

Brief Announcement: On the Parallel Undecided-State Dynamics with Two Colors*

Andrea Clementi¹, Luciano Gualà¹, Francesco Pasquale¹, and Giacomo Scornavacca²

- 1 Università Tor Vergata di Roma, Rome, Italy
clementi@mat.uniroma2.it
- 2 Università Tor Vergata di Roma, Rome, Italy
guala@mat.uniroma2.it
- 3 Università Tor Vergata di Roma, Rome, Italy
pasquale@mat.uniroma2.it
- 4 Università degli Studi dell'Aquila, L'Aquila, Italy
giacomo.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 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. A 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 $O(\log^2 n)$ rounds, with high probability (w.h.p.). 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 $O(\log n)$ round, w.h.p. At the heart of our approach there is a new analysis of the symmetry-breaking phase that the process must perform in order to escape from (almost-)unbiased configurations. Previous symmetry-breaking analysis of consensus dynamics essentially concern sequential communication models (such as Population Protocols) and/or symmetric updated rules (such as majority rules).

1998 ACM Subject Classification F.2 Theory of Computation – Analysis of Algorithms and Problem Complexity

Keywords and phrases Distributed Consensus, Dynamics, Gossip model, Markov chains

Digital Object Identifier 10.4230/LIPIcs.DISC.2017.47

1 Introduction

Strong research interest has been recently focussed on the study of simple, local mechanisms for *Consensus* problems in distributed systems [3, 2, 11, 12, 16, 17]. In a basic setting of the consensus problem, the system consists of a set of n anonymous nodes that run elementary operations and interact by exchanging messages. Every node initially supports a *color* chosen

* The full version of the paper can be found at <https://arxiv.org/abs/1707.05135>.



from an alphabet Σ and a *Consensus Protocol* is a local procedure that, starting from any color configuration, let the system converge to a monochromatic configuration. The consensus is *valid* if the *winning* color is one among those initially supported by at least one node. Moreover, once the system reaches a consensus configuration it will stay there forever.

We study the consensus problem in the *PULL model* [8, 10, 14] in which, at every round, each active node contacts one neighbor uniformly at random to pull information. A well-studied consensus protocol is the *Undecided-State Dynamics* (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: 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. As for the sequential model, [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), *w.h.p.*¹ As for the parallel *PULL* model, even though it is easy to verify that the U-Dynamics achieves consensus in the complete graph (w.h.p.), 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. A 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, 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 analyze the U-Dynamics in the parallel *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, 15], 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 configuration² on the complete graph, the U-Dynamics reaches a monochromatic configuration (thus consensus) within $\mathcal{O}(\log^2 n)$ rounds, w.h.p. This bound is almost 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 challenges arise from (almost) *balanced* initial configurations where the system needs to break symmetry. A key ingredient of our analysis is a suitable application of the *martingale optional stopping theorem*. While the use of that theorem is standard in the analysis of *sequential* processes of interacting particles that can be modeled as *birth-and-death* chains, our new approach allows

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

² Our analysis also considers initial configurations with undecided nodes.

us to analyze the process yielded by running the U-Dynamics in synchronous parallel rounds, that is a somewhat “wild” process where an *unbounded* number of particles may change state at every round. 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* behaviour of the key random variables (such as the bias and the number of undecided nodes at the next round). 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.

Other related work. 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 [7, 12]. For such reasons, the convergence time of this dynamics has been analyzed on different communication models [1, 3, 4, 6, 9, 11, 13, 15, 17]. Concerning the sequential model, [15] recently analyzes 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. In [4, 6, 13, 17], the same dynamics for the binary case has been analyzed in further sequential communication models.

References

- 1 M. A. Abdullah et al. Majority consensus on random graphs of a given degree sequence. *arXiv:1209.5025*, 2012.
- 2 Mohammed Amin Abdullah et al. Global majority consensus by local majority polling on graphs of a given degree sequence. *Discrete Applied Mathematics*, 180, 2015.
- 3 Dana Angluin et al. A Simple Population Protocol for Fast Robust Approximate Majority. *Distributed Computing*, 21(2), 2008.
- 4 A. Babae et al. Distributed multivalued consensus. In *Computer and Information Sciences III*. Springer, 2013.
- 5 Luca Becchetti et al. Plurality consensus in the gossip model. In *ACM-SIAM SODA'15*, 2015.
- 6 F. Bénézit et al. Interval consensus: from quantized gossip to voting. In *IEEE ICASSP'09*, 2009.
- 7 L. Cardelli et al. The cell cycle switch computes approximate majority. *Scientific Reports*, Vol. 2, 2012.
- 8 K. Censor-Hillel et al. Global computation in a poorly connected world: Fast rumor spreading with no dependence on conductance. In *ACM STOC'12*, 2012.
- 9 C. Cooper et al. The power of two choices in distributed voting. In *ICALP'14*, 2014.
- 10 A. Demers et al. Epidemic algorithms for replicated database maintenance. In *ACM PODC'87*, 1987.
- 11 Benjamin Doerr et al. Stabilizing consensus with the power of two choices. In *ACM SPAA'11*, 2011.
- 12 David Doty. Timing in chemical reaction networks. In *ACM-SIAM SODA'14*, 2014.

47:4 Brief Announcement: On the Parallel Undecided-State Dynamics with Two Colors

- 13 Moez Draief et al. Convergence speed of binary interval consensus. *SIAM Journal on Control and Optimisation*, 50(3), 2012.
- 14 D. Kempe et al. Gossip-based computation of aggregate information. In *IEEE FOCS'03*, 2003.
- 15 George B. Mertzios et al. Determining majority in networks with local interactions and very small local memory. In *ICALP'14*, 2014.
- 16 Elchanan Mossel et al. Majority dynamics and aggregation of information in social networks. *Autonomous Agents and Multi-Agent Systems*, 28(3), 2014. doi:10.1007/s10458-013-9230-4.
- 17 Etienne Perron et al. Using Three States for Binary Consensus on Complete Graphs. In *INFOCOM'09*, 2009.