

Stabilizing Consensus with Many Opinions*

L. Becchetti¹, A. Clementi², E. Natale¹, F. Pasquale², and L. Trevisan³

¹*Sapienza* Università di Roma, becchetti@dis.uniroma1.it, natale@di.uniroma1.it

²Università *Tor Vergata* di Roma, clementi@mat.uniroma2.it, pasquale@mat.uniroma2.it

³U.C. Berkeley, luca@berkeley.edu

October 16, 2015

Abstract

We consider the following distributed consensus problem: Each node in a complete communication network of size n initially holds an *opinion*, which is chosen arbitrarily from a finite set Σ . The system must converge toward a consensus state in which all, or almost all nodes, hold the same opinion. Moreover, this opinion should be *valid*, i.e., it should be one among those initially present in the system. This condition should be met even in the presence of a malicious adversary who can modify the opinions of a bounded subset of nodes, adaptively chosen in every round.

We consider the *3-majority dynamics*: At every round, every node pulls the opinion from three random neighbors and sets his new opinion to the majority one (ties are broken arbitrarily). Let k be the number of valid opinions. We show that, if $k \leq n^\alpha$, where α is a suitable positive constant, the 3-majority dynamics converges in time polynomial in k and $\log n$ with high probability even in the presence of an adversary who can affect up to $o(\sqrt{n})$ nodes at each round.

Previously, the convergence of the 3-majority protocol was known for $|\Sigma| = 2$ only, with an argument that is robust to adversarial errors. On the other hand, no anonymous, uniform-gossip protocol that is robust to adversarial errors was known for $|\Sigma| > 2$.

Keywords: Distributed Consensus, Byzantine Agreement, Gossip Model, Majority Rules, Markov Chains.

1 Introduction

We study the following probabilistic, *synchronous* process on a complete network of n anonymous nodes: At the beginning, each node holds an “opinion” which is an element of an arbitrary finite set Σ . We call an opinion *valid* if it is held by at least one node at the beginning. Then, in each round, the following happens: 1) every node pulls the opinion from three random nodes and sets its new opinion to the majority one (ties are broken arbitrarily), and 2) an adaptive *dynamic adversary* can arbitrarily change the opinions of a subset of the nodes, possibly choosing different subsets over different rounds. We consider *F-dynamic adversaries* that, at every round, can change the opinions of up to F nodes, possibly introducing non-valid opinions.

Let the system start from any configuration having k valid opinions with $k \leq n^\alpha$ for a suitable constant $\alpha < 1$ and consider any *F-dynamic adversary* with $F = \mathcal{O}(\sqrt{n}/(k^{5/2} \log n))$. We prove that the process converges to a configuration in which all but $O(\sqrt{n})$ nodes hold the same valid opinion within $O((k^2 \sqrt{\log n} + k \log n)(k + \log n))$ rounds, with high probability. This shows that the *3-majority dynamics* provides an efficient solution to the *stabilizing-consensus* problem in the *uniform-gossip* model. Previously, this was known only for the binary case, i.e. $|\Sigma| = 2$, while for any $|\Sigma| \geq 3$, it has been an important open question for several years [3, 14]. Furthermore, still for any $|\Sigma| \geq 3$, $o(n)$ -time convergence of the 3-majority dynamics was open even in the absence of an adversary whenever the initial bias toward some plurality opinion is not large.

In the remainder of this section, we will describe in more detail the consensus problem and various network scenarios in which it is of interest, our result in this setting, and a comparison with previous related results.

1.1 Consensus (or Byzantine agreement). The *consensus* problem in a distributed network is defined

*Partially supported by Italian MIUR-PRIN 2010-11 Project *ARS TechnoMedia*, the EU FET Project *MULTIPLEX* 317532, the National Science Foundation under Grant No. 1540685, and by the US-Israel Binational Science Foundation under Grant No. 2010451.

as follows: A collection of agents, each holding a piece of information (an element of a set Σ), interact with the goal of agreeing on one of the elements of Σ initially held by at least one agent, possibly in the presence of an adversary that is trying to disrupt the protocol. The consensus problem in the presence of an adversary (known as Byzantine agreement) is a fundamental primitive in the design of distributed algorithms [25, 27]. The goal is to design a distributed, local protocol that brings the system into a configuration that meets the following conditions: (1) *Agreement*: All non-corrupted nodes support the same opinion v ; (2) *Validity*: The opinion v must be a *valid* one, i.e., an opinion which was initially declared by at least one (non-corrupted) node; (3) *Termination*: Every non-corrupted node can correctly decide to stop running the protocol at some round.

Recently, there has been considerable interest in the design of consensus algorithms in models that severely restrict both communication and computation [3, 6, 14], both for efficiency consideration and because such models capture aspects of the way consensus is reached in social networks, biological systems, and other domains of interest in network science [2, 4, 8, 9, 16, 17, 18].

In particular, we assume an anonymous network in which nodes possess no unique IDs, nor do they have any static binding of their local link ports (i.e., nodes cannot keep track of *who sent what*). From the point of view of computation, the most restrictive setting is to assume that each node only has $\mathcal{O}(\log |\Sigma|)$ bits of memory available, i.e., it just suffices to store a constant number of opinions. We further assume that this bound extends to link bandwidth available in each round. Finally, communication capabilities are severely constrained and non-deterministic: Every node can communicate with at most a (small) constant number of random neighbors in each round. These constraints are well-captured by the *uniform-gossip* communication model [12, 21, 22]: At every round, every node can exchange a (short) message (say, $\Theta(\log(|\Sigma|))$ bits) with each of at most h random neighbors, where h is a (small) absolute constant¹. A more recent, sequential variant of the uniform-gossip model is the (*random*) *population-protocols* model [3, 1, 2] in which, in each round, a single interaction between a pair of randomly selected nodes occurs.

The classic notion of consensus is too strong and unrealistic in the aforementioned distributed settings, that

instead rely on *weaker* forms of consensus, deeply investigated in [3, 4, 5, 14]. In this paper, we consider a variant of the *stabilizing-consensus* problem [4] considered in [3]: There, a solution is required to converge to a stable *regime* in which the above three properties are guaranteed in a relaxed, still useful form². More precisely:

DEFINITION 1.1. *Starting from any initial configuration with k valid opinions, a stabilizing almost-consensus protocol must ensure the following properties:*

- Almost agreement. *In a finite number of rounds, the system must reach a regime of configurations where all but a negligible “bad” subset (i.e. having size $\mathcal{O}(n^\gamma)$ for constant $\gamma < 1$) of the nodes support the same opinion.*
- Almost validity. *The system is required to converge w.h.p. to an almost-agreement regime where all but a negligible bad set of nodes keep the same valid opinion.*
- Non termination. *In dynamic distributed systems, nodes represent simple and anonymous computing units which are not necessarily able to detect any global property.*
- Stability. *The convergence toward such a weaker form of agreement is only guaranteed to hold with high probability (in short, w.h.p.³) and only over a long period (i.e. for any arbitrarily-large polynomial number of rounds).*

We remark that no stabilizing almost-consensus protocol is known for $|\Sigma| > 2$ even in the complete graph. Here, we consider the 3-majority dynamics, maybe the simplest gossip dynamics that exhibits positive probabilistic drift toward consensus [10]. The main result of this paper is on its convergence properties in the uniform-gossip model in the presence of the adaptive F -dynamic adversary defined above.

THEOREM 1.1. *Let $k \leq n^\alpha$, for a suitable constant $\alpha < 1$, and $F = \beta\sqrt{n}/(k^{\frac{5}{2}} \log n)$ for some constant $\beta > 0$. The 3-majority dynamics is a stabilizing almost-consensus protocol in the presence of any F -dynamic adversary and its convergence time is $\mathcal{O}((k^2\sqrt{\log n} + k \log n)(k + \log n))$, w.h.p.*

As a simple consequence of our analysis, we also get the same bound on the convergence time in the presence of any F -static adversary with a larger bound on F , i.e., $F = n/k - \sqrt{kn} \log n$. The F -static adversary is the one

¹In fact, $h = 1$ in the standard uniform-gossip model. It is easy to verify that all our results still hold in this more restricted model at the cost of a constant slow-down in convergence time and local memory size.

²These relaxed convergence properties are described in detail in Section 7 of [3].

³According to the standard definition, we say that a sequence of events \mathcal{E}_n , $n = 1, 2, \dots$ holds with high probability if $\mathbf{P}(\mathcal{E}_n) = 1 - \mathcal{O}(1/n^\lambda)$ for some positive constant $\lambda > 0$.

that looks at the initial configuration, then changes the opinion of up to F arbitrary nodes and, after that, no further adversary's actions are allowed.

We remark that in [7] an $\Omega(k \log n)$ bound on the convergence-time of the 3-majority dynamics is derived which holds even when the system starts from biased configurations.

Not assuming a large initial bias of the plurality opinion considerably complicates the analysis. Indeed, the major open challenge is the analysis from (almost) uniform configurations, where the system needs to break the initial symmetry in the absence of significant drifts towards any of the initial opinions. So far, the symmetry breaking in the 3-majority dynamics has never been analyzed even in the non-adversarial case. Moreover, the phase before symmetry breaking is the one in which the adversary has more chances to cause undesired behaviours: Long delays and/or convergence towards non-valid opinions. In Section 2, after providing some preliminaries, we shall discuss the above technical challenges.

Finally, though not the focus of this paper, one may wonder whether it is possible to provide guarantees about the color that eventually achieves majority. As to this point, the results of this paper and of our previous work [7] together imply that a color is not going to become majority unless it is a *near-plurality*, i.e., (1) its size is not below the average by an excessive amount and (2) it is close to the size of the plurality color.

1.2 Previous results. Consensus problems in distributed systems have been the focus of a large body of work in several research areas, such as distributed computing [19], communication networks [28], social networks and voting systems [24, 30], distributed databases [12, 13], biological systems and chemical reaction networks [9]. For brevity's sake, we here focus on results that are closest in spirit to our work.

In [3], the authors show that w.h.p. n agents that meet at random can reach valid stabilizing almost-consensus in $\mathcal{O}(n \log n)$ pairwise interactions against an $F = o(\sqrt{n})$ -bounded dynamic adversary. The adopted protocol is the well-studied third-state protocol [3, 26]. However, their analysis (and, thus, their result) only holds for the binary case and for the *population-protocol* model: At every round only one pair of nodes can interact. The authors left the existence of protocols for the multi-valued Byzantine case as a final open question [3].

In the uniform-gossip model, in [14] 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 3-majority dynamics, thus their result implies that 3-majority is an $(F = \sqrt{n})$ -stabilizing consensus with $\mathcal{O}(\log n)$ convergence time. However, in the non-binary case, it requires Σ to be a totally-ordered set and this order to be *consistent*, i.e. all agents agree on it: This may be a strong restriction when these processes are used to model emerging behavior and self-organization in complex agent systems such as biological ones.

Furthermore, even assuming an ordered opinion set (Σ, \leq) , the 3-median rule does not guarantee the crucial property of *validity* against both F -static (and, clearly, dynamic) adversaries even for very-small bounds on F (say $F = \text{polylog}(n)$).

We strongly believe that the validity property of consensus plays a crucial role in several realistic scenarios, such as monitoring sensor networks, bio-inspired dynamic systems, and voting systems [9, 24, 30].

More recently, the 3-majority rule in the multi-opinion case (i.e. for $|\Sigma| \geq 3$) has been studied for a stronger goal than consensus, namely, *stabilizing plurality consensus* [7]. In this task, the goal is to reach an almost-stable consensus towards the valid opinion initially supported by *the plurality* of the nodes. However, the initial configuration is assumed to have a large bias towards the plurality opinion. Then, let k be the number of valid opinions, and let s be the initial difference between the largest and the second-largest opinion: By strongly exploiting the assumption $s \geq \sqrt{kn \log n}$, the authors in [7] proved that, w.h.p., the system converges to the plurality opinion within time $\Theta(k \log n)$.

Another version of binary stabilizing almost-consensus is the one studied by Yildiz et al in [30]: Here, corrupted nodes are *stubborn* agents of a social network who influence others but never change their opinions. They prove negative results under a generalized variant of the classic voter dynamics in the (Poisson-clock) population-protocol model.

2 The Process and its Analysis in a Nutshell

Preliminaries. We assume a distributed system consisting of n nodes that communicate with each other over a complete graph via the synchronous *uniform-gossip* mechanism: In every round, each node can pull information from (at most) h random neighbors, where h is an absolute constant (in this work, $h = 3$). At the onset, every node chooses an arbitrary item, called *opinion*, from an arbitrary finite set Σ . A simple dynamics for consensus is the *3-majority protocol* [7]:

In each round, every node samples three nodes uniformly at random (including itself and with repetitions) and revises its opinion according to the majority of the opinions it sees. If it sees three different opinions, it picks the first one.

Clearly, in the case of three different opinions, choosing the second or the third one would not make any difference, nor would choosing one of the observed opinions uniformly at random.

Since the communication graph is complete and nodes are anonymous, the overall system state at any round can be described by a *configuration* $\mathbf{c} := (c_1, \dots, c_{|\Sigma|})$, where the *support* c_i of opinion i is the number of nodes holding opinion i in that system's state. Given configuration \mathbf{c} , we say that an opinion i is *active* in \mathbf{c} if $c_i > 0$ and, for any set of active opinions $W \subseteq \Sigma$, we define $m(W) := \arg \min_{i \in W} c_i$. For any variable x of the process, we write $x^{(t)}$ if we are considering its value at round t and $X^{(t)}$ to denote the corresponding random variable. Furthermore, following [23], considered a configuration \mathbf{c} and a random variable X defined over the process, we write $\mathbf{P}_{\mathbf{c}}(X^{(t)} = x)$ for $\mathbf{P}(X^{(t)} = x \mid \mathbf{C}^{(0)} = \mathbf{c})$. Analogously, we write $\mathbf{E}_{\mathbf{c}}[X^{(t)}]$ for the associated conditional expectation.

The next lemma provides the expected number of nodes supporting a given opinion at round $t + 1$ (and a general upper bound to it), given the configuration at round t . The simple proof of the first equality is in [7]. It is also included in Appendix A to make the paper self-contained.

LEMMA 2.1. (SEE [7]) *Let \mathbf{c} be the configuration at round t and let $W \subseteq \Sigma$ be the subset of active opinions in \mathbf{c} . Then, for any opinion $i \in W$,*

$$\begin{aligned} \mathbf{E} \left[C_i^{(t+1)} \mid \mathbf{C}^{(t)} = \mathbf{c} \right] &= c_i \left(1 + \frac{c_i}{n} - \frac{\sum_{j \in W} c_j^2}{n^2} \right) \\ &\leq c_i \left(1 + \frac{c_i}{n} - \frac{1}{|W|} \right) \end{aligned}$$

The above upper bound easily implies that opinions whose supports fall below the average $n/|W|$ decrease in expectation. This expected drift is a key-ingredient of our analysis and, as we will see in the next paragraph, it provides useful intuitions about the process. On the other hand, when \mathbf{c} is almost uniform, the above *drift* turns out to be negligible and symmetry breaking is due to the inherent variance of the random process.

Failed attempts. When the 3-majority dynamics starts from configurations that exhibit a large initial

support bias between the largest and the second-largest opinions, the approach adopted in [7] successfully exploits the fact that the initial plurality is preserved throughout the evolution of the random process, with an expected positive drift that is also preserved, w.h.p. An intuition of this fact can be achieved from simple manipulations of (2.1). However, the aforementioned drift is only preserved if the largest opinion never changes (w.h.p.), *no matter which the second-largest opinion is*: a condition that is not met by uniform configurations. A promising attempt to cope with uniform configurations is to consider the r.v. $S^{(t)} = C_{\mathbb{M}(t)}^{(t)} - C_{2\mathbb{M}(t)}^{(t)}$ where $\mathbb{M}(t)$ and $2\mathbb{M}(t)$ are the r.v.s that take the index of (one of) the largest opinion and of (one of) the second-largest ones, respectively, in round t . For any *fixed* pair i, j , such that $c_i > c_j$, (2.1) implies that the difference $C_i^{(t+1)} - C_j^{(t+1)}$ in the next round is positive in expectation, so a suitable submartingale argument (similarly to those in [23]), seemed to work in order to show that the system (rather quickly) achieves a “sufficiently-large” bias toward the plurality as to allow fast convergence. This approach would work if the *random* indices \mathbb{M} and $2\mathbb{M}$ maintained their initial values across the entire duration of the process. Unfortunately, starting from uniform configurations, in the next round, the expected difference between the *new* largest opinion and the *new* second largest one may have no positive drift at all. Roughly speaking, in the next round, the r.v. $C_{2\mathbb{M}(t+1)}^{(t+1)}$ can be much larger than the r.v. $C_{\mathbb{M}(t)}^{(t+1)}$.

A promising dynamics for the stabilizing almost-consensus problem is the one introduced in [14], in which nodes revise their opinions (assumed to be totally ordered) by taking the median between the currently held opinion and those held by two randomly sampled nodes. However, while we do not assume opinions to be integers (or totally ordered), their analysis strongly relies on the fact that the median opinion (or any good approximation of it) exhibits a strong increasing drift, even when starting from almost-uniform configuration, whereas no opinion is “special” to a majority rule when the starting configuration is uniform. The adoption of an inherently biased function as the median can have important consequences. To get an intuition, the reader may consider the following simple instance: $\Sigma = \{1, 2, 3\}$, with the system starting in configuration $c_1 = n/2, c_2 = 0, c_3 = n/2$. At the end of the first round, a static adversary changes the values of $F = \log n$ nodes, equally distributed in c_1 and c_3 , to value 2. The (non-valid) value 2 is the *global median* and standard counting arguments show that, while values 1 and 2 have no positive expected drift, the median has a strong expected drift that holds w.h.p. whenever $c_1, c_3 = \Theta(n)$.

This might fool the system into the configuration in which $c_2 = n$, thus converging to a non-valid value.

Our New Approach: An Overview. Our analysis significantly departs from the above approaches. It is important to remark that, for $|\Sigma| \geq 3$, no analysis of the 3-majority dynamics with almost-uniform initial configurations is known, even in the simpler non-adversarial case. On the other hand, while simpler, the analysis of the non-adversarial case still has *per-se* interest and it requires to address some of the main technical challenges that also arise in the adversarial case. Section 3 will be thus devoted to the analysis of the non-adversarial case, while an outline is given in the paragraphs that follow.

When the configuration is (approximately) uniform, Lemma 2.1 tells us that the process exhibits no significant drift toward any *fixed* opinion. Interestingly, things change if we consider the random variable $C_m^{(t)}$, indicating the size of the smallest opinion supported at round t . Let $j \leq k$ be the number of active opinions in a given round t , we first prove that the expected value of $C_m^{(t)}$ always exhibits a non-negligible negative drift:

$$(2.1) \quad \mathbf{E} \left[C_m^{(t+1)} \mid \mathbf{C}^{(t)} = \hat{\mathbf{c}} \right] \leq \hat{c}_m - \varepsilon \frac{\sqrt{n}}{j^{3/2}},$$

for some constant $\varepsilon > 0$. This drift is essentially a consequence of Lemma 2.1 *and* of the standard deviation of r.v.s $C_i^{(t)}$ s (see the proof of Lemma 3.2). The analysis then proceeds along consecutive phases, each consisting of a suitable number of consecutive rounds. If the number of active opinions at the beginning of the generic phase is j , we prove that, with positive constant probability, $C_m^{(t)}$ vanishes within the end of the phase, so that the next phase begins with (at most) $j - 1$ active opinions.

We clearly need a good bound on the length of a phase beginning with at most j opinions. To this aim, we derive a new upper bound - stated in Lemma 3.1 - on the *hitting time* of stochastic processes with expected drift that are defined by finite-state Markov chains [23]. Thanks to this result, we can use the negative drift in (2.1) to prove that, from any configuration with $j \leq k$ active opinions, $C_m^{(t)}$ drops below the threshold $n/j - \sqrt{jn \log n}$ within $\mathcal{O}(\text{poly}(j, \log n))$ rounds, with constant positive probability: This “hitting” event represents the exit condition from the symmetry-breaking stage of the phase. Indeed, once it occurs, we can consider *any fixed* active opinion i having support size c_i below the above threshold (thanks to the previous stage, we know that there is a good chance this opinion exists): We then show that C_i has a negative drift of order $\Omega(c_i/j)$. This allows us to prove that C_i drops from

$n/j - \sqrt{jn \log n}$ to zero within $\mathcal{O}(\text{poly}(j, \log n))$ further rounds, with positive constant probability. This interval of rounds is the dropping stage of the phase.

Ideally, the process proceeds along k consecutive phases, indexed as $j = k, k - 1, \dots, 2$, such that we are left with at most $j - 1$ active opinions at the end of Phase j . In practice, we only have a constant probability that at least one opinion disappears during Phase j . However, using standard probabilistic arguments, we can prove that, w.h.p., for every j , the transition from j to $j - 1$ active opinions takes a constant (amortized) number of phases, each requiring $\mathcal{O}(\text{poly}(j, \log n))$ rounds.

The presence of a dynamic, adaptive adversary makes the above analysis technically more complex. A major issue is that a different definition of *Phase* must be considered, since the adversary might permanently feed any opinion so that the latter never dies. So the number of active opinions might not decrease from one phase to the next one. Essentially, we need to manage the persistence of “small” (valid or not) opinions: The end of a phase is now characterized by one “big” valid color that becomes “small” and, moreover, we need to show that, in general, “small” colors never become “big”, no matter what the dynamic F -bounded adversary does. The dynamic-adversary case is described in Subsection 4.

3 The 3-Majority Dynamics without Adversary

Let $\mathcal{C} \subseteq \Sigma$ be the subset of valid opinions, i.e. those supported by at least one node in the initial configuration, and denote by $k = |\mathcal{C}|$ its size. This section is devoted to the proof of the following result.

THEOREM 3.1. (THE ADVERSARY-FREE CASE.)

Starting from any initial configuration with $k \leq n^{1/3-\varepsilon}$ active opinions, where $\varepsilon > 0$ is an arbitrarily-small constant, the 3-majority dynamics reaches consensus within $\mathcal{O}\left((k^2 \log^{1/2} n + k \log n)(k + \log n)\right)$ rounds, w.h.p.

We first provide the lemmas required for the process analysis and then we give the formal proof of the above theorem.

The next lemma shows an upper bound on the time it takes a stochastic process with values in $N = \{0, 1, \dots, n\}$ to reach or exceed a target value m , under mild hypotheses on the process. We give here only an idea of the full proof, which is deferred to Appendix A.

LEMMA 3.1. *Let $\{X_t\}_t$ be a Markov chain with finite state space Ω , let $f : \Omega \rightarrow N$ be a function mapping states of the chain to non-negative integer numbers, and*

let $\{Y_t\}_t$ be the stochastic process over N defined by $Y_t = f(X_t)$. Let $m \in N$ be a “target value” and let

$$\tau = \inf\{t \in \mathbb{N} : Y_t \geq m\}$$

be the random variable indicating the first time Y_t reaches or exceeds value m . Assume that, for every state $x \in \Omega$ with $f(x) \leq m - 1$, it holds that

1. (Positive drift). $\mathbf{E}[Y_{t+1} | X_t = x] \geq f(x) + \lambda$ for some $\lambda > 0$
2. (Bounded jumps). $\mathbf{P}_x(Y_\tau \geq \alpha m) \leq \alpha m/n$, for some $\alpha > 1$.

Then, for every starting state $x \in \Omega$, it holds that

$$\mathbf{E}_x[\tau] \leq 2\alpha \frac{m}{\lambda}$$

Idea of the proof. From Hypothesis 1 it follows that $Z_t = Y_t - \lambda t$ is a submartingale that satisfies the hypotheses of the Doob’s *Optional Stopping Theorem* [15] (see e.g. Corollary 17.8 in [23] or Theorem 10.10 in [29]), thus

$$0 \leq f(x) = \mathbf{E}_x[Z_0] \leq \mathbf{E}_x[Z_\tau] = \mathbf{E}_x[Y_\tau] - \lambda \mathbf{E}_x[\tau]$$

And from Hypothesis 2 it follows that $\mathbf{E}_x[Y_\tau] \leq 2\alpha m$. \square

We next use the above lemma to bound the time required by the *symmetry-breaking* stage.

LEMMA 3.2. (SYMMETRY-BREAKING STAGE) *Let \mathbf{c} be any configuration with j active opinions. Within $t = \mathcal{O}(j^2 \log^{1/2} n)$ rounds it holds that*

$$\mathbf{P}_{\mathbf{c}}\left(\exists i \text{ such that } C_i^{(t)} \leq n/j - \sqrt{jn \log n}\right) \geq \frac{1}{2}$$

Sketch of Proof. Let J be the set of j active opinions in \mathbf{c} and let $\mathbf{C}^{(t)} = (C_i^{(t)} : i \in J)$ be the random variable indicating the opinion configuration at round t , where we assume $\mathbf{C}^{(0)} = \mathbf{c}$. Let $C_m^{(t)} = \min\{C_i^{(t)} : i \in J\}$ be the minimum among all $C_i^{(t)}$ s and consider the stochastic process $\{Y_t\}_t$ defined as $Y_t = \lfloor n/j \rfloor - C_m^{(t)}$. Observe that Y_t takes values in $\{0, 1, \dots, \lfloor n/j \rfloor\}$ and it is a function of $\mathbf{C}^{(t)}$. We are interested in the first time Y_t becomes at least as large as $\sqrt{jn \log n}$, i.e.

$$\tau = \inf\left\{t \in \mathbb{N} : Y_t \geq \sqrt{jn \log n}\right\}$$

We now show that $\{Y_t\}_t$ satisfies Hypotheses 1 and 2 of Lemma 3.1, with $\lambda = \varepsilon \sqrt{n}/j^{3/2}$, for a suitable constant $\varepsilon > 0$.

1. Let $\hat{\mathbf{c}} = (\hat{c}_i : i \in J)$ be any configuration with j active opinions such that $\hat{c}_m > n/j - \sqrt{jn \log n}$. We want to prove that

$$(3.2) \quad \mathbf{E}\left[C_m^{(t+1)} | \mathbf{C}^{(t)} = \hat{\mathbf{c}}\right] \leq c_m - \varepsilon \frac{\sqrt{n}}{j^{3/2}}$$

Two cases may arise.

Case $\hat{c}_m > n/j - 2\varepsilon \sqrt{n/j}$: Observe that, in this case, r.v.s $\{C_i^{(t+1)} : i \in J\}$ conditional on $\{\mathbf{C}^{(t)} = \hat{\mathbf{c}}\}$ have standard deviation $\Omega(\sqrt{n/j})$. Moreover, they are binomial and negatively associated. Hence, by choosing ε small enough, from the Central Limit Theorem we have that

$$\mathbf{P}\left(i \in J \text{ exists : } C_i^{(t+1)} \leq \frac{n}{j} - 6\varepsilon \cdot \sqrt{\frac{n}{j}}\right) \geq 1/2$$

We thus get

$$(3.3) \quad \begin{aligned} \mathbf{E}\left[C_m^{(t+1)} | \mathbf{C}^{(t)} = \hat{\mathbf{c}}\right] &\leq \\ &\leq \frac{1}{2} \left(\frac{n}{j} - 6\varepsilon \cdot \sqrt{\frac{n}{j}}\right) + \frac{1}{2} \cdot \frac{n}{j} \\ &= \frac{n}{j} - 3\varepsilon \sqrt{\frac{n}{j}} \leq c_m - \varepsilon \sqrt{\frac{n}{j}} \\ &\leq c_m - \varepsilon \frac{\sqrt{n}}{j^{3/2}} \end{aligned}$$

Case $\hat{c}_m \leq n/j - 2\varepsilon \sqrt{n/j}$: Equation (3.2) easily follows from Lemma 2.1. Indeed, let $i \in J$ be an opinion such that $\hat{c}_i = \hat{c}_m$, then

$$(3.4) \quad \begin{aligned} \mathbf{E}\left[C_m^{(t+1)} | \mathbf{C}^{(t)} = \hat{\mathbf{c}}\right] &\leq \mathbf{E}\left[C_i^{(t+1)} | \mathbf{C}^{(t)} = \hat{\mathbf{c}}\right] \\ &\leq \hat{c}_i \left(1 + \frac{\hat{c}_i}{n} - \frac{1}{j}\right) \leq \hat{c}_i \left(1 - \frac{2\varepsilon}{\sqrt{nj}}\right) \\ &\leq \hat{c}_i - \frac{\varepsilon \sqrt{n}}{j^{3/2}} = \hat{c}_m - \varepsilon \frac{\sqrt{n}}{j^{3/2}} \end{aligned}$$

where we used the case’s condition and the fact that $\hat{c}_i = \hat{c}_m \geq n/(2j)$.

2. Since random variables $\{C_i^{(t+1)} : i \in J\}$ are binomial, conditional on the configuration at round t , it is possible to apply Chernoff bound (though with some care) to prove that

$$(3.5) \quad \mathbf{P}_{\mathbf{c}}\left(Y_\tau \geq \alpha \sqrt{jn \log n}\right) \leq \frac{1}{n},$$

for some constant $\alpha > 1$. Though this result seems intuitive, its formal proof is less obvious, since τ is a

stopping time and thus itself a random variable. Lemma B.1 in Appendix B offers a formal proof of the above statement.

From (3.2) and (3.5), we have that $\{Y_t\}_t$ satisfies the hypotheses of Lemma 3.1 with $m = \sqrt{jn \log n}$ and $\lambda = \varepsilon \sqrt{n}/j^{3/2}$. Hence $\mathbf{E}_{\mathbf{c}}[\tau] < j^2 \sqrt{\log n}$ and, from Markov inequality, for $t = 2j^2 \sqrt{\log n}$, we finally get

$$(3.6) \quad \mathbf{P}_{\mathbf{c}} \left(\forall i \in J : C_i^{(t)} \geq n/j - \sqrt{jn \log n} \right) \leq \\ \leq \mathbf{P}_{\mathbf{c}} \left(\tau > 2j^2 \sqrt{\log n} \right) \leq \frac{1}{2}$$

□

We now provide the analysis of the *dropping* stage: More precisely, we show that, if the system starts with up to j active opinions and one of them (say i) is below the threshold $n/j - \sqrt{jn \log n}$, then i drops to the smaller threshold $j^2 \log n$ within $\mathcal{O}(j \log n)$ additional rounds. This bound can be proved w.h.p. since, in this regime, C_i is still sufficiently large to apply the Chernoff bound. This concentration result is not necessary to the purpose of proving Theorem 3.1, while it is a key ingredient in the analysis of the adversarial case (Theorem 4.1). The next lemma can be proved by standard concentration arguments - applied in an iterative way - on the r.v. $C_i^{(t)}$ (see Appendix B).

LEMMA 3.3. (DROPPING STAGE 1) *Let \mathbf{c} be any configuration with $j \leq n^{1/3-\varepsilon}$ active opinions, where $\varepsilon > 0$ is an arbitrarily-small positive constant, and such that an opinion i exists with $c_i \leq n/j - \sqrt{jn \log n}$. Within $t = \mathcal{O}(j \log n)$ rounds opinion i becomes $\mathcal{O}(j^2 \log n)$ w.h.p.*

In the next lemma we prove that once c_i becomes smaller than $n/(2j)$, then opinion i disappears within further $\mathcal{O}(j \log n)$ rounds with constant probability. We only give an outline of the proof (the full proof is presented in Appendix B).

LEMMA 3.4. (DROPPING STAGE 2) *Let \mathbf{c} be any configuration with $j \leq n^{1/3-\varepsilon}$ active opinions, where $\varepsilon > 0$ is an arbitrarily-small positive constant, and such that an opinion i exists with $c_i \leq n/(2j)$. Within $t = \mathcal{O}(j \log n)$ rounds opinion i disappears with probability at least $1/2$.*

Idea of the proof. If $c_i \leq n/(2j)$ in configuration \mathbf{c} , then from Lemma 2.1 it follows that

$$\mathbf{E} \left[C_i^{(t+1)} \mid \mathbf{C}^{(t)} = \mathbf{c} \right] \leq c_i \left(1 - \frac{1}{2j} \right)$$

Moreover, since $C_i^{(t+1)}$ conditional on $\{\mathbf{C}^{(t)} = \mathbf{c}\}$ is binomial, if $j \leq n^{1/3-\varepsilon}$, from the Chernoff bound it

follows that $\mathbf{P} \left(C_i^{(t+1)} > n/(2j) \mid \mathbf{C}^{(t)} = \mathbf{c} \right) \leq e^{-\Theta(n^\varepsilon)}$. Hence, it is easy to check that for any initial configuration \mathbf{c} with $c_i \leq n/(2j)$ the following recursive relation holds

$$\mathbf{E}_{\mathbf{c}} \left[C_i^{(t)} \right] \leq \left(1 - \frac{1}{2j} \right) \mathbf{E}_{\mathbf{c}} \left[C_i^{(t-1)} \right] + e^{-n^{\varepsilon/2}}$$

that for some $t = \mathcal{O}(j \log n)$ gives $\mathbf{E}_{\mathbf{c}} \left[C_i^{(t)} \right] \leq 1/2$.

Since $C_i^{(t)}$ is a non-negative integer-valued r.v., the thesis then follows from the Markov inequality. □

Proof of Theorem 3.1. From Lemmas 3.2, 3.3, and 3.4 it follows that from any configuration with $j \leq k$ active opinions, within $\mathcal{O}(k^2 \sqrt{\log n} + k \log n)$ rounds at least one of the opinions disappears with probability at least $1/4$. Thus, within $\mathcal{O}((k^2 \sqrt{\log n} + k \log n)(k + \log n))$ rounds, all opinions but one disappear w.h.p. □

4 Convergence Time of 3-Majority with Adversary

In this section we consider the presence of a Byzantine adversary that can adaptively change the opinions of a subset of nodes in order to (i) delay the convergence time toward a valid consensus, or (ii) make the system converge to a non-valid opinion. We consider two different kinds of adversaries: A static one and a stronger, dynamic one.

The F -static adversary. Let \mathbf{c} be the initial configuration: At the beginning of the process the adversary looks at \mathbf{c} and can replace the opinions of at most $F = n/k - \sqrt{kn \log n}$ nodes with arbitrary opinions in Σ . Then, the adversary is not allowed to perform any further action during the execution of the protocol. Since any opinion the adversary may introduce has size less than $n/k - \sqrt{kn \log n}$, as a simple consequence of the dropping stage (see Lemmas 3.3 and 3.4), the static adversarial case easily reduces to the non-adversarial one. We thus get the following

COROLLARY 4.1. *Starting from any initial configuration with $k \leq n^\alpha$ active opinions, where $\alpha > 0$ is a suitable constant, the 3-majority dynamics reaches almost-consensus within $\mathcal{O}(k^2 \sqrt{\log n} + k \log n)(k + \log n)$ rounds, in the presence of any F -static adversary with $F = n/k - \sqrt{kn \log n}$, w.h.p.*

The F -dynamic adversary. The actions of this adversary over the studied process can be described as follows.

DEFINITION 4.1. At every round t , after nodes have updated their opinions (i.e. once the configuration $\mathbf{C}^{(t)} = \mathbf{c}^{(t)}$ is realized), the F -dynamic adversary looks at the current configuration and replaces the opinion of up to F nodes with any opinion in Σ . We define $\tilde{\mathbf{C}}^{(t)}$ as the configuration that results from the adversary's action on $\mathbf{c}^{(t)}$ and $D_i^{(t)} = D_i^{(t)}(\mathbf{c}^{(0)}, \tilde{\mathbf{c}}^{(0)}, \dots, \mathbf{c}^{(t-1)}, \tilde{\mathbf{c}}^{(t-1)}, \mathbf{c}^{(t)})$ as the r.v. corresponding to the number of nodes that the adversary adds or removes from c_i (note that $\sum_{i \in \Sigma} |D_i| \leq 2F$) at the end of the t -th round, based on all the past history of the process, i.e.

$$\tilde{\mathbf{C}}^{(t)} = \left(C_1^{(t)} + D_1^{(t)}, \dots, C_{|\Sigma|}^{(t)} + D_{|\Sigma|}^{(t)} \right)$$

In what follows we consider an F -dynamic adversary with $F \leq \beta\sqrt{n}/(k^{\frac{5}{2}} \log n)$ for a suitable positive constant β . As we will show in the proof of Lemma 4.1, this bound on F turns out to be almost tight if the goal is to converge to an almost-consensus regime in polynomial time, w.h.p.

The presence of the adversary requires us to distinguish between valid and non valid opinions. So, we recall that the set of valid opinions $\mathcal{C} \subseteq \Sigma$ is the subset of active opinions in the initial configuration and, in the sequel, we denote k as the number of valid opinions, i.e., $k := |\mathcal{C}|$ and define $\bar{\mathcal{C}} := \Sigma - \mathcal{C}$.

We are now ready to state our main result in the presence of the dynamic adversary.

THEOREM 4.1. (THE DYNAMIC-ADVERSARY CASE.) Starting from any initial configuration with $k \leq n^\alpha$ active opinions, where $\alpha > 0$ is a suitable constant, the 3-majority dynamics reaches (valid) almost-consensus within $\mathcal{O}((k^2\sqrt{\log n} + k \log n)(k + \log n))$ rounds, in the presence of any F -dynamic adversary with $F = \beta\sqrt{n}/(k^{\frac{5}{2}} \log n)$, where $\beta > 0$ is a suitable constant, w.h.p.

In order to prove the above theorem, we need to “improve” the technical lemmas shown in the previous section for the non-adversarial case. Informally speaking, the adversary can introduce “small” non-valid opinions and it can keep small valid opinions active that, we know, they would otherwise disappear. These facts lead us to the problem of managing “small” opinions. In the next subsection, we thus provide analyze this issue and prove the main result above.

4.1 Proof of Theorem 4.1. The rigorous definition of “small opinion” is determined by the minimal negative drift for $C_m^{(t)}$ we derived in the proof of Lemma 3.2 (see (3.4)).

DEFINITION 4.2. Let $S := \{i \mid c_i \leq \gamma\sqrt{n}/k^{\frac{3}{2}}\}$ be the set of the small opinions, where γ is some constant such

that $\gamma > \beta$, and let its complement $B := \bar{S} = \{i \mid c_i > \gamma\sqrt{n}/k^{\frac{3}{2}}\}$ be the set of the big opinions.

It turns out that we cannot define the end of a phase as we did in the non-adversarial case, namely, at least one (valid) opinion dies. Rather and without loss of generality, we assume that all k valid opinions are big when the process begins. Then, our new definition of a phase is the following: phase j is an interval of consecutive rounds, in each of which exactly j big valid opinions are present. Our goal then is to show that at the end of phase j , one of the j initially big colors becomes small and, moreover, this color (and no other small color) will never get big again.

In the symmetry-breaking stage of each phase, we thus need to show that the negative drift of $C_m^{(t)}$ (notice that the latter now denotes the minimum among the j big colors) cannot be opposed by the actions of the F -dynamic adversary, provided that $F \leq \beta\sqrt{n}/(k^{\frac{5}{2}} \log n)$.

LEMMA 4.1. (SYMMETRY-BREAKING STAGE) Let $\tilde{\mathbf{c}}$ be any configuration such that $|B| = j$ and $\sum_{i \in \bar{\mathcal{C}}} \tilde{c}_i \leq \gamma\sqrt{n}/k^{\frac{3}{2}}$. Within $t = \mathcal{O}(j^2 \log^{1/2} n)$ rounds, with probability at least $1/2$ it holds that:

- i) $|B| = j$, $\sum_{i \in \bar{\mathcal{C}}} \tilde{C}_i \leq \gamma\sqrt{n}/k^{\frac{3}{2}}$, and
- ii) $\exists i \in B^{(t)}$ such that $\tilde{C}_i^{(t)} \leq n/j - \sqrt{jn \log n}$.

The formal proof of the above lemma is given in Subsection C.1. Informally, the proof is obtained via two different technical steps:

- i) A new bound on the expected negative drift for $C_m^{(t)}$ that considers both the presence of small good opinions and the adversary's opposing action (for its proof see Subsection C.2).

LEMMA 4.2. Let $\tilde{\mathbf{c}}$ be any configuration such that $|B| \leq j$ and $\sum_{i \in \bar{\mathcal{C}}} \tilde{c}_i^{(t)} \leq \gamma\sqrt{n}/k^{\frac{3}{2}}$. For some constant $\alpha > 0$, for any opinion i such that $\tilde{c}_i \geq \gamma\sqrt{n}/k^{\frac{3}{2}}$, it holds

$$(4.7) \quad \mathbf{E} \left[C_i^{(t+1)} \mid \tilde{\mathbf{C}}^{(t)} = \tilde{\mathbf{c}} \right] \leq \tilde{c}_i \left(1 - \frac{1}{j} + \frac{\tilde{c}_i + \alpha\sqrt{n/k}}{n} \right)$$

$$(4.8) \quad \mathbf{E} \left[\tilde{C}_i^{(t+1)} \mid \tilde{\mathbf{C}}^{(t)} = \tilde{\mathbf{c}} \right] \leq \tilde{c}_i (1 - \eta(i, j)), \text{ where}$$

$$\eta(i, j) = \min \left\{ \frac{1}{j} - \frac{\tilde{c}_i + \alpha\sqrt{n/k}}{n}, \frac{1}{2} \left(\frac{1}{j} - \frac{\tilde{c}_i}{n} \right) \right\}$$

- ii) A novel use of Lemma 3.1 on the hitting time of random processes in order to bound the expect time of the symmetry-breaking stage. We in fact need to define a new stopping condition that also includes some “bad”

event: Some small (valid or not) color become big. More precisely, In Lemma 4.3 (its formal proof can be found in Subsection C.3) we prove the following key-properties of the process in the presence of the dynamic adversary: w.h.p., it is *never* the case that

1. if in a given round a valid opinion is small then it gets big at a later time, i.e. $S^{(t-1)} \subseteq S^{(t)}$;
2. the size of the overall set of non valid opinions grows beyond $\gamma\sqrt{n}/k^{\frac{3}{2}}$, i.e. $\sum_{i \in \bar{C}} c_i \geq \gamma\sqrt{n}/k^{\frac{3}{2}}$.

LEMMA 4.3. *If $\tilde{c}^{(t)}$ is such that $\sum_{i \in \bar{C}} \tilde{c}_i^{(t)} \leq \gamma\sqrt{n}/k^{\frac{3}{2}}$, then $\sum_{i \in \bar{C}} \tilde{C}_i^{(t+1)} \leq \gamma\sqrt{n}/k^{\frac{3}{2}}$ and $S^{(t)} \subseteq S^{(t+1)}$, w.h.p.*

The dropping stage of phase j is now defined as the interval of rounds in which $C_m^{(t)}$ drops from the symmetry-breaking threshold $n/j - \sqrt{jn \log n}$ to the size of small colors i.e. $\gamma\sqrt{n}/k^{\frac{3}{2}}$. Similarly to the non-adversarial case, we can here fix the big opinion i that is dropped below the symmetry-breaking threshold and look at its negative drift derived in Lemma 4.2. The drift turns out to be strong enough to compensate the possible actions of any F -bounded adversary and implies an $O(j \log n)$ bound on the time required by this second stage of phase j . This result is stated in the following Lemma (its proof is given in Subsection C.4).

LEMMA 4.4. (DROPPING STAGE) *Assume that, at round t' , $\tilde{c}^{(t')}$ is such that $\sum_{i \in \bar{C}} c_i \leq \gamma\sqrt{n}/k^{\frac{3}{2}}$, $|B^{(t')}| = j$, and an $i \in B^{(t')}$ exists such that $\gamma\sqrt{n}/k^{\frac{3}{2}} \leq c_i^{(t')} \leq n/j - \sqrt{kn \log n}$. Then, w.h.p., a round $t'' = t' + O(k \log n)$ exists such that $\sum_{i \in \bar{C}} \tilde{C}_i^{(t'')} \leq \gamma\sqrt{n}/k^{\frac{3}{2}}$, $i \in S^{(t'')}$ and $|B^{(t'')}| \leq j - 1$.*

Finally, after k phases, we are left with one (valid) opinion that accounts for $n - O(\sqrt{n})$ nodes, while the remaining nodes can have any (possibly non valid) opinion and reflect the presence of the adversary. In fact, this is what happens with high probability. \square

5 Future Work

We believe that our upper bound on the convergence time of the 3-majority dynamics is not tight w.r.t. k . We think that at least a factor k can be saved. To this aim, we would need to show that “more” opinions get small during a phase. This number should also depend on the current number of big colors. Another idea would be that of (also) considering the growth of the maximal opinion. Unfortunately, differently from the minimal opinion (see (2.1) in Section 2), we have no good bound on the expected drift for the maximal opinion that holds from *any* configuration. So, we don't

see how to efficiently adapt our approach without this crucial ingredient.

A more general open question is to analyze stabilizing almost-consensus dynamics, such as the 3-majority one, in some interesting graph topologies. We believe that a suitable combination of our analysis and some previous analysis for the binary case [11] might result useful in expander graphs [20] and some classes of random evolving graphs [10].

Acknowledgement

We wish to thank Paolo Baldi and Benedetto Scoppola for useful discussions on some probability issues and the anonymous reviewers for competent and constructive comments.

References

- [1] Dana Angluin, James Aspnes, and Eisenstat David. Stably computable predicates are semilinear. In *Proc. of the 25th Ann. ACM SIGACT-SIGOPS Symp. on Principles of Distributed Computing (PODC'06)*, pages 292–299. ACM, 2006.
- [2] Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and Peralta René. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, 18(4):235–253, 2006.
- [3] 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).
- [4] Dana Angluin, Michael J. Fischer, and Hong Jiang. Stabilizing consensus in mobile networks. In *Proc. of Distributed Computing in Sensor Systems (DCOSS'06)*, volume 4026 of LNCS, pages 37–50, 2006.
- [5] James Aspnes. Faster Randomized Consensus with an Oblivious Adversary. In *Proc. of the 31st Ann. ACM SIGACT-SIGOPS Symp. on Principles of Distributed Computing (PODC'12)*, pages 1–8. ACM, 2012.
- [6] Luca Becchetti, Andrea Clementi, Emanuele Natale, Francesco Pasquale, and Riccardo Silvestri. Plurality Consensus in the Gossip Model. In *Proc. of the 26th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'15)*, pages 371–390. SIAM, 2015.
- [7] Luca Becchetti, Andrea Clementi, Emanuele Natale, Francesco Pasquale, Riccardo Silvestri, and Luca Trevisan. Simple dynamics for plurality consensus. In *Proc. of the 26th ACM Symp. on Parallelism in Algorithms and Architectures (SPAA'14)*, pages 247–256. ACM, 2014.
- [8] Ohad Ben-Shahar, Shlomi Dolev, Andrey Dolgin, and Michael Segal. Direction election in flocking swarms. In *Proc. of the 6th Int. Workshop on Foundations of Mobile Computing (DIALM-POMC'10)*, pages 73–80. ACM, 2010.

- [9] Luca Cardelli and Attila Csikász-Nagy. The Cell Cycle Switch Computes Approximate Majority. *Scientific Reports*, Vol. 2, 2012.
- [10] Andrea Clementi, Claudio Macci, Angelo Monti, Francesco Pasquale, and Riccardo Silvestri. Flooding time of edge-markovian evolving graphs. *SIAM J. on Discrete Mathematics*, 24(4):1694–1712, 2010.
- [11] Colin Cooper, Robert Elsasser, and Thomas Radzik. The Power of Two Choices in Distributed Voting. In *Proc. of the 41st Int. Coll. on Automata, Languages, and Programming (ICALP'14)*, volume 8573 of *LNCS*, pages 435–446. Springer, 2014.
- [12] 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 *Proc. of the 6th Ann. ACM Symposium on Principles of Distributed Computing (PODC'12)*, pages 1–12. ACM, 1987.
- [13] Martin Dietzfelbinger, Andreas Goerdts, Michael Mitzenmacher, Andrea Montanari, Rasmus Pagh, and Michael Rink. Tight thresholds for cuckoo hashing via XORSAT. In *Proc. of the 37th Int. Coll. on Automata, Languages, and Programming (ICALP'10)*, volume 6198 of *LNCS*, pages 213–225. Springer, 2010.
- [14] Benjamin Doerr, Leslie A. Goldberg, Lorenz Minder, Thomas Sauerwald, and Christian Scheideler. Stabilizing consensus with the power of two choices. In *Proc. of the 23rd Ann. ACM Symp. on Parallelism in Algorithms and Architectures (SPAA'11)*, pages 149–158. ACM, 2011.
- [15] Joseph L. Doob. *Stochastic Processes*. John Wiley & Sons Inc., 1953.
- [16] David Doty. Timing in chemical reaction networks. In *Proc. of 25th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'14)*, pages 772–784. SIAM, 2014.
- [17] Ofer Feinerman, Bernhard Haeupler, and Amos Korman. Breathe Before Speaking: Efficient Information Dissemination Despite Noisy, Limited and Anonymous Communication. In *Proc. of the ACM Symp. on Principles of Distributed Computing (PODC '14)*, pages 114–123. ACM, 2014.
- [18] Nigel R. Franks, Stephen C. Pratt, Eamonn B. Mallon, Nicholas F. Britton, and David J.T. Sumpter. Information flow, opinion polling and collective intelligence in house-hunting social insects. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 357(1427):1567–1583, 2002.
- [19] Seth Gilbert and Dariusz Kowalski. Distributed agreement with optimal communication complexity. In *Proc. of 21st Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'10)*, pages 965–977. SIAM, 2010.
- [20] S. Hoory, Linial N., and Wigderson W. Expander graphs and their applications. *Bull. Amer. Math. Soc. (N.S.)*, 43:439–561, 2006.
- [21] Richard Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vocking. Randomized rumor spreading. In *Proc. of the 41th Ann. IEEE Symp. on Foundations of Computer Science (FOCS'00)*, pages 565–574. IEEE, 2000.
- [22] David Kempe, Alin Dobra, and Johannes Gehrke. Gossip-Based Computation of Aggregate Information. In *Proc. of 43rd Ann. IEEE Symp. on Foundations of Computer Science (FOCS'03)*, pages 482–491. IEEE, 2003.
- [23] David Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2008.
- [24] 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.
- [25] Marshall Pease, Robert Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *Journal of the ACM*, 27(2):228–234, 1980.
- [26] Etienne Perron, Dinkar Vasudevan, and Milan Vojnovic. Using Three States for Binary Consensus on Complete Graphs. In *Proc. of the 28th IEEE Conf. on Computer Communications (INFOCOM'09)*, pages 2527–1535. IEEE, 2009.
- [27] Michael O. Rabin. Randomized byzantine generals. In *Proc. of the 24th Ann. Symp. on Foundations of Computer Science (SFCS'83)*, pages 403–409. IEEE, 1983.
- [28] Yongxiang Ruan and Yasamin Mostofi. Binary Consensus with Soft Information Processing in Cooperative Networks. In *Proc. of the 47th IEEE Conf. on Decision and Control (CDC'08)*, pages 3613–3619. IEEE, 2008.
- [29] David Williams. *Probability with Martingales*. Cambridge University Press, 1991.
- [30] Mehmet E. Yildiz, Asuman E. Ozdaglar, Daron Acemoglu, Amin Saberi, and Anna Scaglione. Binary Opinion Dynamics with Stubborn Agents. *ACM Trans. Econ. Comput.*, 4(1):Article n. 19, 2013.

Appendix

A Preliminary Results

Proof of Lemma 2.1. According to the 3-majority protocol, a node u gets opinion i if it chooses 3 times opinion i , or if it chooses two times i and one time a different opinion, or if it chooses the first time opinion i and then, the second and third time, two different distinct opinions. Hence, if we denote by $X_{i,u}^{(t)}$ the indicator random variable of the event “Node u gets opinion i at time t ”, we have that

$$\begin{aligned} & \mathbf{P} \left(X_{i,u}^{(t+1)} = 1 \mid \mathbf{C}^{(t)} = \mathbf{c} \right) \\ &= \left(\frac{c_i}{n} \right)^3 + 3 \left(\frac{c_i}{n} \right)^2 \left(\frac{n - c_i}{n} \right) \\ & \quad + \left(\frac{c_i}{n} \right) \left[1 - \left(\frac{\sum_{\ell \in S} c_\ell^2}{n^2} + 2 \left(\frac{c_i}{n} \right) \left(\frac{n - c_i}{n} \right) \right) \right] \\ &= \left(\frac{c_i}{n^3} \right) \left(n^2 + c_i n - \sum_{\ell \in S} c_\ell^2 \right). \end{aligned}$$

Then the inequality in (2.1) is obtained by observing that the sum $\sum_{\ell \in S} c_\ell^2$ is minimized for $c_\ell = n/|S|$. \square

Proof of Lemma 3.1. Consider the stochastic process $Z_t = Y_t - \lambda t$ and observe that for any state $x \in \Omega$ with $f(x) \leq m - 1$ it holds that

$$\begin{aligned} \mathbf{E} [Z_{t+1} \mid X_t = x] &= \mathbf{E} [Y_{t+1} \mid X_t = x] - \lambda(t+1) \\ &\geq f(x) + \lambda - \lambda(t+1) \\ &\geq f(x) - \lambda t, \end{aligned}$$

where in the inequality we used Hypotheses 1. Thus Z_t is a *submartingale* up to the stopping time τ , i.e. $\mathbf{E} [Z_{t+1} \mid X_t] \geq Z_t$ for any $t < \tau$. Moreover, since $|Y_t| \leq n$ the *jumps* of Z_t can be bounded by a value independent of t

$$|Z_{t+1} - Z_t| = |Y_{t+1} - \lambda(t+1) - Y_t + \lambda t| \leq n + \lambda,$$

and it is easy to see that Hypotheses 1 implies $\mathbf{E}_x [\tau] < \infty$, thus we can apply *Doob's Optional Stopping Theorem* [15] (see also, e.g., Corollary 17.8 in [23] and Theorem 10.10 in [29]). It then follows that $\mathbf{E}_x [Z_\tau] \geq \mathbf{E}_x [Z_0] = f(x)$ and, since $\mathbf{E}_x [Z_\tau] = \mathbf{E}_x [Y_\tau] - \lambda \mathbf{E}_x [\tau]$, we have that

$$\mathbf{E}_x [\tau] \leq \frac{\mathbf{E}_x [Y_\tau] - f(x)}{\lambda} \leq \frac{\mathbf{E}_x [Y_\tau]}{\lambda}.$$

Finally, we get

$$\begin{aligned} \mathbf{E}_0 [Y_\tau] &= \sum_{j=1}^n j \mathbf{P}_0 (Y_\tau = j) \\ &= \sum_{j=1}^{\lfloor \alpha m \rfloor} j \mathbf{P}_0 (Y_\tau = j) + \sum_{j=\lfloor \alpha m \rfloor + 1}^n j \mathbf{P}_0 (Y_\tau = j) \\ &\leq (\alpha m) + n \mathbf{P}_0 (Y_\tau > \alpha m) \leq 2(\alpha m), \end{aligned}$$

where in the last inequality we used Hypothesis 2. \square

B Proofs for the Non-Adversarial Case

We first provide a formal proof for (3.5).

LEMMA B.1. *Let \mathbf{c} be any configuration with j active opinions. Consider the stochastic process $\{Y_t\}_t$ defined as $Y_t = \lfloor \frac{n}{j} \rfloor - C_m^{(t)}$ and define the stopping time $\tau = \inf \{t \in \mathbb{N} : Y_t \geq \sqrt{jn \log n}\}$. Then:*

$$\mathbf{P}_{\mathbf{c}} (Y_\tau > \alpha \sqrt{jn \log n}) \leq \frac{1}{n}.$$

Proof. Observe that $\mathbf{E}_{\mathbf{c}} [\tau] < \infty$, since $C_m^{(t)}$ has a negative drift (see the proof of Lemma 3.2). Next, from the definition of Y_t :

$$\begin{aligned} & \mathbf{P}_{\mathbf{c}} (Y_\tau > \alpha \sqrt{jn \log n}) \\ &= \mathbf{P}_{\mathbf{c}} \left(C_m^{(\tau)} < \left\lfloor \frac{n}{j} \right\rfloor - \alpha \sqrt{jn \log n} \right) \\ &= \mathbf{P}_{\mathbf{c}} \left(\exists \ell : C_\ell^{(\tau)} < \left\lfloor \frac{n}{j} \right\rfloor - \alpha \sqrt{jn \log n} \right) \\ &\leq \sum_{\ell=1}^j \mathbf{P}_{\mathbf{c}} \left(C_\ell^{(\tau)} < \left\lfloor \frac{n}{j} \right\rfloor - \alpha \sqrt{jn \log n} \right). \end{aligned}$$

Given any opinion l , any comparison operator $\odot \in \{<, \leq, \geq, >\}$ and any round t , let

$$\mathcal{E}_{l \odot}^{(t)} = C_l^{(t)} \odot \left\lfloor \frac{n}{j} \right\rfloor - \sqrt{jn \log n}.$$

From the definition of the stopping time τ :

$$\begin{aligned}
& \text{(B.1)} \\
& \mathbf{P}_{\mathbf{c}} \left(C_{\ell}^{(\tau)} < \left\lfloor \frac{n}{j} \right\rfloor - \alpha \sqrt{jn \log n} \right) \\
&= \sum_{t=1}^{\infty} \mathbf{P}_{\mathbf{c}} \left(C_{\ell}^{(t)} < \left\lfloor \frac{n}{j} \right\rfloor - \alpha \sqrt{jn \log n} \wedge \tau = t \right) \\
&= \sum_{t=1}^{\infty} \mathbf{P}_{\mathbf{c}} \left(C_{\ell}^{(t)} < \left\lfloor \frac{n}{j} \right\rfloor - \alpha \sqrt{jn \log n} \wedge \mathcal{E}_{\mathbf{m} \leq}^{(t)} \mid \bigwedge_{s=1}^{t-1} \mathcal{E}_{\mathbf{m} >}^{(s)} \right) \\
&\quad \cdot \mathbf{P}_{\mathbf{c}} \left(\bigwedge_{s=1}^{t-1} \mathcal{E}_{\mathbf{m} >}^{(s)} \right) \\
&= \sum_{t=1}^{\infty} \mathbf{P}_{\mathbf{c}} \left(C_{\ell}^{(t)} < \left\lfloor \frac{n}{j} \right\rfloor - \alpha \sqrt{jn \log n} \mid \bigwedge_{s=1}^{t-1} \mathcal{E}_{\mathbf{m} >}^{(s)} \right) \\
&\quad \cdot \mathbf{P}_{\mathbf{c}} \left(\bigwedge_{s=1}^{t-1} \mathcal{E}_{\mathbf{m} >}^{(s)} \right),
\end{aligned}$$

where the last equality follows since

$$C_{\ell}^{(t)} < \left\lfloor \frac{n}{j} \right\rfloor - \alpha \sqrt{jn \log n}$$

implies

$$C_{\mathbf{m}}^{(t)} < \left\lfloor \frac{n}{j} \right\rfloor - \sqrt{jn \log n}.$$

We next consider $\mathbf{P}_{\mathbf{c}} \left(\bigwedge_{s=1}^{t-1} \mathcal{E}_{\mathbf{m} >}^{(s)} \right)$. We can write:

$$\begin{aligned}
\mathbf{P}_{\mathbf{c}} \left(\bigwedge_{s=1}^{t-1} \mathcal{E}_{\mathbf{m} >}^{(s)} \right) &= \prod_{s=1}^{t-1} \mathbf{P}_{\mathbf{c}} \left(\mathcal{E}_{\mathbf{m} >}^{(s)} \mid \bigwedge_{r=1}^{s-1} \mathcal{E}_{\mathbf{m} >}^{(r)} \right) \\
&= \prod_{s=1}^{t-1} \mathbf{P}_{\mathbf{c}} \left(\mathcal{E}_{\mathbf{m} >}^{(s)} \mid \mathcal{E}_{\mathbf{m} >}^{(s-1)} \right),
\end{aligned}$$

where the last equality follows since the 3-majority process is Markovian. We next give an upper bound on $\mathbf{P}_{\mathbf{c}}(C_{\mathbf{m}}^{(s)} > \lfloor n/j \rfloor - \sqrt{jn \log n} \mid C_{\mathbf{m}}^{(s-1)} > \lfloor n/j \rfloor - \sqrt{jn \log n})$:

$$\begin{aligned}
& \mathbf{P}_{\mathbf{c}} \left(\mathcal{E}_{\mathbf{m} >}^{(s)} \mid \mathcal{E}_{\mathbf{m} >}^{(s-1)} \right) \\
&= \sum_{\hat{\mathbf{c}}: \hat{\mathbf{c}}_{\mathbf{m}} > \lfloor \frac{n}{j} \rfloor - \sqrt{jn \log n}} \mathbf{P}_{\hat{\mathbf{c}}} \left(\mathcal{E}_{\mathbf{m} >}^{(1)} \right) \cdot \mathbf{P}_{\mathbf{c}} \left(\mathbf{C}^{(s-1)} = \hat{\mathbf{c}} \mid \mathcal{E}_{\mathbf{m} >}^{(s-1)} \right) \text{(B.4)} \\
&\leq \sum_{\hat{\mathbf{c}}: \hat{\mathbf{c}}_{\mathbf{m}} > \lfloor \frac{n}{j} \rfloor - \sqrt{jn \log n}} \mathbf{P}_{\hat{\mathbf{c}}} \left(\mathcal{E}_{\mathbf{m} >}^{(1)} \right) \cdot \mathbf{P}_{\mathbf{c}} \left(\mathbf{C}^{(s-1)} = \hat{\mathbf{c}} \mid \mathcal{E}_{\mathbf{m} >}^{(s-1)} \right),
\end{aligned}$$

where $l = \arg \hat{\mathbf{c}}_{\mathbf{m}}$ (ties broken arbitrarily). We can give an upper bound on $\mathbf{P}_{\hat{\mathbf{c}}} \left(C_l^{(1)} > \left\lfloor \frac{n}{j} \right\rfloor - \sqrt{jn \log n} \right)$ using a “reverse” Chernoff bound⁴. In particular, it is possible

⁴A folklore example with complete proofs can be found at <http://cstheory.stackexchange.com/questions/14471/reverse-chernoff-bound>.

to show that

$$\mathbf{P}_{\hat{\mathbf{c}}} \left(C_l^{(1)} > (1 - \delta) \mathbf{E}_{\hat{\mathbf{c}}} \left[C_l^{(1)} \right] \right) \leq 1 - e^{-\beta \delta^2 \mathbf{E}_{\hat{\mathbf{c}}} \left[C_l^{(1)} \right]}$$

for a suitable constant β . We use $\delta = \sqrt{jn \log n} / \mathbf{E}_{\hat{\mathbf{c}}} \left[C_l^{(1)} \right]$ and note that $n/2j \leq \mathbf{E}_{\hat{\mathbf{c}}} \left[C_l^{(1)} \right] \leq n/j$, so that

$$\text{(B.2)} \quad \mathbf{P}_{\hat{\mathbf{c}}} \left(C_l^{(1)} > (1 - \delta) \mathbf{E}_{\hat{\mathbf{c}}} \left[C_l^{(1)} \right] \right) \leq 1 - e^{-4\beta j^2 \log n}.$$

Saturating with respect to $\hat{\mathbf{c}}$ yields

$$\mathbf{P}_{\mathbf{c}} \left(\mathcal{E}_{\mathbf{m} >}^{(s)} \mid \mathcal{E}_{\mathbf{m} >}^{(s-1)} \right) \leq 1 - e^{-4\beta j^2 \log n}.$$

On the other hand, using standard concentration techniques and recalling that $\mathbf{E}_{\hat{\mathbf{c}}} \left[C_l^{(1)} \right] \leq n/2j$, we can prove that:

$$\text{(B.3)} \quad \mathbf{P}_{\mathbf{c}} \left(C_{\ell}^{(t)} < \left\lfloor \frac{n}{j} \right\rfloor - \alpha \sqrt{jn \log n} \mid \bigwedge_{s=1}^{t-1} \mathcal{E}_{\mathbf{m} >}^{(s)} \right) \leq e^{-\frac{\alpha^2}{6} j^2 \log n}.$$

Next, substituting (B.2) and (B.3) into (B.1), the result follows after simple calculations and by suitably choosing α . \square

Proof of Lemma 3.3. We first prove that the decreasing rate of C_i depends on its value at the end of the previous round. More formally, if we are in a configuration satisfying the hypotheses of the lemma:

$$\begin{aligned}
& \mathbf{P} \left(C_i^{(t)} > c_i^{(t-1)} \left(1 - \frac{1}{2} \left(\frac{1}{j} - \frac{c_i^{(t-1)}}{n} \right) \right) \right) \\
&= \mathbf{P} \left(C_i^{(t)} > c_i^{(t-1)} \left(1 - \left(\frac{1}{j} - \frac{c_i^{(t-1)}}{n} \right) \right) (1 + \delta) \right),
\end{aligned}$$

where $\delta = \frac{1}{2} \left(\frac{1}{j} - \frac{c_i^{(t-1)}}{n} \right) / 1 - \left(\frac{1}{j} - \frac{c_i^{(t-1)}}{n} \right)$

Using Lemma 2.1 and applying Chernoff bound we have:

$$\begin{aligned}
& \mathbf{P} \left(C_i^{(t)} > c_i^{(t-1)} \left(1 - \frac{1}{2} \left(\frac{1}{j} - \frac{c_i^{(t-1)}}{n} \right) \right) \right) \\
&\leq \exp \left\{ -\frac{\delta^2}{3} \left(1 - \left(\frac{1}{j} - \frac{c_i^{(t-1)}}{n} \right) \right) c_i^{(t-1)} \right\} \\
&= \exp \left\{ -\frac{\delta}{3} \left(\frac{1}{2} \left(\frac{1}{j} - \frac{c_i^{(t-1)}}{n} \right) \right) c_i^{(t-1)} \right\} \\
&< \exp \left\{ -\frac{1}{3} \left(\frac{1}{2} \left(\frac{1}{j} - \frac{c_i^{(t-1)}}{n} \right) \right)^2 c_i^{(t-1)} \right\} \\
&\text{(B.5)} \quad = n^{-\Theta(1)}.
\end{aligned}$$

The second equality in (B.5) follows from the definition of δ , while the third inequality follows by (upper) bounding the denominator of δ by 1, which is always possible since $c_i/n - 1/j < 0$ from the hypotheses. Finally, to prove the last equality, we used the fact that $c_i \geq j^2 \log n$ and that the function $x(1-x)^2$ is decreasing iff $x \in (1/3, 1)$, with $x = jc_i/n$.

Finally, we can iteratively apply (B.5) as long as we have at most j active opinions and $C_i^{(t)}$ keeps not smaller than $j^2 \log n$. By standard concentration arguments we get that the time to reach this threshold is $\mathcal{O}(j \log n)$, w.h.p. \square

Proof of Lemma 3.4. Let J be the set of active opinions. By conditioning on all the configurations $\hat{\mathbf{c}} = (\hat{c}_\ell : \ell \in J)$ that the system can take at round $t-1$, we can bound the expectation of $C_i^{(t)}$ as follows

$$\begin{aligned} & \mathbf{E}_{\mathbf{c}} \left[C_i^{(t)} \right] \\ &= \sum_{\hat{\mathbf{c}}} \mathbf{E} \left[C_i^{(t)} \mid \mathbf{C}^{(t-1)} = \hat{\mathbf{c}} \right] \mathbf{P}_{\mathbf{c}} \left(\mathbf{C}^{(t-1)} = \hat{\mathbf{c}} \right) \\ &\leq \left(1 - \frac{1}{2j} \right) \sum_{\hat{\mathbf{c}}: \hat{c}_i \leq n/(2j)} \hat{c}_i \cdot \mathbf{P}_{\mathbf{c}} \left(\mathbf{C}^{(t-1)} = \hat{\mathbf{c}} \right) \\ &\quad + n \cdot \sum_{\hat{\mathbf{c}}: \hat{c}_i > n/(2j)} \mathbf{P}_{\mathbf{c}} \left(\mathbf{C}^{(t-1)} = \hat{\mathbf{c}} \right) \\ &\leq \left(1 - \frac{1}{2j} \right) \mathbf{E}_{\mathbf{c}} \left[C_i^{(t-1)} \right] + n \cdot \mathbf{P}_{\mathbf{c}} \left(C_i^{t-1} > \frac{n}{2j} \right), \end{aligned}$$

where we used that, for any configuration $\hat{\mathbf{c}}$ with $\hat{c}_i \leq n/(2j)$, Lemma 2.1 gives the bound $\mathbf{E} \left[C_i^{(t)} \mid \mathbf{C}^{(t-1)} = \hat{\mathbf{c}} \right] \leq \hat{c}_i \left(1 - \frac{1}{2j} \right)$. Moreover, if $j \leq n^{1/3-\varepsilon}$, from Chernoff bound it follows that

$$\mathbf{P} \left(C_i^{(t)} > \frac{n}{2j} \mid \mathbf{C}^{(t-1)} = \hat{\mathbf{c}} \right) \leq e^{-\Theta(n^\varepsilon)}$$

for any such configuration $\hat{\mathbf{c}}$. Hence, for any t we have that $\mathbf{P}_{\mathbf{c}} \left(C_i^{(t)} > \frac{n}{2j} \right) \leq te^{-\Theta(n^\varepsilon)}$. Indeed,

$$\begin{aligned} & \mathbf{P}_{\mathbf{c}} \left(C_i^{(t)} > \frac{n}{2j} \right) \\ &\leq \mathbf{P}_{\mathbf{c}} \left(\exists \bar{t} = 1, \dots, t : C_i^{(\bar{t})} > \frac{n}{2j} \wedge C_i^{(\bar{t}-1)} \leq \frac{n}{2j} \right) \\ &\leq \sum_{\bar{t}=1}^t \mathbf{P}_{\mathbf{c}} \left(C_i^{(\bar{t})} > \frac{n}{2j} \wedge C_i^{(\bar{t}-1)} \leq \frac{n}{2j} \right) \\ &= \sum_{\bar{t}=1}^t \sum_{\hat{\mathbf{c}}: \hat{c}_i \leq \frac{n}{2j}} \mathbf{P} \left(C_i^{(\bar{t})} > \frac{n}{2j} \mid \mathbf{C}^{(\bar{t}-1)} = \hat{\mathbf{c}} \right) \mathbf{P}_{\mathbf{c}} \left(\mathbf{C}^{(\bar{t}-1)} = \hat{\mathbf{c}} \right) \\ &\leq te^{-\Theta(n^\varepsilon)}. \end{aligned}$$

Thus for any $t = \text{poly}(n)$ the following recursive relation holds

$$\mathbf{E}_{\mathbf{c}} \left[C_i^{(t)} \right] \leq \left(1 - \frac{1}{2j} \right) \mathbf{E}_{\mathbf{c}} \left[C_i^{(t-1)} \right] + e^{-n^{\varepsilon/2}},$$

and it gives

$$\mathbf{E}_{\mathbf{c}} \left[C_i^{(t)} \right] \leq \left(1 - \frac{1}{2j} \right)^t \frac{n}{2j} + e^{-n^{\varepsilon/3}}.$$

Hence, for $t = 2j(\log n + 1)$ we have that $\mathbf{E}_{\mathbf{c}} \left[C_i^{(t)} \right] \leq 1/2$ and since $C_i^{(t)}$ takes non-negative integer values, the thesis follows from Markov inequality. \square

C Omitted Proofs for the Adversarial Case

The proofs for the adversarial case need the following preliminary results.

LEMMA C.1. *Let $\tilde{\mathbf{c}}$ be any configuration such that $|B| = j$ and $\sum_{i \in \tilde{\mathcal{C}}} \tilde{c}_i \leq \gamma \sqrt{n}/k^{\frac{3}{2}}$. Consider the stochastic process $\{\tilde{Y}_t\}_t$ defined as $\tilde{Y}_t = \lfloor \frac{n}{j} \rfloor - \tilde{C}_m^{(t)}$ and define the stopping time*

$$\begin{aligned} \tau = \inf \{ t \in \mathbb{N} : \tilde{Y}_t \geq \sqrt{jn \log n} \vee \left(\sum_{i \in \tilde{\mathcal{C}}} \tilde{C}_i \geq \gamma \sqrt{n}/k^{\frac{3}{2}} \right) \right. \\ \left. \vee (S^{(t-1)} \not\subseteq S^{(t)}) \right\}. \end{aligned}$$

Then, it holds that

$$\mathbf{P}_{\mathbf{c}} \left(\tilde{Y}_\tau > \alpha \sqrt{jn \log n} \right) \leq \frac{1}{n}.$$

Sketch of Proof. The proof of this Lemma follows from minor modifications of the proof of Lemma B.1. In particular, the argument is based on the following observations:

1. The event defining the stopping time τ is in this case

$$\begin{aligned} \mathcal{E}^{(t)} = \left(\tilde{Y}_t \geq (\sqrt{jn \log n}) \vee \left(\sum_{i \in \tilde{\mathcal{C}}} \tilde{C}_i \geq \gamma \sqrt{n}/k^{\frac{3}{2}} \right) \right. \\ \left. \vee (S^{(t-1)} \not\subseteq S^{(t)}) \right). \end{aligned}$$

The negated of this event is

$$\begin{aligned} -\mathcal{E}^{(t)} = \left(\tilde{Y}_t \leq (\sqrt{jn \log n}) \wedge \left(\sum_{i \in \tilde{\mathcal{C}}} \tilde{C}_i \leq \gamma \sqrt{n}/k^{\frac{3}{2}} \right) \right. \\ \left. \wedge (S^{(t-1)} \subseteq S^{(t)}) \right), \end{aligned}$$

which implies the event $\left(\tilde{Y}_t \leq \sqrt{jn \log n} \right)$.

2. Proceeding like in the proof of Lemma B.1, we can write an expression that is similar to (B.1), with the generic conditioning event

$$C_m^{(s)} > \left\lfloor \frac{n}{j} \right\rfloor - \sqrt{jn \log n},$$

replaced by $-\mathcal{E}^{(s)}$. The conditioned event

$$C_\ell^{(t)} < \left\lfloor \frac{n}{j} \right\rfloor - \alpha\sqrt{jn \log n},$$

is instead replaced by the event

$$\left(C_\ell^{(t)} < \left\lfloor \frac{n}{j} \right\rfloor - \alpha\sqrt{jn \log n} \right) \wedge \mathcal{E}^{(t)}.$$

Now, note that the event

$$C_\ell^{(t)} < \left\lfloor \frac{n}{j} \right\rfloor - \alpha\sqrt{jn \log n},$$

again implies $\mathcal{E}^{(t)}$. Hence, we can still write (B.1), from which the proof requires some non-hard adaptations w.r.t. the case without adversary. \square

Since the adversary, at round t , may decide what to do based on the full history of the process up to time t , the stochastic process $\{\tilde{C}^{(t)}\}_t$ may not be a Markov process anymore. Thus, we need a more general version of Lemma 3.1.

LEMMA C.2. *Let $\{X_t\}_t$ be a discrete time stochastic process with a finite state space Ω , let $f_t : \Omega^t \rightarrow N$ be a function mapping histories of the process in non-negative integer numbers, and let $\{Y_t\}_t$ be the stochastic process over N defined by $Y_t = f_t(X_0, \dots, X_t)$. Let $m \in N$ be a “target value”, let $A \subseteq \Omega$ be an arbitrary subset of states, and let*

$$\tau = \inf\{t \in \mathbb{N} : Y_t \geq m \text{ or } X_t \notin A\}$$

be the random variable indicating the first time X_t exits from set A or Y_t reaches or exceeds value m . Assume that, for every sequence of states $x_0, \dots, x_t \in A$ with $f_t(x_0, \dots, x_t) \leq m - 1$, it holds that

1. (Positive drift). $\mathbf{E}[Y_{t+1} | X_0 = x_0, \dots, X_t = x_t] \geq f_t(x_0, \dots, x_t) + \lambda$ for some $\lambda > 0$
2. (Bounded jumps). $\mathbf{P}_x(Y_\tau \geq \alpha m) \leq \alpha m/n$, for some $\alpha > 1$.

Then, for every starting state $x \in A$, it holds that

$$\mathbf{E}_x[\tau] \leq 2\alpha \frac{m}{\lambda}$$

Proof. The proof is a straight adaptation of the proof of Lemma 3.1, in which we take into account the full history of the process.

Consider the stochastic process $Z_t = Y_t - \lambda t$. For any sequence of states $x_0, \dots, x_t \in A$ with

$f_t(x_0, \dots, x_t) \leq m - 1$ it holds that

$$\begin{aligned} \mathbf{E}[Z_{t+1} | X_0 = x_0, \dots, X_t = x_t] &= \mathbf{E}[Y_{t+1} | X_0 = x_0, \dots, X_t = x_t] - \lambda(t+1) \\ &\geq f_t(x_0, \dots, x_t) + \lambda - \lambda(t+1) \\ &\geq f_t(x_0, \dots, x_t) - \lambda t, \end{aligned}$$

where in the inequality we used Hypotheses 1. Thus, Z_t is a *submartingale* up to the stopping time τ . Moreover, since $|Y_t| \leq n$ then $|Z_{t+1} - Z_t| \leq n + \lambda$ and, together with Hypotheses 1 this implies $\mathbf{E}_x[\tau] < \infty$. Thus, we can apply *Doob's Optional Stopping Theorem* [15]. It follows that $\mathbf{E}_x[Z_\tau] \geq \mathbf{E}_x[Z_0] = f_0(x)$ and, since $\mathbf{E}_x[Z_\tau] = \mathbf{E}_x[Y_\tau] - \lambda \mathbf{E}_x[\tau]$, we have that

$$\mathbf{E}_x[\tau] \leq \frac{\mathbf{E}_x[Y_\tau] - f_0(x)}{\lambda} \leq \frac{\mathbf{E}_x[Y_\tau]}{\lambda}$$

Finally, we get

$$\begin{aligned} \mathbf{E}_0[Y_\tau] &= \sum_{j=1}^{\lfloor \alpha m \rfloor} j \mathbf{P}_0(Y_\tau = j) + \sum_{j=\lfloor \alpha m \rfloor + 1}^n j \mathbf{P}_0(Y_\tau = j) \\ &\leq (\alpha m) + n \mathbf{P}_0(Y_\tau > \alpha m) \leq 2(\alpha m), \end{aligned}$$

where in the last inequality we used Hypothesis 2. \square

C.1 Proof of Lemma 4.1. We proceed by adapting the proof of Lemma 3.2. Let $\tilde{C}^{(0)} = \tilde{c}$ be the initial configuration. Let us consider the stochastic process $\{\tilde{Y}_t\}_{t \geq 0}$ defined as

$$\tilde{Y}_t = \left\lfloor \frac{n}{j} \right\rfloor - \tilde{C}_m^{(t)}$$

where $\tilde{C}_m^{(t)} = \min\{\tilde{C}_i^{(t)} : i \in B^{(t)}\}$. We are interested in round

$$\begin{aligned} \tau &= \inf\{t \in \mathbb{N} : \tilde{Y}_t \geq (\sqrt{jn \log n}) \vee \left(\sum_{i \in \tilde{C}} \tilde{C}_i \geq \gamma \sqrt{n}/k^{\frac{3}{2}} \right) \vee \right. \\ &\quad \left. \vee (S^{(t-1)} \not\subseteq S^{(t)})\} \end{aligned}$$

Now we show that $\{\tilde{Y}_t\}_t$ satisfies the Hypotheses 1 and 2 of Lemma 3.1 with $A = \left(\sum_{i \in \tilde{C}} \tilde{C}_i \leq \gamma \sqrt{n}/k^{\frac{3}{2}} \right) \vee (S^{(t-1)} \subseteq S^{(t)})$ and $\lambda = \varepsilon \sqrt{n}/j^{3/2}$, for a suitable constant $\varepsilon > \alpha$.

1. Let \tilde{c} be any configuration such that $\tilde{c}_m > n/j - \sqrt{jn \log n}$. Now we prove that

$$(C.6) \quad \mathbf{E} \left[\tilde{C}_m^{(t+1)} | \tilde{C}^{(t)} = \tilde{c} \right] \leq \tilde{c}_m - \varepsilon \frac{\sqrt{n}}{j^{3/2}}$$

Case $\tilde{c}_m > n/j - 2\varepsilon \sqrt{n}/j$: Observe that, in this case, random variables $\{\tilde{C}_i^{(t+1)} : i \in B\}$ have standard deviation is $\Omega(\sqrt{n}/j)$. Moreover they are binomial and negatively associated. Hence, by choosing ε small enough,

from the Central Limit Theorem we have that

$$\mathbf{P} \left(\exists i \in B \text{ such that } C_i^{(t+1)} \leq \frac{n}{j} - 6\varepsilon \cdot \sqrt{\frac{n}{j}} \right) \geq 1/2$$

We thus get

$$\begin{aligned} \mathbf{E} \left[\tilde{C}_m^{(t+1)} \mid \tilde{\mathbf{C}}^{(t)} = \tilde{\mathbf{c}} \right] &\leq \frac{1}{2} \left(\frac{n}{j} - 6\varepsilon \cdot \sqrt{\frac{n}{j}} \right) + \frac{1}{2} \cdot \frac{n}{j} + \frac{\beta\sqrt{n}}{k^{\frac{5}{2}} \log n} \\ \text{(C.7)} \quad &= \frac{n}{j} - 2\varepsilon \sqrt{\frac{n}{j}} + \frac{\beta\sqrt{n}}{k^{\frac{5}{2}} \log n} \leq \tilde{c}_m - \varepsilon \sqrt{\frac{n}{j}} \leq \tilde{c}_m - \varepsilon \frac{\sqrt{n}}{j^{3/2}} \end{aligned}$$

Case $\tilde{c}_m \leq n/j - 2\varepsilon \sqrt{n/j}$: Equation (C.6) easily follows from Lemma 4.2. Indeed, let $i \in B$ be an opinion such that $\hat{c}_i = \hat{c}_m$, then

$$\begin{aligned} \mathbf{E} \left[\tilde{C}_m^{(t+1)} \mid \tilde{\mathbf{C}}^{(t)} = \tilde{\mathbf{c}} \right] &\leq \mathbf{E} \left[\tilde{C}_i^{(t+1)} \mid \tilde{\mathbf{C}}^{(t)} = \tilde{\mathbf{c}} \right] \leq \tilde{c}_i \left(1 + \frac{\tilde{c}_i + \alpha\sqrt{n/k}}{n} - \frac{1}{j} \right) \\ &\leq \tilde{c}_i \left(1 - \frac{2\varepsilon}{\sqrt{nj}} + \frac{\alpha}{\sqrt{kn}} \right) \leq \tilde{c}_i - \frac{\varepsilon\sqrt{n}}{j^{3/2}} = \tilde{c}_m - \varepsilon \frac{\sqrt{n}}{j^{3/2}}, \end{aligned}$$

where we used the case's condition and $\tilde{c}_i = \tilde{c}_m \geq n/(2j)$.

2. Since random variables $\{\tilde{C}_i^{(t)} : i \in B^{(t)}\}$ are binomial conditional on the configuration at round $t-1$, from the Chernoff bound it follows that

$$\mathbf{P}_{\tilde{\mathbf{c}}} \left(\tilde{Y}_\tau \geq \alpha \sqrt{jn \log n} \right) \leq \frac{1}{n}, \text{ for some constant } \alpha > 1$$

See Lemma C.1 for the formal statement of the last fact.

From (C.6) and (C.8) we have that $\{\tilde{Y}_t\}_t$ satisfies the hypotheses of Lemma C.2 with $m = \sqrt{jn \log n}$, $\lambda = \varepsilon\sqrt{n}/j^{3/2}$ and $A = (\sum_{i \in \tilde{\mathcal{C}}} \tilde{C}_i \leq \gamma\sqrt{n}/k^{\frac{3}{2}}) \vee (S^{(t-1)} \subseteq S^{(t)})$. Moreover, by iteratively applying Lemma 4.3, we have that, for any $t = \mathcal{O}(n^2)$, it holds w.h.p. that $(\sum_{i \in \tilde{\mathcal{C}}} \tilde{C}_i^{(t)} \leq \gamma\sqrt{n}/k^{\frac{3}{2}}) \vee (S^{(t-1)} \subseteq S^{(t)})$. Thus, from Markov's inequality, for $t = 2j^2\sqrt{\log n}$, we have that

$$\begin{aligned} \mathbf{P}_{\tilde{\mathbf{c}}} (\forall i \in B : (C_i^{(t)} \leq n/j - \sqrt{jn \log n}) \wedge (\sum_{i \in \tilde{\mathcal{C}}} \tilde{C}_i^{(t)} \leq \gamma\sqrt{n}/k^{\frac{3}{2}}) \\ \wedge (S^{(0)} \subseteq S^{(t)})) \geq \mathbf{P}_{\tilde{\mathbf{c}}} \left(\hat{\tau} \leq 2j^2\sqrt{\log n} \right) \geq \frac{1}{3} \end{aligned}$$

where $\hat{\tau} = \inf\{t \in \mathbb{N} : \tilde{Y}_t \geq \sqrt{jn \log n}\}$. \square

C.2 Proof of Lemma 4.2. Similarly to the proof of Lemma 2.1 we have

$$\begin{aligned} \mathbf{E} \left[C_i^{(t+1)} \mid \tilde{\mathbf{C}}^{(t)} = \tilde{\mathbf{c}} \right] &\leq \tilde{c}_i \left(1 + \frac{\tilde{c}_i}{n} - \frac{\sum_j \tilde{c}_j^2}{n^2} \right) \leq \tilde{c}_i \left(1 + \frac{\tilde{c}_i}{n} - \frac{\sum_{j \in B} \tilde{c}_j^2}{n^2} \right) \\ &\leq \tilde{c}_i \left(1 + \frac{\tilde{c}_i}{n} - \frac{\sum_{j \in B} \left(\frac{n - (k-j+1)\gamma\sqrt{n/k}^{\frac{3}{2}}}{j} \right)^2}{n^2} \right) \\ &\leq \tilde{c}_i \left(1 + \frac{\tilde{c}_i}{n} - \frac{\sum_{j \in B} (n - \alpha/4\sqrt{n/k})^2}{j^2 n^2} \right) \\ &\leq \tilde{c}_i \left(1 + \frac{\tilde{c}_i}{n} - \frac{1}{j} + \frac{\alpha/2\sqrt{n/k}}{jn} \right) \\ &\leq \tilde{c}_i \left(1 - \frac{n/j - \tilde{c}_i - \alpha/2\sqrt{n/k}}{n} \right). \end{aligned}$$

Taking into account any possible action of the adversary, we thus get that

(C.9)

$$\begin{aligned} \mathbf{E} \left[\tilde{C}_i^{(t+1)} \mid \tilde{\mathbf{C}}^{(t)} = \tilde{\mathbf{c}} \right] &= \mathbf{E} \left[C_i^{(t+1)} \mid \tilde{\mathbf{C}}^{(t)} = \tilde{\mathbf{c}} \right] + \mathbf{E} \left[D_i^{(t+1)} \mid \tilde{\mathbf{C}}^{(t)} = \tilde{\mathbf{c}} \right] \\ &\leq \tilde{c}_i \left(1 - \frac{n/j - \tilde{c}_i - \alpha/2\sqrt{n/k}}{n} \right) + F \\ \text{(C.10)} \quad &\leq \tilde{c}_i \left(1 - \frac{n/j - \tilde{c}_i}{n} + \frac{2 \max \left\{ \alpha/2\sqrt{n/k}, Fn/\tilde{c}_i \right\}}{n} \right). \end{aligned}$$

By distinguishing the cases $\tilde{c}_i \geq n/(3j)$ or $\tilde{c}_i < n/(3j)$, from (C.10) we get (4.8). \square

C.3 Proof of Lemma 4.3. From Lemma 4.2, for each $i \in S^{(t)}$ we have that

$$\mathbf{E} \left[C_i^{(t+1)} \mid \tilde{\mathbf{C}}^{(t)} = \tilde{\mathbf{c}} \right] \leq \tilde{c}_i \left(1 + \frac{\tilde{c}_i}{n} - \frac{1}{k} \right).$$

From a direct application of the Chernoff bound to $C_i^{(t+1)}$, and taking into account any possible action of the adversary, we thus get that w.h.p.

$$\tilde{C}_i^{(t+1)} = C_i^{(t+1)} + D_i^{(t+1)} \leq \gamma \frac{\sqrt{n}}{k^{\frac{3}{2}}} \left(1 - \frac{1}{4k} \right) + F \leq \gamma \frac{\sqrt{n}}{k^{\frac{3}{2}}},$$

that is, $i \in S^{(t)}$ w.h.p. Analogously, we have

$$\begin{aligned} & \mathbf{E} \left[\sum_{i \in \bar{C}^{(t)}} C_i^{(t+1)} \mid \tilde{C}^{(t)} = \tilde{\mathbf{c}} \right] \\ & \leq \sum_{i \in \bar{C}^{(t)}} \tilde{c}_i^{(t)} \left(1 + \frac{\tilde{c}_i}{n} - \frac{1}{k} \right) \leq \gamma \frac{\sqrt{n}}{k^{\frac{3}{2}}} \left(1 - \frac{1}{2k} \right), \end{aligned}$$

and then, by applying the Chernoff bound, we get that w.h.p.

$$\begin{aligned} \sum_{i \in \bar{C}^{(t)}} \tilde{C}_i^{(t+1)} &= \sum_{i \in \bar{C}^{(t)}} C_i^{(t+1)} + \sum_i D_i^{(t+1)} \\ &\leq \gamma \frac{\sqrt{n}}{k^{\frac{3}{2}}} \left(1 - \frac{1}{4k} \right) + F \leq \gamma \frac{\sqrt{n}}{k^{\frac{3}{2}}}, \end{aligned}$$

concluding the proof. \square

C.4 Proof of Lemma 4.4. By iteratively applying Lemma 4.3, we have that, w.h.p., for each $t \in \{t', \dots, t'' - 1\}$ it holds $\sum_{i \in \bar{C}} \tilde{C}_i^{(t)} \leq \gamma \sqrt{n}/k^{\frac{3}{2}}$ and $i \in S^{(t)}$.

To prove that $|B^{(t'')}| \leq j - 1$, we first prove that, for each round $t \in \{t' + 1, \dots, t''\}$, w.h.p. $\tilde{C}_i^{(t)}$ decreases by a certain extent that depends on $\tilde{c}_i^{(t-1)}$, regardless of what the adversary does.

Let $\psi = (1/j - (\tilde{c}_i^{(t-1)} + \alpha\sqrt{n/k})/n)$. If we are in a configuration satisfying the hypotheses of the lemma, we have

$$\begin{aligned} \text{(C.11)} \quad \mathbf{P} \left(C_i^{(t)} > \tilde{c}_i^{(t-1)} \left(1 - \frac{\psi}{2} \right) \right) &= \\ &= \mathbf{P} \left(C_i^{(t)} > \tilde{c}_i^{(t-1)} (1 - \psi(1 + \delta)) \right), \end{aligned}$$

where $\delta = \frac{1}{2}\psi/(1 - \psi)$. Thus, using Lemma 4.2 and applying the Chernoff bound we have

$$\begin{aligned} \mathbf{P} \left(C_i^{(t)} > \tilde{c}_i^{(t-1)} \left(1 - \frac{\psi}{2} \right) \right) &\leq \exp \left\{ -\frac{\delta^2}{3} \psi \tilde{c}_i^{(t-1)} \right\} \\ \text{(C.12)} \quad &< \exp \left\{ -\frac{1}{3} \left(\frac{1}{2} \psi \right)^2 \tilde{c}_i^{(t-1)} \right\} = n^{-\Theta(1)}, \end{aligned}$$

where the second inequality follows from the definition of δ and the fact that its denominator is smaller than 1, and the equality in (C.12) follows by minimizing $\psi^2 \tilde{c}_i^{(t-1)}$ for $\gamma \sqrt{n}/k^{\frac{3}{2}} \leq c_i^{(t)} \leq n/j - \sqrt{kn \log n}$.

It follows that, w.h.p.

$$\text{(C.13)} \quad \tilde{C}_i^{(t)} = C_i^{(t)} + D_i^{(t)} \leq \tilde{c}_i^{(t-1)} \left(1 - \frac{\psi}{2} \right) + F \leq \tilde{c}_i^{(t-1)}$$

Thus, w.h.p., we can iteratively apply (C.13) until $\tilde{c}_i^{(t-1)} \leq \gamma \sqrt{n}/k^{\frac{3}{2}}$. We next prove that this happens

within $\mathcal{O}(k \log n)$ rounds, w.h.p. Interestingly, showing that, within $\mathcal{O}(k \log n)$ rounds, C_i decreases to a constant fraction of its value at the beginning of the dropping stage does not seem obvious. For this reason, we consider the evolution of the displacement $\frac{n}{j} - C_i$, which seems analytically more tractable. To this purpose, note that (C.12) implies that, w.h.p.

$$\begin{aligned} \text{(C.14)} \quad \frac{n}{j} - C_i^{(t)} &\geq \\ &\geq \frac{n}{j} - c_i^{(t-1)} + \frac{c_i^{(t-1)}}{2} \left(\frac{1}{j} - \frac{c_i^{(t-1)} + \alpha\sqrt{n/k}}{n} \right) \\ &= \frac{n}{j} - c_i^{(t-1)} + \frac{c_i^{(t-1)}}{2} \left(\frac{1}{j} - \frac{c_i^{(t-1)}}{n} \right) \left(1 - \frac{\alpha\sqrt{n/k}}{\frac{1}{j} - \frac{c_i^{(t-1)}}{n}} \right) \\ &= \frac{n}{j} - c_i^{(t-1)} + \frac{c_i^{(t-1)}}{2} \left(\frac{1}{j} - \frac{c_i^{(t-1)}}{n} \right) \left(1 + \frac{\alpha}{\log n} \right) \\ &= \left(\frac{n}{j} - c_i^{(t-1)} \right) \left(1 + \alpha_1 \frac{c_i^{(t-1)}}{2n} \right), \end{aligned}$$

for some constant $\alpha_1 > 0$, where in the first equality of (C.14) we have used that $n/j - c_i^{(t-1)} \geq \sqrt{kn \log n}$.

We can now conclude the proof of Lemma 4.4. We first prove that $C_i \leq n/(2j)$ within $\mathcal{O}(k \log n)$ steps, w.h.p. To this purpose, note that $\frac{n}{j} - c_i \geq \sqrt{kn \log n}$ at the beginning of the dropping stage from the hypotheses. Furthermore, for some positive constants α_2 and α_3 , as long as $C_i \geq \alpha_3 n/j$, it holds $1 + \alpha_1 c_i/n \geq 1 + \alpha_2/j$. Hence, after $\mathcal{O}(k \log n)$ steps, w.h.p. we have $\frac{n}{j} - c_i \geq (1 - \alpha_3) \frac{n}{j}$, which in turn implies $c_i \leq \alpha_3 n/j$. Once $c_i \leq \alpha_3 n/j$, using again (C.14) we have that C_i decreases by a factor $1 - \Omega(1/j)$ in every round w.h.p. By standard concentration arguments we obtain that eventually $c_i \leq \gamma \sqrt{n}/k^{\frac{3}{2}}$ within $\mathcal{O}(k \log n)$ more steps, w.h.p. \square