sqrt{n \log n}$ but we can get the claim with a simple application of the Markov inequality) and Claim 8 and we get that the process reaches a configuration \mathbf{X}'' such that $q(\mathbf{X}'') \leq 12\sqrt{n \log n}$ and $b(\mathbf{X}'') \leq 2\sqrt{n \log n}$ in $\mathcal{O}(\log n)$ rounds and now we apply Claim 11 and the process reaches the monochromatic configuration w.h.p. Since every step of the iterations holds w.h.p. and the number of steps is $\mathcal{O}(\log n)$, we easily obtain the thesis by a simple application of the Union Bound. ■

Phases H_5 and H_7 : Starters: We show that if the process is in a configuration where the number of the undecided nodes is relatively small with respect to the bias, then in the next round the number of the undecided nodes becomes large while the bias does not decrease too much, w.h.p. This essentially implies that if the process starts in H_5 then in the next round the process moves to a configuration belonging to H_5 or H_6 (Lemma 8), while if it starts in H_7 then in the next round it moves to H_4 or H_5 or H_6 (Lemma 9).

Claim 12. Let γ, ε be any two positive constants and $\mathbf{x} \in \mathcal{C}$ any configuration such that $s \geq \gamma\sqrt{n \log n}$ then it holds that $S \geq (\gamma - \varepsilon)\sqrt{n \log n}$, w.h.p.

The above claim together with Claim 1 immediately implies the following

Lemma 8 (Phase H_5). Starting from any configuration $\mathbf{x} \in H_5$ with $a > b$, the U-Process reaches a configuration $\mathbf{X}' \in (H_4 \cup H_6)$ with $a > b$ in one round, w.h.p.

Concerning phase H_7 , we have

Lemma 9 (Phase H_7). Starting from any configuration $\mathbf{x} \in H_7$ with $a > b$, the U-Process reaches a configuration $\mathbf{X}' \in (H_4 \cup H_5 \cup H_6)$ with $a > b$ in one round, w.h.p.

Proof: Note that Claim 12 implies that in the next round the process does not enter in H_1 , H_2 or H_3 w.h.p. The hypothesis that $s \leq n - 5\sqrt{n \log n}$ and $q \leq \sqrt{n \log n}$ implies that $b \geq 2\sqrt{n \log n}$ and thus we can apply the Claim 9 and get that the process leaves H_7 because of the grown of the undecided nodes. ■

VI. CONCLUSIONS

We provided a full analysis of the U-Dynamics in the parallel *PULL* model for the binary case showing that the resulting process converges quickly, regardless of the initial configuration. Besides giving tight bounds on the convergence time, our set of results well-clarifies the main aspects of the process evolution and the crucial role of the undecided nodes in each phase of this evolution.

An interesting open question is that of considering the same process in the multi-color case and to derive bounds on the time required to break symmetry from balanced configurations, as well.

Finally, we believe our analysis can be suitably adapted in order to show that the U-Dynamics can efficiently stabilize to a valid consensus regime even in presence of a dynamic adversary that can change the state of a subset of nodes of size $o(\sqrt{n})$. So, borrowing the notions from [3], [6], we should be able to show the U-Dynamics is a *stabilizing almost-consensus protocol* in presence of an $o(\sqrt{n})$ -dynamic adversary.

REFERENCES

- [1] M. A. Abdullah and M. Draief. Majority consensus on random graphs of a given degree sequence. *arXiv:1209.5025*, 2012.
- [2] M. A. Abdullah and M. Draief. Global majority consensus by local majority polling on graphs of a given degree sequence. *Discrete Applied Mathematics*, 180:1–10, 2015.
- [3] D. Angluin, J. Aspnes, and D. Eisenstat. A Simple Population Protocol for Fast Robust Approximate Majority. *Distributed Computing*, 21(2):87–102, 2008. (Preliminary version in DISC'07).
- [4] A. Babae 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, and R. Silvestri. Plurality consensus in the gossip model. In *ACM-SIAM SODA'15*, pages 371–390, 2015.
- [6] L. Becchetti, A. Clementi, E. Natale, F. Pasquale, and L. Trevisan. Stabilizing consensus with many opinions. In *ACM-SIAM SODA'16*, pages 620–635, 2016.
- [7] L. Becchetti, A. E. F. Clementi, E. Natale, F. Pasquale, R. Silvestri, and L. Trevisan. Simple dynamics for plurality consensus. In *ACM SPAA'14*, pages 247–256, 2014.
- [8] F. Bénézit, P. Thiran, and M. Vetterli. Interval consensus: from quantized gossip to voting. In *IEEE ICASSP'09*, 2009.
- [9] P. Berenbrink, A. Clementi, P. Kling, R. Elsässer, F. Mallmann-Trenn, and E. Natale. Ignore or comply? on breaking symmetry in consensus. In *ACM PODC'17*, 2017. to appear (Tech. Rep. in arXiv:1702.04921 [cs.DC]).
- [10] P. Berenbrink, G. Giakkoupis, A.-M. Kermerrec, and F. Mallmann-Trenn. Bounds on the voter model in dynamic networks. In *ICALP'16*, 2016.
- [11] L. Cardelli and A. Csikász-Nagy. The cell cycle switch computes approximate majority. *Scientific Reports*, Vol. 2, 2012.
- [12] 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 *ACM STOC'12*, 2012.
- [13] F. R. K. Chung and L. Lu. Survey: Concentration inequalities and martingale inequalities: A survey. *Internet Mathematics*, 3(1):79–127, 2006.
- [14] C. Cooper, R. Elsässer, and T. Radzik. The power of two choices in distributed voting. In *ICALP'14*, 2014.

- [15] C. Cooper, R. Elsässer, T. Radzik, N. Rivera, and T. Shiraga. Fast consensus for voting on general expander graphs. In *DISC'15*, 2015.
- [16] 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 *ACM PODC'87*, 1987.
- [17] B. Doerr, L. A. Goldberg, L. Minder, T. Sauerwald, and C. Scheideler. Stabilizing consensus with the power of two choices. In *ACM SPAA'11*, pages 149–158, 2011.
- [18] D. Doty. Timing in chemical reaction networks. In *ACM-SIAM SODA'14*, pages 772–784, 2014.
- [19] M. Draief and M. Vojnović. Convergence speed of binary interval consensus. *SIAM Journal on Control and Optimisation*, 50(3):1087–1109, 2012.
- [20] R. Karp, C. Schindelhauer, S. Shenker, and B. Vocking. Randomized rumor spreading. In *Proc. of 41st IEEE FOCS*, pages 565–574. IEEE, 2000.
- [21] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *IEEE FOCS'03*, 2003.
- [22] 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 *ICALP'14*, pages 871–882, 2014.
- [23] E. Mossel, J. Neeman, and O. Tamuz. Majority dynamics and aggregation of information in social networks. *Autonomous Agents and Multi-Agent Systems*, 28(3):408–429, 2014.
- [24] E. Perron, D. Vasudevan, and M. Vojnovic. Using Three States for Binary Consensus on Complete Graphs. In *INFOCOM'09*, pages 2527–1535, 2009.

APPENDIX

A. Chernoff Bound multiplicative form

Let X_1, \dots, X_n be independent 0-1 random variables. Let $X = \sum_{i=1}^n X_i$ and $\mu \leq \mathbf{E}[X] \leq \mu'$. Then, for any $0 < \delta < 1$ the following Chernoff bounds hold:

$$\mathbf{P}(X \geq (1 + \delta)\mu) \leq e^{-\mu\delta^2/3}. \quad (8)$$

$$\mathbf{P}(X \leq (1 - \delta)\mu) \leq e^{-\mu'\delta^2/2}. \quad (9)$$

B. Chernoff Bound additive form

Let X_1, \dots, X_n be independent 0-1 random variables. Let $X = \sum_{i=1}^n X_i$ and $\mu = \mathbf{E}[X]$. Then the following Chernoff bounds hold:

for any $0 < \lambda < n - \mu$

$$\mathbf{P}(X \leq \mu - \lambda) \leq e^{-2\lambda^2/n}, \quad (10)$$

for any $0 < \lambda < \mu$

$$\mathbf{P}(X \geq \mu + \lambda) \leq e^{-2\lambda^2/n}. \quad (11)$$

C. Reverse Chernoff Bound

Let X_1, \dots, X_n be independent 0-1 random variables, $X = \sum_{i=1}^n X_i$, $\mu = \mathbf{E}[X]$ and $\delta \in (0, 1/2]$. Assuming that $\mu \leq \frac{1}{2}n$ and $\delta^2\mu \geq 3$ then the following bounds hold:

$$\mathbf{P}(X \geq (1 + \delta)\mu) \geq e^{-9\delta^2\mu}, \quad (12)$$

$$\mathbf{P}(X \leq (1 - \delta)\mu) \geq e^{-9\delta^2\mu}. \quad (13)$$

D. Proof of Claim 1

From (4) we get

$$\mathbf{E}[Q|X_t = \mathbf{x}] \geq \frac{n}{3} - \frac{s^2}{2n} \geq \frac{1}{3}n - \frac{2}{9}n = \frac{1}{9}n.$$

By applying the additive form of the Chernoff Bound (see (10)) to the random variable Q we easily get the claim, i.e.,

$$\begin{aligned} \mathbf{P}\left(Q \leq \frac{1}{18}n\right) &= \mathbf{P}\left(Q \leq \frac{1}{18}n + \frac{1}{18}n - \frac{1}{18}n\right) \\ &\leq \mathbf{P}\left(Q \leq \mathbf{E}[Q] - \frac{1}{18}n\right) \\ &\leq e^{-2n^2/18^2n} \\ &= e^{-\Theta(n)}. \end{aligned}$$

E. Proof of Claim 2

We recall (1) and we get

$$\begin{aligned} \mathbf{P}(A|X_t = \mathbf{x}) &= a\left(\frac{a+2q}{n}\right) \\ &\geq a\left(\frac{\frac{n-q}{2} + 2q}{n}\right) \\ &= a\left(\frac{1}{2} + \frac{3q}{2n}\right) \geq a\left(\frac{1}{2} + 3\right) = a\frac{7}{2}, \end{aligned} \quad (14)$$

where in (14) we used that $a = \frac{n-q+s}{2} \geq \frac{n-q}{2}$. We apply the multiplicative form of the Chernoff Bound ((9) in Appendix A) with $\delta = \frac{1}{2}$

$$\begin{aligned} \mathbf{P}\left(A \leq a\frac{7}{2}\left(1 - \frac{1}{2}\right)\right) &\leq \exp^{-a\frac{7}{2}\frac{1}{4}/3} \\ &\leq \exp^{-\log n \frac{7}{2}\frac{1}{4}/3} \\ &= \frac{1}{n^{\Theta(1)}}. \end{aligned}$$

Thus w.h.p.

$$A > a\frac{7}{4} = a\left(1 + \frac{3}{4}\right).$$

F. Proof of Lemma 3

Thanks to Claim 3 we can say that $a \geq \log n$ within $\mathcal{O}(\log n)$. Then we can iteratively apply the Claim 2 in order to say that within $\mathcal{O}(\log n)$ rounds the number of undecided nodes has to drop below $\frac{1}{2}n$ and the process enters in H_2 . Note that can also happen that the process directly enters in H_4 because in these rounds the bias is increased. Thanks to Claim 1 the process does not enter into H_3 or H_5 . Since every step of the iterations holds w.h.p. and the number of steps is $\mathcal{O}(\log n)$, we easily obtain the thesis by a simple application of the Union Bound. ■

G. Sketch of proof of Lemma 4

For any round t and any $X_t = x \in \Omega$ such that $f(x) \geq h\sqrt{n}$ we define t as a *successful* round if $f(X_{t+1}) \geq (1 + \varepsilon)f(X_t)$. Let's assume that each time a round is not successful then the Markov Chain restarts from any state $x \in \Omega$ such that $f(x) \geq h\sqrt{n}$. We define the random variable $Z_t = \frac{f(X_t)}{\sqrt{n}}$. Then we define a potential function $Y_t = \exp(m - Z_t)$ and we compute its expectation at the next round:

$$\begin{aligned} \mathbf{E}[Y_{t+1}|Z_t = z_t] &\leq \mathbf{P}(f(X_{t+1}) < (1 + \varepsilon)z_t) e^m \\ &\quad + \mathbf{P}(f(X_{t+1}) \geq (1 + \varepsilon)z_t) e^{m-(1+\varepsilon)z_t} \\ &\leq e^{-c_2 z_t^2} \cdot e^m + 1 \cdot e^{m-(1+\varepsilon)z_t} \\ &= e^{m-c_2 z_t^2} + e^{m-x_t-\varepsilon z_t} \\ &\leq e^{m-3c_2 z_t} + e^{m-x_t-\varepsilon z_t} \\ &= e^{m-z_t}(e^{-(3c_2-1)z_t} + e^{-\varepsilon z_t}) \\ &= y_t(e^{-(3c_2-1)h} + e^{-\varepsilon h}), \end{aligned}$$

$$= \frac{y_t}{e},$$

where we used that $z_t \geq h$, $z_t^2 \geq 3z_t$ and that $(e^{-(3c-1)h} + e^{-\varepsilon h}) \leq (e^{-2} + e^{-2}) < e^{-1}$. Note that if $Y_T \leq 1$ then $X_T \geq m$. Thanks to the Markov inequality:

$$\mathbf{P}(Y_T > 1) \leq \frac{\mathbf{E}[Y_T]}{1} \leq \frac{\mathbf{E}[Y_{T-1}]}{e} \leq \dots \leq \frac{\mathbf{E}[Y_0]}{e^T} \leq \frac{e^m}{e^T}.$$

Choosing $T = m + \log n \leq 2 \log n$ we have that w.h.p. in T rounds the process has reached the target value.

We assumed that at each not successful round the Markov Chain restarts from a state $x \in \Omega$ such that $f(x) \geq h\sqrt{n}$. In the case the Markov Chain instead jumps in the set of states such that $f(x) < h\sqrt{n}$ we need to count also the rounds needed to let the Markov Chain exit from such set.

We know that, starting from a balanced configuration the number of rounds needed to reach a configuration such that $f(x) \geq h\sqrt{n}$ is dominated by a geometric variable of parameter c_1 . In the worst case each couple of rounds in our sequence $\{Y_t\}_t$ can be interval by a geometric random variable. How many can be these random variables? At most the number $T = 2 \log n$ needed to have that $Y_t \leq 1$ w.h.p. We know that the sum of $O(\log n)$ random geometric random variable is $\mathcal{O}(\log n)$ w.h.p. (Chernoff bound for geometric variable, see Theorem 3.7 in [13]). ■

H. Proof of Claim 5

The upper bound is a consequence of Claim 1. In order to show that $Q \leq \frac{1}{2}n$ w.h.p., from (4) we get

$$\mathbf{E}[Q | X_t = \mathbf{x}] = \frac{2q^2 + (n-q)^2 - s^2}{2n} \leq \frac{2q^2 + (n-q)^2}{2n}.$$

Note that if $\frac{1}{18}n \leq q \leq \frac{1}{2}n$ then the maximum of $2q^2 + (n-q)^2$ is in $q = \frac{1}{18}n$. Then we get

$$\begin{aligned} \frac{2q^2 + (n-q)^2}{2n} &\leq \frac{\frac{2}{18^2}n^2 + (n - \frac{1}{18}n)^2}{2n} \\ &= \frac{n^2(\frac{2}{18^2} + \frac{17^2}{18^2})}{2n} \\ &= \frac{2 + 17^2}{2 \cdot 18^2}n \\ &= \frac{291}{648}n \\ &= (\frac{1}{2} - \frac{66}{648})n. \end{aligned}$$

By using the additive form of the Chernoff bound (see (11) in Appendix B) with $\lambda = \frac{66}{648}n$, we obtain

$$\begin{aligned} \mathbf{P}\left(Q \geq \frac{1}{2}n\right) &\leq \mathbf{P}\left(Q \geq \mathbf{E}[Q | X_t = \mathbf{x}] - \lambda\right) \\ &\leq e^{-2 \cdot 66^2 n^2 / 648^2 n} \\ &= e^{-\Theta(n)}. \end{aligned}$$

I. Proof of Claim 6

Recall that $S = A - B$. In order to show that $S > s(1 + \frac{1}{36})$ w.h.p., we provide two independent bounds to the values of A and B , respectively. We use the additive form of the Chernoff bound ((10) and (11) in Appendix B) with $\lambda = \frac{\gamma\sqrt{n\log n}}{72}$. Hence, we have

$$\begin{aligned} \mathbf{P}(A \leq \mathbf{E}[A | X_t = \mathbf{x}] - \lambda) &\leq e^{-2\lambda^2/n} \\ &= e^{-2\gamma^2 \log n / 72^2} \\ &= \frac{1}{n^{\Theta(1)}}, \end{aligned}$$

and

$$\begin{aligned} \mathbf{P}(B \geq \mathbf{E}[B | X_t = \mathbf{x}] + \lambda) &\leq e^{-2\lambda^2/n} \\ &= e^{-2\gamma^2 \log n / 72^2} \\ &= \frac{1}{n^{\Theta(1)}}. \end{aligned}$$

Then w.h.p.

$$\begin{aligned} S &> \mathbf{E}[A | X_t = \mathbf{x}] - \lambda - [B | X_t = \mathbf{x}] - \lambda \\ &= \mathbf{E}[A - B | X_t = \mathbf{x}] - 2\lambda \\ &= \mathbf{E}[S | X_t = \mathbf{x}] - 2\lambda \\ &= s(1 + \frac{q}{n}) - \frac{\gamma\sqrt{n\log n}}{36} \\ &\geq s(1 + \frac{q}{n}) - s/36 \\ &\geq s(1 + \frac{1}{18} - \frac{1}{36}) \\ &= s(1 + \frac{1}{36}). \end{aligned}$$

We now show that $S < 2s$ w.h.p. using similar arguments as above. Once again, we use the additive form of the Chernoff bound with $\lambda = \frac{\gamma\sqrt{n\log n}}{4}$. We have

$$\begin{aligned} \mathbf{P}(A \geq \mathbf{E}[A | X_t = \mathbf{x}] + \lambda) &\leq e^{-2\lambda^2/n} \\ &= e^{-2\gamma^2 \log n / 16} \\ &= \frac{1}{n^{\Theta(1)}}, \end{aligned}$$

and

$$\begin{aligned} \mathbf{P}(B \leq \mathbf{E}[B | X_t = \mathbf{x}] - \lambda) &\leq e^{-2\lambda^2/n} \\ &= e^{-2\gamma^2 \log n / 16} \\ &= \frac{1}{n^{\Theta(1)}}. \end{aligned}$$

As a consequence, we have that w.h.p.

$$\begin{aligned} S &< \mathbf{E}[A | X_t = \mathbf{x}] + \lambda - [B | X_t = \mathbf{x}] + \lambda \\ &= \mathbf{E}[A - B | X_t = \mathbf{x}] + 2\lambda \\ &= \mathbf{E}[S | X_t = \mathbf{x}] + 2\lambda \end{aligned}$$

$$\begin{aligned}
&= s\left(1 + \frac{q}{n}\right) + \frac{\gamma}{2}\sqrt{n \log n} &&= \frac{1}{n^{\Theta(1)}}. \\
&< s\left(1 + \frac{1}{2}\right) + \frac{1}{2}s \\
&= 2s.
\end{aligned}$$

J. Proof of Claim 7

From (5), since $s \geq \frac{2}{3}n$, we have that

$$\begin{aligned}
\mathbf{E}[B | X_t = \mathbf{x}] &= b \left(1 - \frac{2s + 3b - n}{n}\right) \\
&\leq b \left(1 - \frac{2s - n}{n}\right) \\
&\leq b \left(1 - \frac{\frac{4}{3}n - n}{n}\right) \\
&= b \left(1 - \frac{1}{3}\right).
\end{aligned}$$

Thus, we apply the multiplicative form of the Chernoff Bound ((8) in Appendix A) with $\delta = \frac{1}{3}$, and we obtain

$$\begin{aligned}
\mathbf{P}\left(B \geq (1 + \delta) \left(1 - \frac{1}{3}\right) b\right) &\leq e^{-b(1 - \frac{1}{3})\delta^2/3} \\
&\leq e^{-\log n(1 - \frac{1}{3})\delta^2/3} \\
&= \frac{1}{n^{\Theta(1)}}.
\end{aligned}$$

As a consequence, we have that w.h.p.

$$\begin{aligned}
B &\leq b(1 + \delta) \left(1 - \frac{1}{3}\right) \\
&= b \left(1 + \frac{1}{3}\right) \left(1 - \frac{1}{3}\right) \\
&= b \left(1 - \frac{1}{9}\right).
\end{aligned}$$

K. Proof of Claim 8

We recall that $S = A - B$, thus we provide two independent bounds to the values of A and B respectively. We use the additive form of the Chernoff bound ((10) and (11) in Appendix B) with $\lambda = \varepsilon \frac{\sqrt{n \log n}}{2}$. We have

$$\begin{aligned}
\mathbf{P}(A \leq \mathbf{E}[A | X_t = x] - \lambda) &\leq e^{-2\lambda^2/n} \\
&= e^{\varepsilon^2 \log n/2} \\
&= \frac{1}{n^{\Theta(1)}},
\end{aligned}$$

and

$$\begin{aligned}
\mathbf{P}(B \geq \mathbf{E}[B | X_t = x] + \lambda) &\leq e^{-2\lambda^2/n} \\
&= e^{\varepsilon^2 \log n/2}
\end{aligned}$$

Then it holds that, w.h.p.

$$\begin{aligned}
S &\geq \mathbf{E}[A | X_t = x] - \lambda - \mathbf{E}[B | X_t = x] - \lambda \\
&= \mathbf{E}[A - B | X_t = x] - 2\lambda \\
&= \mathbf{E}[S | X_t = x] - 2\lambda \\
&= s\left(1 + \frac{q}{n}\right) - \varepsilon\sqrt{n \log n} \\
&\geq s + \frac{2\sqrt{n \log n}}{3} - \varepsilon\sqrt{n \log n} \\
&> s
\end{aligned}$$

L. Proof of Claim 9

The number of Beta-colored nodes is at least $2\sqrt{n \log n}$ and each node has probability at least $2/3$ to pick a Alpha-colored node. Thus $\mathbf{E}[Q] > \frac{4}{3}\sqrt{n \log n}$ and we get the claim by a simple application of the additive form of the Chernoff bound.

M. Proof of Claim 10

From (2), we have:

$$\begin{aligned}
\mathbf{E}[Q | X_t = \mathbf{x}] &= \frac{q^2 + 2ab}{n} \leq \frac{q^2 + 4n\sqrt{n \log n}}{n} \\
&= q\left(\frac{q}{n} + \frac{4\sqrt{n \log n}}{q}\right) \\
&\leq q\left(\frac{1}{3} + \frac{1}{3}\right) \\
&= q\left(1 - \frac{1}{3}\right).
\end{aligned}$$

Thus we apply the multiplicative form of the Chernoff Bound (8 in Appendix A) with $\delta = \frac{1}{3}$

$$\begin{aligned}
\mathbf{P}\left(Q \geq (1 + \delta) \left(1 - \frac{1}{3}\right) q\right) &\leq e^{-q(1 - \frac{1}{3})\delta^2/3} \\
&\leq e^{-\log n(1 - \frac{1}{3})\delta^2/3} \\
&= \frac{1}{n^{\Theta(1)}},
\end{aligned}$$

and thus we get that, w.h.p.

$$Q \leq \left(1 + \frac{1}{3}\right) \left(1 - \frac{1}{3}\right) q = q \left(1 - \frac{1}{9}\right).$$

N. Proof of Claim 11

We first show that in one round the number nodes that support the color Beta becomes logarithmic and the number of undecided nodes does not increase.

$$\mathbf{E}[B | X_t = \mathbf{x}] = b \left(\frac{b + 2q}{n}\right)$$

$$\begin{aligned} &\leq 2\sqrt{n \log n} \left(\frac{2\sqrt{n \log n} + 2\gamma\sqrt{n \log n}}{n} \right) \\ &= 4(\gamma + 1) \log n. \end{aligned}$$

It is immediate concentrate using the multiplicative form of the Chernoff bound and get that $B < 8(\gamma + 1) \log n$ w.h.p. We now show that the number of the undecided nodes is still $\mathcal{O}(\sqrt{n \log n})$. Indeed

$$\begin{aligned} \mathbf{E}[Q | X_t = \mathbf{x}] &= \frac{q^2}{n} + \frac{2ab}{n} \\ &\leq \gamma^2 \log n + 4\sqrt{n \log n}. \end{aligned}$$

Then using the additive form of the chernoff bound we get that $Q \leq 5\sqrt{n \log n}$ w.h.p. In the next round, w.h.p., no undecided node picks a node colored of Beta or viceversa so we can conclude that there are no nodes supporting Beta left (and it easy to show that there is at least one supporter of Alpha w.h.p.). From now on, the stochastic process is equivalent to a classic spreading process via \mathcal{PULL} operations, and thus, in the $\mathcal{O}(\log n)$ rounds, all the nodes will support Alpha w.h.p. ■

O. Proof of Claim 12

We recall that $S = A - B$, thus we provide two independent bounds to the values of A and B respectively. We use the additive form of the Chernoff bound ((10) and (11) in Appendix B) with $\lambda = \varepsilon \frac{\sqrt{n \log n}}{2}$. We have

$$\mathbf{P}(A \leq \mathbf{E}[A|X_t = x] - \lambda) \leq e^{-2\lambda^2/n} = e^{\varepsilon^2 \log n/2} = \frac{1}{n^{\Theta(1)}},$$

and

$$\mathbf{P}(B \geq \mathbf{E}[B|X_t = x] + \lambda) \leq e^{-2\lambda^2/n} = e^{\varepsilon^2 \log n/2} = \frac{1}{n^{\Theta(1)}}.$$

Then it holds that, w.h.p.

$$\begin{aligned} S &\geq \mathbf{E}[A|X_t = x] - \lambda - \mathbf{E}[B|X_t = x] - \lambda \\ &= \mathbf{E}[A - B|X_t = x] - 2\lambda \\ &= \mathbf{E}[S|X_t = x] - 2\lambda \\ &= s(1 + \frac{q}{n}) - \varepsilon\sqrt{n \log n} \\ &\geq s - \varepsilon\sqrt{n \log n} \\ &\geq \gamma\sqrt{n \log n} - \varepsilon\sqrt{n \log n} \\ &= (\gamma - \varepsilon)\sqrt{n \log n}. \end{aligned}$$

P. Tightness of Theorem 2

Claim 13. *An initial configuration exists with $|s| = \Theta(\sqrt{n})$ such that the process converges to the minority color with constant probability*

Let us consider the configuration \mathbf{x} such that $q(\mathbf{x}) = n/3, a(\mathbf{x}) = n/3 + \sqrt{n}$ and $b(\mathbf{x}) = n/3 - \sqrt{n}$. We prove

that in one round there is constant probability that the bias becomes zero or negative. After that, by simple symmetry argument, we get the claim.

We define A^q, B^q, Q^q the random variables counting the nodes that was undecided in the configuration \mathbf{x} , and in the next round are respectively colored of Alpha, Beta or undecided. Similarly A^a (B^b) counts the nodes that was supporting the color Alpha (Beta) in the configuration \mathbf{x} and that are still support the same color in the next round.

Since it is impossible that a node supporting a color in one round could support the other color in the next round, it holds that $A = A^q + A^a$ and $B = B^q + B^b$. Note that, among these random variables, only A^q and B^q are not independents. Thus, for any positive constant δ , it holds that

$$\begin{aligned} \mathbf{P}(S \leq 0) &= \mathbf{P}(B \geq A) \\ &= \mathbf{P}(B^q + B^b \geq A^q + A^a) \\ &\geq \mathbf{P}(B^q \geq A^q, B^b \geq n/3 + \delta\sqrt{n}, A^a \leq n/3 + \delta\sqrt{n}) \\ &= \mathbf{P}(B^q \geq A^q) \cdot \mathbf{P}(B^b \geq n/3 + \delta\sqrt{n}) \\ &\quad \cdot \mathbf{P}(A^a \leq n/3 + \delta\sqrt{n}). \end{aligned}$$

With a simple application of the Reverse Chernoff bound (see (12)) we get that $\mathbf{P}(B^b \geq n/3 + \delta\sqrt{n})$ is atleast constant, whereas the fact that $\mathbf{P}(A^a \leq n/3 + \delta\sqrt{n})$ is atleast constant is an immediate consequence of the additive form of the Chernoff Bound (see (11)).

Thus we need to show that also $\mathbf{P}(B^q \geq A^q)$ is atleast constant. Note that the distribution B^q conditioned to the event $Q^q = k$ is a binomial distribution with parameters $(\frac{n}{3} - k, \frac{b(\mathbf{x})}{a(\mathbf{x})+b(\mathbf{x})})$ and with expectation $\mathbf{E}[B^q | Q^q = k, X_t = \mathbf{x}] = (\frac{n}{3} - k)/2 - (\frac{n}{3} - k)/(6\sqrt{n})$. Thus we get

$$\begin{aligned} \mathbf{P}(B^q \geq A^q) &= \sum_{k=1}^{n/3} \mathbf{P}(B^q \geq A^q | Q^q = k) \mathbf{P}(Q^q = k) \\ &> \sum_{k=n/4}^{n/2} \mathbf{P}(B^q \geq A^q | Q^q = k) \mathbf{P}(Q^q = k) \\ &= \sum_{k=n/4}^{n/2} \mathbf{P}\left(B^q \geq \left(\frac{n}{3} - k\right)/2 \mid Q^q = k\right) \mathbf{P}(Q^q = k) \\ &= \sum_{k=n/4}^{n/2} \mathbf{P}\left(B^q \geq \frac{\mathbf{E}[B^b | Q^q=k]}{+(\frac{n}{3}-k)/(6\sqrt{n})} \mid Q^q = k\right) \mathbf{P}(Q^q = k) \\ &= \sum_{k=n/4}^{n/2} \mathbf{P}\left(B^q \geq \frac{\mathbf{E}[B^b | Q^q=k]}{.(1+\frac{1}{3\sqrt{n}-1})} \mid Q^q = k\right) \mathbf{P}(Q^q = k) \\ &\geq \sum_{k=n/4}^{n/2} \exp\left(\frac{-9}{(3\sqrt{n}-1)^2} \cdot \mathbf{E}[B^b | Q^q = k]\right) \mathbf{P}(Q^q = k) \\ &\geq \sum_{k=n/4}^{n/2} \exp\left(\frac{-9n}{(3\sqrt{n}-1)^2}\right) \mathbf{P}(Q^q = k) \end{aligned} \tag{15}$$

$$\begin{aligned}
&= \exp\left(\frac{-9n}{(3\sqrt{n}-1)^2}\right) \sum_{k=n/4}^{n/2} \mathbf{P}(Q^q = k) \\
&= \Theta(1) \sum_{k=n/4}^{n/2} \mathbf{P}(Q^q = k) \\
&= \Theta(1)(1 - e^{\Theta(n)}) \tag{16}
\end{aligned}$$

Where in (15) we used the reverse Chernoff bound (see (12)) and in (16) we used that $\mathbf{E}[Q^q] \approx \frac{n}{3}$ and the additive form of the Chernoff bound. ■