

Brief Announcement: Phase Transitions of the k -Majority Dynamics in a Biased Communication Model

Emilio Cruciani¹, Hlafo Alfie Mimun², Matteo Quattropiani², Sara Rizzo³

¹Inria Sophia Antipolis Méditerranée, I3S, Université Côte d’Azur, CNRS, France

²Department of Economics and Finance, LUISS, Italy

³Gran Sasso Science Institute, Italy

Abstract

We analyze the binary-state (either \mathcal{R} or \mathcal{B}) k -MAJORITY dynamics in a biased communication model where nodes have some fixed probability p , independent of the dynamics, of being seen in state \mathcal{B} by their neighbors. In this setting we study how p , as well as the initial unbalance between the two states, impact on the speed of convergence of the process, identifying sharp phase transitions.

1 Introduction

Designing distributed algorithms that let the nodes of a graph reach a *consensus*, i.e., a configuration of states where all the nodes agree on the same state, is a fundamental problem in distributed computing and multi-agent systems. Consensus algorithms are used in protocols for other tasks, such as leader election and atomic broadcast, and in real-world applications such as clock synchronization tasks and blockchains.

Recently there has been a growing interest in the analysis of simple local *dynamics* as distributed algorithms for the consensus problem [2, 5], inspired by simple mechanisms studied in statistical mechanics for interacting particle systems. In this scenario, nodes are *anonymous* (i.e., they do not have distinct IDs) and they have a state that evolves over time according to some common local interaction with their neighbors. Many dynamics have been investigated in such a setting, for example VOTER, where nodes copy the state of a random neighbor, and 3-MAJORITY, where nodes sample 3 neighbors with replacement and update their states to the most frequent among the samples. The time needed by such dynamics to reach a consensus is very different: for example, on the complete graph, VOTER needs $\Omega(n)$ rounds, while 3-MAJORITY converges in $\mathcal{O}(\log n)$ [2, 5].

2 Communication Model

We focus on the scenario where each of the n nodes of the underlying communication graph G has a *binary state* (either \mathcal{R} or \mathcal{B}) and the communication among the nodes proceeds in *synchronous* rounds. In this setting we analyze the k -MAJORITY dynamics, where nodes sample k neighbors from their neighborhood uniformly at random and with replacement, and then update their state to the majority of the sample; ties are broken uniformly at random.

Differently from most of the previous works, we consider a communication model which is *biased*, w.l.o.g., toward state \mathcal{B} , i.e., whenever nodes sample a neighbor they see state \mathcal{B} with some probability p , regardless of the state of the sampled node, and its true state with probability $1 - p$. This biased communication model has been introduced in [3], where it is used as a tool for the analysis of the 2-CHOICES dynamics on core-periphery networks.

3 Results

Most of the previous works rely on strong topological assumptions, e.g., considering complete graphs or expanders, to prove upper bounds on the consensus time of the dynamics. We move a step forward in this direction by removing all assumptions that depend on the topology of the graph; indeed we only require the underlying communication graph to be *sufficiently dense*, i.e., with minimum degree $\omega(\log n)$. Such a milder assumption, though, comes at the cost of a weaker notion of consensus, that we call *almost-consensus*.

Definition 1. *A process reaches an \mathcal{B} -almost-consensus whenever a fraction $1 - o(1)$ of the volume of the graph is in state \mathcal{B} .*

Pushing our techniques to their limit, we could prove a consensus on state \mathcal{B} in $\mathcal{O}(\log n)$ rounds. We would still require no topological assumptions, but just a stronger condition on the minimum degree (from $\omega(\log n)$ to $\Omega(n)$), thus dramatically restricting the class of graphs taken into account to extremely dense ones. Note that, in the biased communication model we consider, the Markov Chain that models the process has a single absorbing state where all nodes are in state \mathcal{B} . For this reason we first consider an initial configuration where all the nodes are in state \mathcal{R} and study the time needed by the process to reach a \mathcal{B} -almost-consensus. Trivially, if $p = 0$ the process remains stuck in its initial configuration, while if $p = 1$ the process reaches the absorbing state in one single round. More in general, it is intuitive that the process will converge slowly to the absorbing state if p is small and quickly if p is large. With such an intuition in mind, we prove a first *phase transition* phenomenon.

Theorem 1. *Consider the k -MAJORITY dynamics in the communication model with bias p toward state \mathcal{B} , in any graph G with minimum degree $\omega(\log n)$, and where initially all nodes are in state \mathcal{R} . For every $k \geq 3$ there exists a constant $p_k^* \in [\frac{1}{9}, \frac{1}{2}]$ such that:*

- Slow convergence: *If $p < p_k^*$, a \mathcal{B} -almost-consensus is reached in $n^{\omega(1)}$ rounds, a.a.s.¹*
- Fast convergence: *If $p > p_k^*$, a \mathcal{B} -almost-consensus is reached in $\mathcal{O}(1)$ rounds, a.a.s.*

The proof looks at the expected evolution of the fraction of neighbors in state \mathcal{B} of every node and then applies concentration of probability arguments, akin to what has been done in [3]. However, compared to [3], we have a more comprehensive scenario (k -MAJORITY for any k vs. 2-CHOICES) with substantially more precise results.

Another important innovation with respect to [3] is the presence of a second phase transition on the initial unbalance between the states. Clearly in the *fast convergence regime*, i.e., when $p > p_k^*$, the initial configuration in which none of the nodes is in state \mathcal{B} is the hardest one for the process to reach a \mathcal{B} -almost-consensus. However, the scenario changes in the *slow convergence regime*, where a different initial configuration with some of the nodes already in state \mathcal{B} could speed up the process.

Theorem 2. *Consider the k -MAJORITY dynamics in the communication model with bias $p < p_k^*$ toward state \mathcal{B} , in any graph G with minimum degree $\omega(\log n)$, and where initially every node is in state \mathcal{B} with probability $1 - q$, independently of the others. For every $k \geq 3$ there exists a constant $q_{p,k}^*$ such that:*

- Fast convergence: *If $q < q_{p,k}^*$, a \mathcal{B} -almost-consensus is reached in $\mathcal{O}(1)$ rounds, a.a.s.*
- Slow convergence: *If $q > q_{p,k}^*$, a \mathcal{B} -almost-consensus is reached in $n^{\omega(1)}$ rounds, a.a.s.*

4 Applications

Our results show that adding a bias to k -MAJORITY affects the dynamics in a non-trivial way. In particular, the arise of a metastable phase makes the framework suitable to design distributed algorithms to recover planted partitions in networks [1, 4, 6]. Consider a graph $G = ((V_1, V_2), E)$ with two clusters and such

¹We say that an event \mathcal{E}_n happens *asymptotically almost surely* (in short, *a.a.s.*) if $\Pr(\mathcal{E}_n) = 1 - o(1)$.

that, for every pair of nodes in each cluster, their fraction of neighborhood toward the other cluster is equal to some constant z , as in [1]. Let G run the k -MAJORITY and suppose that the two clusters reach a local almost-consensus on different states, say nodes in V_1 agrees on \mathcal{R} and nodes in V_2 on \mathcal{B} . The evolution of such a process can be described by the biased k -MAJORITY dynamics. In fact, note that the k -MAJORITY dynamics with bias $p = z$ toward \mathcal{B} performed by the nodes of the subgraph induced by V_1 describes the local evolution of the nodes in V_1 , when considering a worst-case scenario in which the nodes in V_2 never change color. If G is such that $z = p < p_k^*$, Theorem 1 implies that V_1 remains in an almost-consensus configuration, a.a.s. Since the same reasoning can be done for V_2 , it follows that the graph would stay in a configuration that highlights its clustered structure for $n^{\omega(1)}$ rounds, hence making k -MAJORITY a suitable dynamics for the design a distributed community detection protocol.

References

- [1] Luca Becchetti, Emilio Cruciani, Francesco Pasquale, and Sara Rizzo. Step-by-step community detection in volume-regular graphs. In *ISAAC 2019*, pages 20:1–20:23, 2019.
- [2] Colin Cooper, Robert Elsässer, Hirotaka Ono, and Tomasz Radzik. Coalescing random walks and voting on connected graphs. *SIAM J. Discrete Math.*, 27(4):1748–1758, 2013.
- [3] Emilio Cruciani, Emanuele Natale, André Nusser, and Giacomo Scornavacca. Phase transition of the 2-choices dynamics on core-periphery networks. In *AAMAS 2018*, pages 777–785, 2018.
- [4] Emilio Cruciani, Emanuele Natale, and Giacomo Scornavacca. Distributed community detection via metastability of the 2-choices dynamics. In *AAAI 2019*, pages 6046–6053, 2019.
- [5] Mohsen Ghaffari and Johannes Lengler. Nearly-tight analysis for 2-choice and 3-majority consensus dynamics. In *PODC 2018*, pages 305–313, 2018.
- [6] Nobutaka Shimizu and Takeharu Shiraga. Phase transitions of best-of-two and best-of-three on stochastic block models. In *DISC 2019*, pages 32:1–32:17, 2019.