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

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

Mohsen Ghaffari
ETH Zürich, Switzerland
ghaffari@inf.ethz.ch

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

Emanuele Natale
Max Planck Institute for Informatics, Germany
enatale@mpi-inf.mpg.de

Francesco Pasquale[1]
Università *Tor Vergata* di Roma, Italy
pasquale@mat.uniroma2.it

Giacomo Scornavacca
Università degli Studi dell'Aquila, Italy
giacomo.scornavacca@graduate.univaq.it

Abstract

The *Undecided-State Dynamics* is a well-known protocol for distributed consensus. We analyze it in the parallel $\mathcal{PULL}$ communication model on the complete graph with n nodes for the *binary* case (every node can either support one of *two* possible colors, or be in the undecided state).

An interesting open question is whether this dynamics is an efficient *Self-Stabilizing* protocol, namely, starting from an arbitrary initial configuration, it reaches consensus *quickly* (i.e., within a polylogarithmic number of rounds). 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.

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

2012 ACM Subject Classification "F.2 Theory of Computation - Analysis of Algorithms and Problem Complexity"

Keywords and phrases Distributed Consensus, Self-Stabilization, $\mathcal{PULL}$ Model, Markov Chains.

Digital Object Identifier 10.4230/LIPIcs.MFCS.2018.28

---

[1] Partly supported by the University of "Tor Vergata" under research programme "Mission: Sustainability" project ISIDE (grant no. E81I18000110005)

1 Introduction

Simple local mechanisms for *Consensus* problems in distributed systems recently received a lot of attention [2, 1, 18, 19, 28, 31]. In one of the basic versions of the consensus problem the system consists of anonymous entities (*nodes*) each one initially supporting a *color* out of a finite set of colors Σ. Nodes run elementary operations and interact by exchanging messages. A *Consensus Protocol* is a local procedure that makes the system converge to a *monochromatic* configuration, where all nodes support the same color. Consensus has to be *valid*, i.e., the "winning" color must be one of those initially supported by at least one node. A crucial property of a consensus protocol is *self-stabilization* [6, 18, 30]: Informally, if the system is "perturbed" by some external event and moved to an arbitrary configuration, then the protocol must bring the system back to a valid consensus and, moreover, once the system reaches consensus, it must remain in that configuration forever, unless a further external event takes place[2]. Self-stabilizing consensus processes are fundamental building-blocks that play an important role in coordination tasks and self-organizing behavior in population systems [11, 13, 19, 29].

We study the consensus problem in the $\mathcal{PULL}$ *communication model* [12, 17, 24] where, at every round, each active node of a communication network contacts one neighbor uniformly at random to pull information. A natural consensus protocol in this model is the *Undecided-State Dynamics*[3] (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 neighbor and updates its state according to the following updating rule (see Table 1): 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 previously studied in both *sequential* [2] and *parallel* [4] models: Informally, in the former only one random node is activated at every round and it updates its state according to the local rule, while in the latter all nodes are activated at every round and they update their state, synchronously.

As for the sequential model[4], [2] provides an unconditional analysis showing (among other results) that the U-Dynamics is a self-stabilizing protocol for *binary* consensus (i.e., when $|\Sigma| = 2$) in the complete graph with n nodes. They show the convergence time is $\mathcal{O}(n \log n)$ (and, thus, work per node is $\mathcal{O}(\log n)$), *with high probability*[5]. This result also clarifies the algorithmic interest for this process. Indeed, the U-Dynamics can be seen as a variant of the popular *Voter Model* [9, 23, 26] where every active node simply takes the color it pulls at every round. On one hand, the Voter Model uses minimal number of node states (i.e. $|\Sigma|$) and takes $\Theta(n)$ work per node to reach consensus (see for instance [25]). On the other hand, the U-Dynamics exponentially improves the work complexity by using one additional state, only. Further motivations on the U-Dynamics are discussed in Subsection 1.2.

We remark that the stochastic process induced by the parallel dynamics significantly departs from the one induced by the sequential dynamics. As a simple evidence of such

---

[2] Notice that, according to previous work [6, 18], we require self-stabilization to hold *with high probability*.

[3] In some previous papers [31] 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.

[4] [2] in fact considers the *Population-Protocol* model which is, in our specific context, equivalent to the sequential $\mathcal{PULL}$ model.

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

qualitative differences, observe that, starting from a configuration with no undecided nodes, in the parallel case the system might end up in the *non-valid*, monochromatic configuration where all nodes are undecided (this would happen if, for example, at the first round every node pulled a node with the other color). On the other hand, it is easy to see that in the sequential case the process always ends up in a monochromatic configuration with no undecided nodes, unless it starts from a configuration with all nodes undecided. The crucial difference lies in the random number of nodes that may change color at every round: In the sequential model, this is at most one[6], while in the parallel one, *all* nodes may change state in one round and, for most phases of the process, the expected number of changes is indeed linear in n. The above difference is one of the main reasons why no general techniques are currently available to extend any quantitative analysis for the sequential process to the corresponding parallel one (and vice versa). In particular, the analysis in [2] strongly uses the fact that only one node can change state in one round in order to derive a suitable supermartingale argument to bound the stopping time of the process. It thus fully covers the case of sequential interaction models, but it is not helpful to understand the evolution of the U-Dynamics process on any interaction model in which the number of nodes that may change state in one round is not bounded by some absolute constant.

As for the parallel $\mathcal{PULL}$ model, while 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. Indeed, in [4] the authors analyze the U-Dynamics in the parallel $\mathcal{PULL}$ model on the complete graph for any number $k = o(n^{1/3})$ of colors. However, their analysis requires the initial configuration to 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 [4] it is assumed that $c_1 \geqslant \alpha c_2$, 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.

1.1 Our results

We prove that, starting from any color configuration[7] 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. Previous analysis of this phase consider either *sequential* processes of interacting particles that can be modeled as *birth-and-death* chains [2] or parallel processes whose local rule is fully symmetric w.r.t. the states/colors of the nodes (such as majority rules) [6, 18]. 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. We believe this issue has a *per-se* scientific interest since symmetry-breaking phenomena yielded

[6] This number actually becomes 2 if the sequential communication model activates a random edge per round, rather than one single node [2].
[7] Our analysis also considers initial configurations with undecided nodes.

by simple and local mechanisms plays a central role in key aspects of population systems [10] and, more generally, in the emerging field of natural algorithms [13].

Informally speaking, in Section 4 we deal with almost-balanced starting configurations. By devising a coupling to a "simplified" pruned process, we show that the analysis of this *symmetry-breaking* phase essentially reduces to the analysis of a specific regime where the number q of undecided nodes remains a suitable constant fraction of n *until* the magnitude of the bias s reaches $\Omega(\sqrt{n \log n})$: 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 for two main reasons: Along this regime, (i) the *variance* of the bias s is large (i.e. $\Theta(n)$) and (ii) whenever the bias s is $\Omega(\sqrt{n})$, its drift turns out to be *exponential* with non-negligible, increasing probability (w.r.t. s itself). Then, we prove a variant of a general Lemma [18] that provides a logarithmic bound on the hitting time of Markov chains satisfying 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 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* behavior 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.

1.2 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 modeled as so-called chemical reaction networks [11, 19]. For such reasons, the convergence time of this dynamics has been analyzed on different communication models [2, 3, 4, 27, 31]. As previously mentioned, the U-Dynamics has been analyzed in the parallel $\mathcal{PULL}$ model in [4] 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 [2] on 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 $\omega\left(\sqrt{n \log n}\right)$.

Still concerning the sequential model, [27] recently analyzes, besides other protocols, the U-Dynamics in arbitrary graphs where in the initial configuration each node samples uniformly at random one out of 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 result relies on an exponentially-small upper bound on the probability that a certain minority can win in the complete graph (see [27] for more details). In [3, 7, 20, 31], the same dynamics for the binary case has been analyzed in other 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 [2, 5, 9, 11, 15, 16, 18, 31].

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). The authors of [5] 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 $k \leqslant (n/\log n)^{1/4}$. 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 [8].

In [18] the authors provide an analysis of the *3-median* rule, in which every node updates its value to the median of its current value and the values of two randomly chosen neighbors. 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 [18].

Very recently, [22] provides an optimal bound $\Theta(k \log n)$ for the 2-Choices dynamics on the complete graph even under some dynamic adversary. In [15, 16], the authors consider the 2-Choices dynamics for plurality consensus in the binary case (i.e. $k = 2$). For random d-regular graphs, [15] 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 [16], these results are extended to general expander graphs.

1.3 Structure of the paper

In Section 2, we provide some preliminaries and an informal description of the expected evolution of the U-Process. In Section 3, we formally state the main results of this paper and describe an outline of the corresponding proofs. Section 4 is devoted to the description of the tight analysis of the symmetry-breaking phase. The analysis of the “majority” phase of the process is given in Section 5. Conclusions and some open questions are discussed in Section 6. Due to lack of space, all the omitted proofs can be found in the full-version of the paper [14].

2 Preliminaries

We analyze the parallel version of the dynamics called U-Dynamics in the (uniform) $\mathcal{PULL}$ model on the complete graph: Starting from an initial configuration where every node supports a color, i.e. a value from a set Σ of k possible colors[8], 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 a

[8] W.l.o.g. we can define $\Sigma = [k]$ where $[k] = \{1, 2, \cdots, k\}$

node is in the undecided state and pulls a color, it gets that color. Finally, a node that pulls either an undecided node or a node with its own color remains in its current state.

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

Table 1 The update rule of the 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 the graph is complete, a configuration $\mathbf{x} \in \mathcal{C}$ is uniquely determined by fixing the number of `Alpha`-colored nodes and the number of `Beta`-colored ones, 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 $s(\mathbf{x}) = a(\mathbf{x}) - b(\mathbf{x})$ called the *bias* of $\mathbf{x}$. 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 $\{\mathbf{X}_t\}_{t \geqslant 0}$ 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 calculate the following expected values (see also Section 3 in [4]):

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

$$\mathbf{E}\left[Q \,|\, \mathbf{X}_t = \mathbf{x}\right] = \frac{q^2 + 2ab}{n}, \tag{2}$$

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

2.1 The expected evolution of the U-Process

Equations (1)-(3) can be used to have a preliminary intuitive idea on the expected evolution of the U-Process. 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, $q \geqslant \delta n$, for some positive constant δ). By rewriting (2) in terms of q and s we have that

$$\mathbf{E}\left[Q \,|\, \mathbf{X}_t = \mathbf{x}\right] = \frac{q^2 + 2ab}{n} = \frac{2q^2 + (n-q)^2 - s^2}{2n} \geqslant \frac{n}{3} - \frac{s^2}{2n}, \tag{4}$$

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

When the magnitude of the bias s reaches $\frac{2}{3}n$, it is easy to see that the expected number of nodes with the *minority* color decreases exponentially. Indeed, suppose w.l.o.g. that `Beta` is the minority color and rewrite (1) for B and in terms of b and s. We get

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

Hence, when $s > \frac{2}{3}n$ we have that $\mathbf{E}[B \mid \mathbf{X}_t = \mathbf{x}] \leqslant \frac{2}{3}b$.

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* looses its predictive power. As an extreme example, observe that when $a = b = \frac{n}{3}$ 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 4. 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.

3 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 $\mathcal{PULL}$ 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. $q \leqslant n-1$), then the process converges to a configuration such that $|s| = n$, w.h.p.*

▶ **Theorem 2** (Plurality consensus). *Let γ be any positive constant. Assume that the U-Process starts at any biased configuration such that $|s| \geqslant \gamma\sqrt{n \log n}$ 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 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[9] 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 $s \leqslant \gamma\sqrt{n \log n}$ and $q \geqslant \frac{1}{2}n$; H_2 is the set of configurations such that $s \leqslant \gamma\sqrt{n \log n}$ and $\frac{1}{18}n \leqslant q \leqslant \frac{1}{2}n$; H_3 is the set of configurations such that

[9] We remark that our analysis focuses on asymptotic bounds and it does not definitely optimize the corresponding constants: However, using technicalities and loosing readability, all such constants can be largely improved.

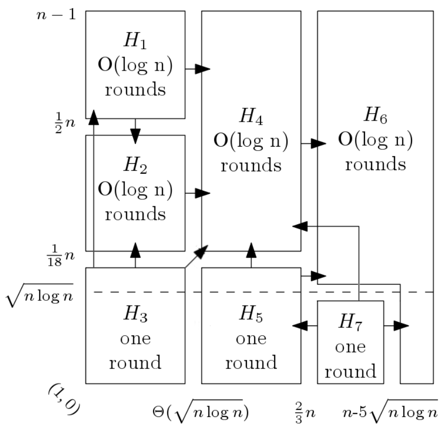

Figure 1 $\{H_1, \ldots, H_7\}$ is the considered partitioning of the configuration space $\mathcal{C}$. On the x axis we represent the bias s, on the y axis the number of undecided nodes q. Missing arrows are transitions that have negligible probabilities.

$s \leqslant \gamma\sqrt{n \log n}$ and $q \leqslant \frac{1}{18}n$. H_4 is the set of configurations such that $\gamma\sqrt{n \log n} \leqslant s \leqslant \frac{2}{3}n$ and $q \geqslant \frac{1}{18}n$; H_5 is the set of configurations such that $\gamma\sqrt{n \log n} \leqslant s \leqslant \frac{2}{3}n$ and $q \leqslant \frac{1}{18}n$; H_7 is the set of configurations such that $\frac{2}{3}n \leqslant s \leqslant n - 5\sqrt{n \log n}$ and $q \leqslant \sqrt{n \log n}$. H_6 is the set of configurations such that $s \geqslant \frac{2}{3}n$ minus H_7.

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

Observe that the scheme highlighted in Fig. 1 can be seen as a directed acyclic graph G having a single sink, H_6, that is reachable from any other region. We remark that, starting from certain configurations, a monochromatic state may be reached via different paths in G. This departs from previous analysis of consensus processes [4, 5, 18] 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 $s(\mathbf{x}) \geqslant \gamma\sqrt{n \log n}$, for some positive constant γ, and assume w.l.o.g. that the majority color in $\mathbf{x}$ is `Alpha`. We first show that, if the initial configuration $\mathbf{x}$ is in H_4, then the bias grows exponentially fast and thus the process enters in H_6 within $\mathcal{O}(\log n)$ rounds. Then we prove that, once in H_6, the process ends in the monochromatic configuration where $a = n$ within $\mathcal{O}(\log n)$ rounds, w.h.p. All the other cases "reduce" to the above ones in at most two rounds. Indeed, we show that, starting from any configuration in H_5, the process falls into H_4 or H_6 in one round and that, starting from any configuration in H_7, the process falls into H_4, H_5 or H_6 in one round. As for the tightness of the result stated in the second part of the theorem, we have that the lower bound (Claim (i)) on the convergence time is an immediate consequence of Claim (ii) of Lemma 12, while the second claim, concerning the lower bound on the initial bias, is proved in the full version of the paper [14].

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 4, 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 8 we then prove that, if the process is in a configuration in H_2, then the bias s will become $\Omega(\sqrt{n \log n})$ within $\mathcal{O}(\log n)$ rounds, w.h.p.

4 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 at 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 8 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** (Phases H_1 and H_3: Starters).

- *Starting from any configuration $\boldsymbol{x} \in H_3$, the U-Process reaches a configuration $\boldsymbol{X}' \in (H_1 \cup H_2 \cup H_4)$ in one round, w.h.p.*
- *Starting from any configuration $\boldsymbol{x} \in H_1$, the U-Process reaches a configuration $\boldsymbol{X}' \in (H_2 \cup H_4)$ within $\mathcal{O}(\log n)$ rounds, 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 the following bound on the hitting time of any Markov chain having suitable drift properties. This result is a variant of Claim 2.9 in [18] that requires a new proof.

▶ **Lemma 4.** *Let $\{X_t\}_{t \in \mathbb{N}}$ 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 c_3 be any positive constant and let $m = c_3\sqrt{n} \log n$ be a target value. Assume the following properties hold:*

1. *For any positive constant h, a positive constant $c_1 < 1$ exists such that, for any $x \in \Omega$ with $f(x) < m$, it holds that*
$$\mathbf{P}\left(f(X_{t+1}) < h\sqrt{n} | X_t = x\right) < c_1 \,,$$
2. *Two positive constants ε, c_2 exist such that, for any $x \in \Omega$ with $f(x) < m$, it holds that*
$$\mathbf{P}\left(f(X_{t+1}) < (1+\varepsilon) f(X_t) | X_t = x\right) < e^{-c_2 f(x)^2/n} \,.$$

Then the process reaches a state x such that $f(x) \geqslant m$ within $\mathcal{O}(\log n)$ rounds, w.h.p.

Proof. We first define a set of hitting times $T = \{\tau(i)\}_{i \in \mathbb{N}}$ where

$$\tau(i) = \inf_{t \in \mathbb{N}} \{t : t > \tau(i-1), f(X_t) \geqslant h\sqrt{n}\}$$

setting $\tau(0) = 0$. By Hypothesis (1), for every $i \in \mathbb{N}$, the expectation of $\tau(i)$ is finite. Then we define the following stochastic process which is a subsequence of $\{X_t\}_{t \in \mathbb{N}}$: $\{R_i\}_{i \in \mathbb{N}} =$

$\{X_{\tau(i)}\}_{i\in\mathbb{N}}$. Observe that $\{R_i\}_{i\in\mathbb{N}}$ is still a Markov Chain. Indeed, let $\{x_1,\dots,x_{i-1}\}$ a set of states in Ω:

$$\begin{aligned}
&\mathbf{P}\left(R_i = x | R_{i-1} = x_{i-1} \wedge \dots \wedge R_1 = x_1\right) = \mathbf{P}\left(X_{\tau(i)} = x | X_{\tau(i-1)} = x_{i-1} \wedge \dots \wedge X_{\tau(1)} = x_1\right)\\
&= \sum_{t(i)\wedge\dots\wedge t(0)\in\mathbb{N}} \mathbf{P}\left(X_{t(i)} = x | X_{t(i-1)} = x_{i-1} \wedge \dots \wedge X_{t(1)} = x_1\right)\\
&\cdot \mathbf{P}\left(\tau(i) = t(i) \wedge \tau(i-1) = t(i-1) \wedge \dots \wedge \tau(1) = t(1)\right)\\
&= \sum_{t(i)\wedge\dots\wedge t(0)\in\mathbb{N}} \mathbf{P}\left(X_{t(i)} = x | X_{t(i-1)} = x_{i-1}\right)\\
&\cdot \mathbf{P}\left(\tau(i) = t(i) \wedge \tau(i-1) = t(i-1) \wedge \dots \wedge \tau(1) = t(1)\right)\\
&= \mathbf{P}\left(X_{\tau(i)} = x | X_{\tau(i-1)} = x_{i-1}\right) = \mathbf{P}\left(R_i = x | R_{i-1} = x_{i-1}\right).
\end{aligned}$$

By definition the state space of R is $\{x\in\Omega : f(x) \geqslant h\sqrt{n}\}$. Moreover Hypothesis (2) still holds for this new Markov Chain. Indeed:

$$\begin{aligned}
&\mathbf{P}\left(f(R_{i+1}) < (1+\varepsilon) f(R_i) | R_i = x\right)\\
&= 1 - \mathbf{P}\left(f(R_{i+1}) \geqslant (1+\varepsilon) f(R_i) | R_i = x\right)\\
&= 1 - \mathbf{P}\left(f(X_{\tau(i+1)}) \geqslant (1+\varepsilon) f(X_{\tau(i)}) | X_{\tau(i)} = x\right)\\
&\leqslant 1 - \mathbf{P}\left(f(X_{\tau(i+1)}) \geqslant (1+\varepsilon) f(X_{\tau(i)}) \wedge \tau(i+1) = \tau(i)+1 | X_{\tau(i)} = x\right)\\
&= 1 - \mathbf{P}\left(f(X_{\tau(i)+1}) \geqslant (1+\varepsilon) f(X_{\tau(i)}) | X_{\tau(i)} = x\right)\\
&= 1 - \mathbf{P}\left(f(X_{t+1} \geqslant (1+\varepsilon) f(X_t) | X_t = x)\right) < e^{-c_2 f(x)^2/n}.
\end{aligned}$$

These two properties are sufficient to study the number of rounds required by the new Markov Chain $\{R_i\}_{i\in\mathbb{N}}$ to reach the target value m. Indeed, by defining the random variable $Z_i = \frac{f(R_i)}{\sqrt{n}}$ and considering the following "potential" function, $Y_i = \exp(\frac{m}{\sqrt{n}} - Z_i)$ we can compute its expectation at the next round as follows. Let us fix any state $x \in \Omega$ such that $h\sqrt{n} \leqslant f(x) < m$ and define $z = \frac{f(x)}{\sqrt{n}}$ and $y = \exp(\frac{m}{\sqrt{n}} - z)$. We get:

$$\begin{aligned}
\mathbf{E}\left[Y_{i+1} | R_i = x\right] &\leqslant \mathbf{P}\left(f(R_{i+1}) < (1+\varepsilon) f(x)\right) e^{m/\sqrt{n}}\\
&+ \mathbf{P}\left(f(R_{i+1}) \geqslant (1+\varepsilon) f(x)\right) e^{m/\sqrt{n}-(1+\varepsilon)z}\\
\text{(from Hypothesis (2))} &\leqslant e^{-c_2 z^2} \cdot e^{m/\sqrt{n}} + 1 \cdot e^{m/\sqrt{n}-(1+\varepsilon)z} \qquad (6)\\
&= e^{m/\sqrt{n}-c_2 z^2} + e^{m/\sqrt{n}-z-\varepsilon z}\\
&= e^{m/\sqrt{n}-z}(e^{z-c_2 z^2} + e^{-\varepsilon z}) \leqslant e^{m/\sqrt{n}-z}(e^{-2} + e^{-2}) \qquad (7)\\
&< \frac{e^{m/\sqrt{n}-z}}{e} < \frac{y}{e},
\end{aligned}$$

where in (7) we used that z is always at least h and thanks to Hypothesis (1) we can choose a sufficiently large h. By applying the Markov inequality and iterating the above bound, we get:

$$\mathbf{P}\left(Y_i > 1\right) \leqslant \frac{\mathbf{E}\left[Y_i\right]}{1} \leqslant \frac{\mathbf{E}\left[Y_{i-1}\right]}{e} \leqslant \dots \leqslant \frac{\mathbf{E}\left[Y_0\right]}{e^{\tau_R}} \leqslant \frac{e^{m/\sqrt{n}}}{e^i}.$$

We observe that if $Y_i \leqslant 1$ then $R_i \geqslant m$, thus by setting $i = m/\sqrt{n} + \log n = (c_3+1)\log n$, we get:

$$\mathbf{P}\left(R_{(c_3+1)\log n} < m\right) = \mathbf{P}\left(Y_{(c_3+1)\log n} > 1\right) < \frac{1}{n}. \qquad (8)$$

Our next goal is to give an upper bound on the hitting time $\tau_{(c_3+1)\log n}$. Note that the event "$\tau_{(c_3+1)\log n} > c_4 \log n$" holds if and only if the number of rounds such that $f(X_t) \geqslant h\sqrt{n}$ (before round $c_4 \log n$) is less than $(c_3+1)\log n$. Thanks to Hypothesis (1), at each round t there is at least probability $1 - c_1$ that $f(X_t) \geqslant h\sqrt{n}$. This implies that, for any positive constant c_4, the probability $\mathbf{P}\left(\tau_{(c_3+1)\log n} > c_4 \log n\right)$ is bounded by the probability that, within $c_4 \log n$ independent Bernoulli trials, we get less than $(c_3+1)\log n$ successes, where the success probability is at least $1 - c_1$. We can thus choose a sufficiently large c_4 and apply the multiplicative form of the Chernoff bound (see e.g. Theorem 1.1 in[21]) and obtain:

$$\mathbf{P}\left(\tau_{(c_3+1)\log n} > c_4 \log n\right) < \frac{1}{n}. \tag{9}$$

We are now ready to prove the Lemma using (8) and (9), indeed:

$$\begin{aligned}
\mathbf{P}\left(\exists t \leqslant c_4 \log n : X_t \geqslant m\right) &> \mathbf{P}\left(R_{(c_3+1)\log n} \geqslant m \wedge \tau_{(c_3+1)\log n} \leqslant c_4 \log n\right) \\
&= 1 - \mathbf{P}\left(R_{(c_3+1)\log n} < m \vee \tau_{(c_3+1)\log n} > c_4 \log n\right) \\
&\geqslant 1 - \mathbf{P}\left(R_{(c_3+1)\log n} < m\right) - \mathbf{P}\left(\tau_{(c_3+1)\log n} > c_4 \log n\right) \\
&> 1 - \frac{2}{n}.
\end{aligned}$$

Hence, choosing a suitable big c_4, we have shown that in $c_4 \log n$ rounds the process reaches the target value m, 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 having bias $\Omega(\sqrt{n \log n})$. To this aim, we first show that, for any configuration in H_2, Properties 1 and 2 in Lemma 4 are satisfied.

▶ Claim 5. Let $\mathbf{x} \in \mathcal{C}$ be any configuration such that $\frac{n}{18} \leqslant q(\mathbf{x}) \leqslant \frac{n}{2}$ and $|s(\mathbf{x})| < c_4\sqrt{n}\log n$ for any positive constant c_4, then it holds:

1. for any constant $h > 0$ a constant $c_1 < 1$ exists such that $\mathbf{P}\left(|S| < h\sqrt{n} \,|\, \mathbf{X}_t = \mathbf{x}\right) < c_1$,
2. two positive constants c_2, ε exist such that $\mathbf{P}\left(|S| \geqslant (1+\varepsilon)s \,|\, \mathbf{X}_t = \mathbf{x}\right) \geqslant 1 - e^{-c_2 s^2/n}$.

It is important to observe that the above claim ensures Properties 1 and 2 of Lemma 4 whenever $\frac{1}{18}n \leqslant q \leqslant \frac{1}{2}n$. 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 4.1, we define a *pruned* process, a variant of U-Process where it is possible to apply Lemma 4. Then, in Subsection 4.2 we show a coupling between the U-Process and the pruned one.

4.1 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 $q \geqslant \frac{1}{2}n$ is negligible (see Claim 7). Thus, intuitively speaking, all the configurations *actually visited* by the process before leaving H_2 do satisfy Lemma 4. In order to make this intuitive argument rigorous, in what follows, we define a suitably *pruned* process by removing from H_2 all the *unwanted* transitions.

Let $\bar{s} \in [n]$ and $\mathbf{z}(\bar{s})$ be the configuration such that $s(\mathbf{z}(\bar{s})) = \bar{s}$ and $q(\mathbf{z}(\bar{s})) = \frac{1}{2}n$. Let $p_{\mathbf{x},\mathbf{y}}$ be 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:

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

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 5 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.

▶ **Corollary 6.** *Starting from any configuration $\boldsymbol{x} \in H_2$, the U-Pruned-Process reaches a configuration $\boldsymbol{X}' \in H_4$ within $\mathcal{O}(\log n)$ rounds, w.h.p.*

4.2 Back to the original process

The definition of the U-Pruned-Process suggests a natural coupling between the original process and the pruned one: If the two processes are in different states then they act independently, while, if they are in the same configuration $\mathbf{x}$, they move together unless the U-Process goes 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 $\mathbf{z}(s(\mathbf{y}))$. Using this coupling, we first show that, if the two processes are in the same configuration, the probability that they get separated is negligible.

▶ Claim 7. 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

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

Then, thanks to the above claim, we can show that the H_2 exit time of the pruned procedure stochastically dominates the H_2 exit time of the original process. Thus, using Corollary 6, we get the main result of this section.

▶ **Lemma 8** (Phase H_2). *Starting from any configuration $\boldsymbol{x} \in H_2$, the U-Process reaches a configuration $\boldsymbol{X}' \in H_4$ within $\mathcal{O}(\log n)$ rounds, w.h.p.*

5 Convergence to the majority

In this section we state the key technical lemmas we use to prove our second main result, namely Theorem 2, which essentially establishes that, starting from any sufficiently-biased configuration, the U-Process converges to the monochromatic configuration where all nodes support the initial majority color within $\Theta(\log n)$ rounds, w.h.p.

The proofs of the technical lemmas, as well as the almost-tightness result on the minimal magnitude of the initial bias stated in Theorem 2, can be found in the full version of the paper [14].

Phases H_5 and H_7 (Starters II)

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_4 or H_6 (Lemma 9), while if it starts in H_7 then in the next round it moves to H_4 or H_5 or H_6 (Lemma 10).

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

▶ **Lemma 10** (Phase H_7). *Starting from any configuration $\boldsymbol{x} \in H_7$ with $a > b$, the U-Process reaches a configuration $\boldsymbol{X}' \in (H_4 \cup H_5 \cup H_6)$ with $a > b$ in one round, w.h.p.*

Phase H_4 (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 11. Let γ be any positive constant and $\mathbf{x} \in \mathcal{C}$ be any configuration such that $s \geqslant \gamma\sqrt{n \log n}$ and $q \geqslant \frac{1}{18}n$. Then, it holds that $s\,(1 + 1/36) < S < 2s$, w.h.p.

The above result allows us to prove the following bounds on the time the process requires to reach Phase H_6.

▶ **Lemma 12** (Phase H_4). *Let $\boldsymbol{x} \in H_4$ be a configuration with $a > b$. Then, (i) starting from $\boldsymbol{x}$, the U-Process reaches a configuration $\boldsymbol{X}' \in H_6$ with $a > b$ within $\mathcal{O}(\log n)$ rounds, w.h.p. Moreover, (ii) an initial configuration $\boldsymbol{y} \in H_4$ exists such that the U-Process stays in H_4 for $\Omega(\log n)$ rounds, w.h.p.*

Phase H_6 (Majority takeover)

This is the phase in which, due to the large bias, 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 and that the bias is preserved round by round. Then, when $b \leqslant 2\sqrt{n \log n}$, the number of undecided nodes starts to decrease exponentially fast as well. At the very end, when there are only few nodes (i.e., $\mathcal{O}(\sqrt{n \log n})$) that do not support the majority color yet, the minority color disappears in few steps and thus the U-Process converges to majority within $\mathcal{O}(\log n)$ rounds

▶ **Lemma 13** (Phase H_6). *Starting from any configuration $\boldsymbol{x} \in H_6$ with $a > b$, the U-Process ends in the monochromatic configuration where $a = n$ within $\mathcal{O}(\log n)$ rounds, w.h.p.*

6 Conclusions

We provided a full analysis of the U-Dynamics in the parallel $\mathcal{PULL}$ model for the binary case showing it is an efficient self-stabilizing consensus protocol. 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 that our analysis can be suitably adapted in order to show that the

U-Dynamics efficiently stabilizes to a valid consensus "regime" [10] even in the presence of a dynamic adversary that can change the state of a subset of nodes of size $o(\sqrt{n})$ provided that the initial number of colored nodes is $\Omega(\sqrt{n})$.

References

1. Mohammed Amin Abdullah and Moez Draief. Global majority consensus by local majority polling on graphs of a given degree sequence. *Discrete Applied Mathematics*, 180:1–10, 2015.
2. Dana Angluin, James Aspnes, and David Eisenstat. A Simple Population Protocol for Fast Robust Approximate Majority. *Distributed Computing*, 21(2):87–102, 2008. (Preliminary version in DISC'07).
3. Arta Babaee and Moez Draief. Distributed Multivalued Consensus. In *Computer and Information Sciences III*, pages 271–279. Springer, 2013.
4. Luca Becchetti, Andrea Clementi, Emanuele Natale, Francesco Pasquale, and Riccardo Silvestri. Plurality Consensus in the Gossip Model. In *ACM-SIAM SODA'15*, pages 371–390, 2015.
5. Luca Becchetti, Andrea Clementi, Emanuele Natale, Francesco Pasquale, Riccardo Silvestri, and Luca Trevisan. Simple dynamics for plurality consensus. In *ACM SPAA'14*, pages 247–256, 2014.
6. Luca Becchetti, Andrea Clementi, Emanuele Natale, Francesco Pasquale, and Luca Trevisan. Stabilizing consensus with many opinions. In *ACM-SIAM SODA'16*, pages 620–635, 2016.
7. Florence Bénézit, Patrick Thiran, and Martin Vetterli. Interval consensus: from quantized gossip to voting. In *IEEE ICASSP'09*, pages 3661–3664, 2009.
8. Petra Berenbrink, Andrea Clementi, Peter Kling, Robert Elsässer, Frederik Mallmann-Trenn, and Emanuele Natale. Ignore or Comply? On Breaking Symmetry in Consensus. In *ACM PODC'17*, pages 335–344, 2017. (Tech. Rep. in arXiv:1702.04921 [cs.DC]).
9. Petra Berenbrink, George Giakkoupis, Anne-Marie Kermarrec, and Frederik Mallmann-Trenn. Bounds on the Voter Model in Dynamic Networks. In *ICALP'16*, 2016.
10. Alan A. Berryman and Pavel Kindlmann. *Population Systems: A General Introduction*. Springer, 2008.
11. Luca Cardelli and Attila Csikász-Nagy. The cell cycle switch computes approximate majority. *Scientific Reports*, Vol. 2, 2012.
12. Keren Censor-Hillel, Bernhard Haeupler, Jonathan Kelner, and Petar Maymounkov. Global Computation in a Poorly Connected World: Fast Rumor Spreading with No Dependence on Conductance. In *ACM STOC'12*, pages 961–970, 2012.
13. Bernard Chazelle. Natural Algorithms and Influence Systems. *Commun. ACM*, 55(12):101–110, 2012.
14. Andrea Clementi, Mohsen Ghaffari, Luciano Gualà, Emanuele Natale, Francesco Pasquale, and Giacomo Scornavacca. A Tight Analysis of the Parallel Undecided-State Dynamics with Two Colors. *CoRR*, abs/1707.05135v3, 2018. URL: `https://arxiv.org/abs/1707.05135v3`.
15. Colin Cooper, Robert Elsässer, and Tomasz Radzik. The Power of Two Choices in Distributed Voting. In *ICALP'14*, pages 435–446, 2014.
16. Colin Cooper, Robert Elsässer, Tomasz Radzik, Nicolas Rivera, and Takeharu Shiraga. Fast Consensus for Voting on General Expander Graphs. In *DISC'15*, pages 248–262, 2015.

[10] According to the notion of *stabilizing almost-consensus protocol* given in [2, 6].

17 Alan Demers, Dan Greene, Carl Hauser, Wes Irish, John Larson, Scott Shenker, Howard Sturgis, Dan Swinehart, and Doug Terry. Epidemic algorithms for replicated database maintenance. In *ACM PODC'87*, pages 1–12, 1987.

18 Benjamin Doerr, Leslie Ann Goldberg, Lorenz Minder, Thomas Sauerwald, and Christian Scheideler. Stabilizing consensus with the power of two choices. In *ACM SPAA'11*, pages 149–158, 2011.

19 David Doty. Timing in chemical reaction networks. In *ACM-SIAM SODA'14*, pages 772–784, 2014.

20 Moez Draief and Milan Vojnović. Convergence speed of binary interval consensus. *SIAM Journal on Control and Optimisation*, 50(3):1087–1109, 2012.

21 Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009.

22 Mohsen Ghaffari and Johannes Lengler. Tight analysis for the 3-majority consensus dynamics. *CoRR*, abs/1705.05583, 2017. URL: `http://arxiv.org/abs/1705.05583`, `arXiv:1705.05583`.

23 Yehuda Hassin and David Peleg. Distributed probabilistic polling and applications to proportionate agreement. *Information and Computation*, 171(2):248–268, 2001.

24 David Kempe, Alin Dobra, and Johannes Gehrke. Gossip-based computation of aggregate information. In *IEEE FOCS'03*, pages 482–491, 2003.

25 David Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2008.

26 Thomas M. Liggett. *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes*. Springer-Verlag, 1999.

27 George B. Mertzios, Sotiris E. Nikoletseas, Christoforos Raptopoulos, and Paul G. Spirakis. Determining Majority in Networks with Local Interactions and Very Small Local Memory. In *ICALP'14*, pages 871–882, 2014.

28 Elchanan Mossel, Joe Neeman, and Omer Tamuz. Majority dynamics and aggregation of information in social networks. *Autonomous Agents and Multi-Agent Systems*, 28(3):408–429, 2014.

29 Saket Navlakha and Ziv Bar-Joseph. Algorithms in nature: the convergence of systems biology and computational thinking. *Mol. Syst. Biol.*, 7, 2011.

30 Marshall Pease, Robert Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *Journal of the ACM*, 27(2):228–234, 1980.

31 Etienne Perron, Dinkar Vasudevan, and Milan Vojnovic. Using Three States for Binary Consensus on Complete Graphs. In *IEEE INFOCOM'09*, pages 2527–1535, 2009.