



On the Limits of Information Spread by Memory-Less Agents

Niccolò D’Archivio  

COATI, INRIA d’Université Côte d’Azur, Sophia-Antipolis, France

Robin Vacus  

Bocconi Institute for Data Science and Analytics, Bocconi University, Milan, Italy

Abstract

We address the self-stabilizing bit-dissemination problem, designed to capture the challenges of spreading information and reaching consensus among entities with minimal cognitive and communication capacities. Specifically, a group of n agents is required to adopt the correct opinion, initially held by a single informed individual, choosing from two possible opinions. In order to make decisions, agents are restricted to observing the opinions of a few randomly sampled agents, and lack the ability to communicate further and to identify the informed individual. Additionally, agents cannot retain any information from one round to the next. According to a recent publication by Becchetti et al. in SODA (2024), a logarithmic convergence time without memory is achievable in the parallel setting (where agents are updated simultaneously), as long as the number of samples is at least $\Omega(\sqrt{n \log n})$. However, determining the minimal sample size for an efficient protocol to exist remains a challenging open question. As a preliminary step towards an answer, we establish the first lower bound for this problem in the parallel setting. Specifically, we demonstrate that it is impossible for any memory-less protocol with constant sample size, to converge with high probability in less than an almost-linear number of rounds. This lower bound holds even when agents are aware of both the exact value of n and their own opinion, and encompasses various simple existing dynamics designed to achieve consensus. Beyond the bit-dissemination problem, our result sheds light on the convergence time of the “minority” dynamics, the counterpart of the well-known majority rule, whose chaotic behavior is yet to be fully understood despite the apparent simplicity of the algorithm.

2012 ACM Subject Classification Theory of computation → Random walks and Markov chains; Theory of computation → Distributed algorithms

Keywords and phrases Opinion Dynamics, Consensus Protocols, Collective Animal Behavior, Lower Bound

Digital Object Identifier 10.4230/LIPIcs.DISC.2024.18

Related Version *Full Version:* <https://arxiv.org/abs/2402.11553>

Funding *Niccolò D’Archivio:* This work has been supported by the AID INRIA-DGA project n°2023000872 “BioSwarm”.

Acknowledgements The authors would like to thank Andrea Clementi for his feedback, and Amos Korman for helpful comments and coming up with the minority dynamics. Thanks also to Luca Trevisan for preliminary discussions, and for everything else.

1 Introduction

Exploring the computational power and limits of well-chosen models – ones that are simple enough for analytical tractability, and yet relevant to specific biological scenarios – can lead to insightful conclusion about the functioning of biological distributed systems [1, 18, 16, 20, 11]. In line with this approach, we consider the *bit-dissemination* problem, introduced in [9] in order to evaluate the possibility to solve two fundamental problems concurrently: reaching agreement efficiently, while ensuring that an information possessed initially by a single



© Niccolò D’Archivio and Robin Vacus;
licensed under Creative Commons License CC-BY 4.0
38th International Symposium on Distributed Computing (DISC 2024).
Editor: Dan Alistarh; Article No. 18; pp. 18:1–18:21



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

individual is propagated to the whole group. In order to fit biological scenarios, the problem features extremely constrained communications. Agents engage in random interactions with just a few individuals at a time, as in the *PULL* model. Furthermore, they can only disclose their current decision and no other information, following an assumption introduced in [23] to model situations in which individuals do not actively communicate [14].

In this paper, we further restrict attention to *memory-less* entities lacking the ability to perform computations over extended periods of time; or at least not in a sufficiently reliable manner. In particular, this assumption precludes the possibility to maintain clocks and counters, or to estimate the tendency of the dynamics. Biological ensembles for which it is plausible include ant colonies [17], slime molds, cells and bacteria [24] or even plants [29].

Although it is always hard to rule out the possibility that species make use of memory (especially in the case of social insects), this modeling choice remains applicable even without presupposing specific cognitive abilities, as many species often stick to simple behavioral rules when attempting to reach a consensus, such as quorum sensing [28] or alignment rules [30]. Furthermore, our goal is not to focus on one truly realistic model, but rather to find out what are the minimal requirements to perform certain tasks.

Informal description of the problem

More precisely, we consider a group of n agents holding binary opinions. One of the agents, referred to as the *source*, knows what opinion is “correct” and remains with it at all times. Execution proceeds in discrete rounds. We assume that agents have no memory of what happened in previous rounds, besides their current opinion. In the *parallel* setting, all non-source agents are activated simultaneously in every round, while in the *sequential* setting, only one non-source agent, selected uniformly at random, is activated. Upon activation, a non-source agent i samples a set S consisting of ℓ other agents drawn uniformly at random (with replacement¹). Then, based only on the opinions of the agents in S and on its own opinion, Agent i may choose to adopt a new opinion. In particular, Agent i does not know whether or not S contains the source. A protocol is successful if every non-source agent eventually adopts the correct opinion and remains with it forever. Finally, a protocol must converge independently of the initial opinions of the agents (including the correct opinion), which can be thought of as being chosen by an adversary.

Previous works

The main parameters of the problems are the activation pattern (which may be parallel or sequential) and the sample size ℓ . To compare protocols across various settings, we are typically interested in their *convergence time*, i.e., the total number of activations required to reach consensus w.r.t. the correct opinion. Here, we will express the convergence time in terms of *parallel rounds*: one parallel round is made up of n activations, which corresponds to 1 round in the parallel setting and n rounds in the sequential setting (although the two settings are not equivalent).

The bit-dissemination problem without memory was first studied in [5], in the sequential setting, where nearly matching lower and upper bounds are given. On the one hand, the authors show that no protocol can converge in less than $\Omega(n)$ parallel rounds in expectation, regardless of the sample size. On the other hand, they show that the well-known Voter

¹ Note that when $\ell \ll n$, sampling with and without replacement are essentially equivalent, and we choose the former for the sake of convenience.

dynamics (Algorithm 1 in Appendix A) achieves consensus in $O(n \log^2 n)$ parallel rounds with high probability. Since the Voter dynamics only needs the sample size to be 1, these results imply that ℓ is not a critical parameter in the sequential setting.

Later, the authors of [7] show that the aforementioned lower bound does not hold in the parallel setting. In fact, they show that the *minority* dynamics (Algorithm 2 in Appendix A) converges in $O(\log^2 n)$ parallel rounds w.h.p., as long as the sample size is at least $\Omega(\sqrt{n \log n})$. This finding reveals that the best convergence times achievable in the sequential versus parallel settings differ by an exponential factor. The bit-dissemination problem is all the more interesting to study as it is one of the most natural ones exhibiting this property.

Part of the explanation behind this curious phenomenon is that the stochastic processes involved are of different mathematical natures, in one setting compared to the other. In the sequential setting, the number of agents with opinion 1 may only vary by at most one unit in every round, since only one agent is activated at a time. Therefore, independently of the protocol being operated, the evolution of the system can always be described by a “birth-death” chain, i.e., a Markov chain whose underlying graph is a path of size n . In fact, all proofs in [5] heavily rely on this observation. In contrast, in the parallel setting, the process may jump from any configuration to any other – albeit with extremely small probability. On the one hand, this characteristic allows for fast convergence. For instance, qualitatively speaking, the minority dynamics succeeds by first reaching a configuration in which an appropriate proportion of the agents hold the wrong opinion; after what all non-source agents, perceiving the same minority, simultaneously adopt the correct opinion. On the other hand, it complicates the analysis of any protocol in the parallel setting, and even more so the task of deriving lower bounds. Following on from these works, we are ultimately interested in the following question:

Is there any protocol achieving a poly-logarithmic convergence time in the parallel setting, when the sample size is $o(\sqrt{n})$?

Based on the findings in [7], the minority dynamics is a natural candidate for this task. Despite its extreme simplicity, the conditions under which it is able to converge quickly have not been identified. While the analysis in [7] relies on a sample size of at least $\Omega(\sqrt{n \log n})$, the authors do not provide a lower bound for this parameter, nor do they justify informally why this quantity is necessary. Therefore, we also consider the following independent question.

What is the minimal sample size for which the minority dynamics converges in poly-logarithmic time?

As a first step towards answering these questions, we focus on the case that the sample size ℓ is constant, or in other words, independent of n . Beyond analytical tractability, several reasons motivate this assumption. First, many existing opinion dynamics that have been traditionally studied in the context of consensus are defined with a small (fixed) number of samples [6]. This is the case for the Voter dynamics and majority dynamics, but also for the undecided states dynamics, as well as all *population protocols* [2], in which agents interact by pairs. In addition, protocols relying on a sample size that increases with n implicitly require the agents to have some knowledge about the size of the population, which is often undesirable or unrealistic in distributed systems. Finally, real biological entities are most likely interacting with few of their conspecifics even when they are part of a larger group. As an illustration, it has been shown empirically that the movement of any bird in a flock depends mostly on its 6 or 7 closest neighbors, regardless of how many individuals are present in the vicinity [3, 8], while other works point at a similar phenomenon in fishes [22].

1.1 Problem Definition

We consider a finite set $I = \{1, \dots, n\}$ of agents. Let $X_t^{(i)} \in \{0, 1\}$ be the *opinion* of Agent i in round t . We assume that Agent 1 is the *source* and holds the *correct* opinion throughout the execution. Denoting the sample size by ℓ (independent of n), a protocol \mathcal{P} is defined as a family of functions

$$g_n^{[b]} : \{0, \dots, \ell\} \rightarrow [0, 1].$$

For a given population size n , an opinion b , and a sample of opinions containing k '1' out of ℓ , $g_n^{[b]}(k)$ gives the probability that an agent adopts opinion 1 in the next round when operating the corresponding protocol. Specifically, the process is obtained by performing the following two steps for every agent i simultaneously, in every round t :

1. A vector $S_t^{(i)} \in I^\ell$ of size ℓ is sampled uniformly at random (u.a.r.) (the same agent may appear several times in $S_t^{(i)}$, and i may appear in it as well).
2. Writing $k_t^{(i)}$ to denote the number of agents with opinion 1 in $S_t^{(i)}$, Agent i updates its opinion according to

$$X_{t+1}^{(i)} \leftarrow 1 \text{ with probability } g_n^{[X_t^{(i)}]}(k_t^{(i)}), \quad 0 \text{ otherwise.}$$

For the sake of clarity, we note a few important consequences of this definition.

- Non-source agents do not know where the opinions that they observe come from. In particular, they do not know if $S_t^{(i)}$ contains the source.
- Non-source agents do not have identifiers, or in other words, all of them must run exactly the same update rule. They are also not aware of the round number (indices are used for analysis purposes only).
- However, non-source agents are aware of their current opinion, as well as the exact value of n .
- Besides their opinion, non-source agents have no memory, in the sense that their behavior cannot depend on any information from previous rounds.

Since agents have no memory besides their opinion, and no identifiers, the configuration of the system in round t can be described simply by a pair (z, X_t) , where $z \in \{0, 1\}$ denotes the correct opinion, and $X_t \in \{0, \dots, n\}$ denotes the number of agents with opinion 1. For a given $n \in \mathbb{N}$, and an initial configuration $C = (z, X_0)$, we define the convergence time of protocol \mathcal{P} as the first round for which all agents have adopted the correct opinion and remain with it forever, that is:

$$\tau_n(\mathcal{P}, C) := \inf\{t \geq 0, \text{ for every } s \geq t, X_s = n \cdot z\}.$$

Given a sequence of events $\{A_n\}_{n \in \mathbb{N}}$, we say that “ A_n happens with high probability (w.h.p.)” if $\mathbb{P}(A_n) = 1 - 1/n^{\Omega(1)}$.

1.2 Our Results

We show that, when the sample size is bounded, any protocol that does not have access to memory needs almost-linear time to solve the bit-dissemination problem.

► **Theorem 1.** *Assume that the sample size ℓ is constant. For every protocol \mathcal{P} , there exists a sequence of initial configuration C_n such that for every $\varepsilon > 0$, the convergence time of \mathcal{P} is greater than $n^{1-\varepsilon}$ w.h.p.:*

$$\mathbb{P}[\tau_n(\mathcal{P}, C_n) < n^{1-\varepsilon}] = \frac{1}{n^{\Omega(1)}}.$$

To the best of our knowledge, this is the first non-trivial lower-bound for this problem in the parallel setting. The proof of Theorem 1 is presented in Section 4, and uses a general result on Markov chains, described in Section 3, as a black box. It consists in studying a kind of “characteristic” function, defined as

$$F_n(p) := -p + \sum_{k=0}^{\ell} \binom{\ell}{k} p^k (1-p)^{\ell-k} \left(p g_n^{[1]}(k) + (1-p) g_n^{[0]}(k) \right). \quad (1)$$

The sum in Equation (1) corresponds to the probability that a non-source agent, taken uniformly at random, adopts opinion 1, given that the current proportion of agents with opinion 1 is p . Informally, $F_n(p)$ measures the “bias” of a protocol \mathcal{P} towards opinion 1, or in other words:

$$\mathbf{E} \left(\frac{X_{t+1}}{n} \mid X_t = x_t \right) \approx \frac{x_t}{n} + F_n \left(\frac{x_t}{n} \right)$$

(see Proposition 5 for a more accurate statement). As a consequence, the sign of the function at p provides information on the trend of the dynamics when the proportion of agents with opinion 1 is p ; moreover, roots correspond to fixed points of the dynamics (either stable or unstable).

Similar functions have already been defined in the literature, where they are typically used to obtain two types of results. On the one hand, they can be used to identify phase transitions [12, 13]. For example, when such characteristic function depends on a parameter α , there might be a critical value α^* at which a new root appears. In that case, the behavior of the dynamics below and above α^* can be significantly different. On the other hand, if a characteristic function has a constant sign over a large interval (and under some additional conditions), the dynamics cannot easily travel in the opposite direction, which can be exploited to obtain lower bounds on the convergence time [12].

In this paper, we use function F_n yet in another way. Specifically, we leverage the fact that a bound on the sample size ℓ implies a bound on the degree of F_n , and therefore a bound on the number of roots within the interval $[0, 1]$. Then, we consider a well-chosen interval of constant length between two roots of F_n , and employ the aforementioned argument to obtain a lower bound: if F_n is negative on this interval, i.e., \mathcal{P} tends to make the proportion of 1-opinions decrease, we show that the process will be slow to reach consensus every time the correct opinion is 1. Conversely, if it is positive, we show that fast convergence fails whenever the correct opinion is 0.

In terms of the dependency on n , we show that our lower bound is nearly tight (up to a sub-polynomial factor) by adapting a well-known result to our setting. Its proof does not introduce any novel argument and is deferred to Appendix C.

► **Theorem 2.** *Consider the Voter dynamics \mathcal{P}^{voter} , with sample size $\ell = 1$. For every sequence of initial configuration C_n , the convergence time of \mathcal{P}^{voter} is less than $2n \log n$ w.h.p.:*

$$\mathbb{P}[\tau_n(\mathcal{P}, C_n) \leq 2n \log n] \geq 1 - \frac{1}{n^2}.$$

When it comes to parameter ℓ , a gap remains between our lower bound and the upper bound in [7], where it is shown that the minority dynamics solves the problem in $O(\log^2 n)$ rounds w.h.p. when ℓ is at least $\Omega(\sqrt{n \log n})$.

Unfortunately, we believe that our techniques cannot be used to extend the lower bound to a higher value of ℓ . Indeed, if $\ell = \Omega(\log n)$, it is already possible for a protocol to converge in just one round w.h.p. from configurations that are arbitrarily far away from the consensus.

This observation destroys any hope of restricting the analysis to a small interval of the configuration space. In contrast, this phenomenon does not happen w.h.p. in our setting (see Proposition 4). However, we have no good reason to think that $\Theta(\sqrt{n \log n})$ is the smallest value of ℓ allowing for an efficient protocol (such as the minority dynamics) – this is left as an open problem.

1.3 Other Related Works

The bit-dissemination problem was also studied under the assumption that agents can use a moderate amount of memory. An efficient protocol is identified in [23], and achieves consensus in $O(\text{polylog } n)$ parallel rounds with high probability. It relies on agents being able to memorize $\log \log n$ bits of information from one round to the next and requires a sample size logarithmic in n . Other candidates are mentioned in [9] but are not analysed. The authors of [15] show that the problem can be solved in the context of population protocols with a memory of only constant size. Importantly however, population protocols do not fit the framework of passive communications. Specifically, interaction rules in this model depend on the exact states of the agents, and not just on their binary opinion.

The bit-dissemination problem is a specific case of the *majority bit-dissemination* problem, introduced in [9] and also addressed in [15]. In this variant, the number of source agents is arbitrarily large, and they may have conflicting *preferences*. The opinions of sources must not necessarily be in line with their preferences, and they can participate to the protocol in the same way as regular agents. The correct opinion is defined as the most widespread preference among sources. On the one hand, an efficient solution can be derived from the results in [15, 4], but require active communications, and relies on memory. On the other hand, the authors of [23] show that the majority bit-dissemination problem is impossible with passive communications.

More generally, many works within the opinion dynamics literature investigate the influence of the presence of “stubborn” or “biased” agents on the behaviour of the system. Typically, these works focus on a single arbitrary process, mainly the Voter dynamics [27, 31, 19, 26, 25], and to the best of our knowledge, do not establish general lower bounds. In contrast, our goal is to better understand the difficulty of spreading information as an algorithmic problem; therefore, we do not want to rule out any imaginable protocol within the constraints of our setting. Furthermore, they often investigate different questions, such as the impact of the number of sources or their position in the network on the convergence time, or assume that sources may have conflicting opinions.

2 Preliminaries

In this section, we make a few general observations that we will use later in our analysis. We assume that ℓ is a constant w.r.t. n .

Conditioning on Agent i sampling exactly k times the opinion “1”, for every $k \in \{0, \dots, \ell\}$, we obtain that

$$\mathbb{P}\left(X_{t+1}^{(i)} = 1 \mid X_t = x_t, X_t^{(i)} = b\right) = \sum_{k=0}^{\ell} \binom{\ell}{k} \left(\frac{x_t}{n}\right)^k \left(1 - \frac{x_t}{n}\right)^{\ell-k} g_n^{[b]}(k). \quad (2)$$

After convergence has happened, i.e., $X_t \in \{0, n\}$, this probability must be equal to 0 or 1 respectively so that a consensus is maintained. This imposes a constraint on any protocol attempting to solve the bit-dissemination problem, which can be formalized as follows.

► **Proposition 3.** Any protocol \mathcal{P} solving the bit-dissemination problem must satisfy $g_n^{[0]}(0) = 0$ and $g_n^{[1]}(\ell) = 1$.

Proof. Consider the case that the correct opinion $z = 0$. If $X_t = 0$ for some round t , then each agent has opinion 0, and receives exactly 0 samples equal to '1'. Following the protocol, all agents adopt opinion 1 in the next round w.p. $g_n^{[0]}(0)$ independently of each-other. If $g_n^{[0]}(0) > 0$, then

$$\mathbb{P}(X_{t+1} = 0 \mid X_t = 0) = \prod_{i \geq 2} \mathbb{P}(X_{t+1}^{(i)} = 0 \mid X_t = 0) = \left(1 - g_n^{[0]}(0)\right)^{n-1} < 1.$$

Therefore, $\inf\{t \geq 0, \text{ for every } s \geq t, X_s = 0\} = +\infty$ almost surely, and the protocol cannot solve the bit-dissemination problem in the sense of Section 1.1. By symmetry, we obtain the other statement about $g_n^{[1]}(\ell)$, which concludes the proof of Proposition 3. ◀

Accordingly, we will always assume that $g_n^{[0]}(0) = 0$ and $g_n^{[1]}(\ell) = 1$. Using this assumption, we can show a general upper bound on the fraction of agents with opinion 0 that can change opinion in a single time step.

► **Proposition 4.** Let $c \in (0, 1)$ and consider a protocol \mathcal{P} solving the bit-dissemination problem. There is a constant $y = y(c, \ell) \in (c, 1)$ s.t. for every n large enough, and $x_t \leq cn$,

$$\mathbb{P}(X_{t+1} \leq yn \mid X_t = x_t) \geq 1 - \exp\left(-2n^{-1/2}\right).$$

Proof. Let $t \in \mathbb{N}$, and $x_t \leq cn$. By Equation (2), and since we assumed $g_n^{[0]}(0) = 0$,

$$\begin{aligned} \mathbb{P}\left(X_{t+1}^{(i)} = 0 \mid X_t = x_t, X_t^{(i)} = 0\right) &= 1 - \sum_{k=0}^{\ell} \binom{\ell}{k} \left(\frac{x_t}{n}\right)^k \left(1 - \frac{x_t}{n}\right)^{\ell-k} g_n^{[0]}(k) \\ &= \sum_{k=0}^{\ell} \binom{\ell}{k} \left(\frac{x_t}{n}\right)^k \left(1 - \frac{x_t}{n}\right)^{\ell-k} \left(1 - g_n^{[0]}(k)\right) \\ &\geq \left(1 - \frac{x_t}{n}\right)^{\ell} \geq (1 - c)^{\ell}. \end{aligned}$$

Let Y be the number of agents with opinion 0 in round t , that keep opinion 0 in round $t + 1$ (conditioning on $X_t = x_t$). By assumption, $n - X_t \geq (1 - c)n$, so the last equation implies the following domination²:

$$Y \succeq \text{Binomial}\left((1 - c)n, (1 - c)^{\ell}\right) := Z.$$

Let $a = a(c, \ell) := (1 - c)^{\ell+1}$, so that $\mathbf{E}(Z) = an$. Let $a' := a - n^{-1/4}$. For n large enough, we have $a' > a/2$. By Hoeffding's bound, we have

$$\mathbb{P}\left(Z \leq \frac{an}{2}\right) \leq \mathbb{P}(Z \leq a'n) = \mathbb{P}\left(Z \leq \mathbf{E}(Z) - n^{3/4}\right) \leq \exp\left(-2n^{1/2}\right).$$

Finally, setting $y = y(c, \ell) := \max(1 - a/2, c)$ (so that $y \in (c, 1)$), we obtain

$$\mathbb{P}(X_{t+1} \geq yn \mid X_t = x_t) \leq \mathbb{P}\left(Y \leq \frac{an}{2}\right) \leq \mathbb{P}\left(Z \leq \frac{an}{2}\right) \leq \exp\left(-2n^{1/2}\right),$$

which concludes the proof of Proposition 4. ◀

² Given two real-valued random variables X and Y , we say that X is *stochastically dominated* by Y , and write $X \preceq Y$, if for every $x \in \mathbb{R}$, $\mathbb{P}(X > x) \leq \mathbb{P}(Y > x)$.

18:8 On the Limits of Information Spread by Memory-Less Agents

Finally, the following proposition justifies the informal claim made in Section 1.2 according to which the function F_n , defined in Equation (1), represents the “bias” of the corresponding protocol towards opinion 1.

► **Proposition 5.** *For every protocol \mathcal{P} , and every $x_t \in [n]$,*

$$\mathbf{E}(X_{t+1} \mid X_t = x_t) \leq x_t + n F_n\left(\frac{x_t}{n}\right) + 1, \quad (3)$$

$$\mathbf{E}(X_{t+1} \mid X_t = x_t) \geq x_t + n F_n\left(\frac{x_t}{n}\right) - 1. \quad (4)$$

Proof. For a non-source agent $i \in I \setminus \{1\}$, an opinion $b \in \{0, 1\}$, and any $p \in [0, 1]$, let

$$P_b := \mathbb{P}\left(X_{t+1}^{(i)} = 1 \mid X_t = np, X_t^{(i)} = b\right) = \sum_{k=0}^{\ell} \binom{\ell}{k} p^k (1-p)^{\ell-k} g_n^{[b]}(k),$$

where the second equality is a restatement of Equation (2). Note that by definition of F_n ,

$$F_n(p) = p P_1 + (1-p) P_0 - p. \quad (5)$$

Denoting by z the correct opinion, one can check that there are $X_t - z$ non-source agents with opinion 1 in round t , and $n - X_t - (1 - z)$ non-source agents with opinion 0. Hence,

$$\begin{aligned} \mathbf{E}(X_{t+1} \mid X_t = np) &= z + (np - z)P_1 + (n - np - (1 - z))P_0 \\ &= n(pP_1 + (1-p)P_0) + z(1 - P_1) - (1 - z)P_0 \\ &= np + n F_n(p) + z(1 - P_1) - (1 - z)P_0. \end{aligned} \quad (\text{by Equation (5)})$$

Note that for any source opinion $z \in \{0, 1\}$, since $P_0, P_1 \in [0, 1]$, we have

$$-1 \leq z(1 - P_1) - (1 - z)P_0 \leq +1,$$

from which Equations (3) and (4) follow by taking $p := x_t/n$. ◀

3 An Intermediate Result on Markov Chains

In this section, we present a result that we will use later as a black box (with $a_1, a_2, a_3 \in [0, 1]$). Informally, the theorem says that if a Markov chain is a super-martingale over an interval of values, and given that it cannot skip the interval entirely, then the time required to cross the interval is at least the time needed by a martingale to escape it.

► **Theorem 6.** *Let $\{X_t\}_{t \in \mathbb{N}}$ be a Markov chain on \mathbb{Z} and $\varepsilon > 0$. If there are $a_1 < a_2 < a_3 \in \mathbb{R}$ s.t.*

(i) *for every $x_t \in \{[a_1 n], \dots, [a_3 n]\}$, $\mathbf{E}(X_{t+1} \mid X_t = x_t) \leq x_t + 1$,*

(ii) *for every $x_t < a_1 n$, $\mathbb{P}(X_{t+1} > a_2 n \mid X_t = x_t) = \exp(-n^{\Omega(1)})$,*

(iii) *$\mathbb{P}(|X_{t+1} - \mathbf{E}(X_{t+1} \mid X_t)| > n^{1/2+\varepsilon/4}) < 2 \exp(-2n^{\varepsilon/2})$,*


then for $X_0 = \frac{a_2 + a_3}{2} \cdot n$ and n large enough, we have w.h.p.

$$\inf\{t \in \mathbb{N}, X_t \geq a_3 n\} \geq n^{1-\varepsilon}.$$

Proof. Let $T = n^{1-\varepsilon}$. Let $Y_t := X_t - t$. We will consider the Doob decomposition of Y_t : for every $t \geq 1$, let

$$A_0 := 0 \quad \text{and for all } t > 0, \quad A_t := \sum_{k=1}^t [\mathbf{E}(Y_k \mid Y_{k-1}) - Y_{k-1}],$$

$$M_0 := Y_0 \quad \text{and for all } t > 0, \quad M_t := Y_0 + \sum_{k=1}^t [Y_k - \mathbf{E}(Y_k \mid Y_{k-1})].$$



■ **Figure 1** Sketch of the proof of Theorem 6. (a) By assumption (ii), with high probability, Y_t cannot jump from below $a_1 n - t$ to above $a_2 n - t$ in a single step, let alone $a_2 n$. (b) In Claim 7, we use assumption (i) and the properties of the Doob's decomposition to show that, if $a_1 n - t \leq Y_t \leq M_t$, then Y_t cannot jump above M_t in a single step. (c) In Claim 8, we use the Azuma-Hoeffding inequality to show that M_t remains in the interval $[a_2 n + T, a_3 n - T]$ for at least T rounds w.h.p. Overall, (a) (b) and (c) implies that Y_t must remain below $a_3 n - T$ for at least T rounds w.h.p., yielding the desired conclusion.

With this definition, one can check that $Y_t = M_t + A_t$, and that $\{M_t\}_{t \in \mathbb{N}}$ is a martingale. The main ideas of the proof are depicted on Figure 1.

First, we show that by construction and Assumption (i), Y_t can never “jump over” M_t in one round, as long as it starts from the interval $\{a_1 n - t, \dots, a_3 n - t\}$.

▷ **Claim 7.** For every $t \in \mathbb{N}$,

$$M_t \geq Y_t \text{ and } Y_t \in \{a_1 n - t, \dots, a_3 n - t\} \implies M_{t+1} \geq Y_{t+1}.$$

Proof. Let $y_t \in \{a_1 n - t, \dots, a_3 n - t\}$ and $m_t \geq y_t$, and consider the event

$$\mathcal{E} := \{Y_t = y_t \cap M_t = m_t\}.$$

We have

$$\begin{aligned} \mathbf{E}(Y_{t+1} \mid Y_t = y_t) &= \mathbf{E}(X_{t+1} - (t+1) \mid X_t = y_t + t) \\ &= \mathbf{E}(X_{t+1} \mid X_t = y_t + t) - (t+1) \\ &\leq (y_t + t + 1) - (t+1) && \text{(by (i))} \\ &= y_t. \end{aligned}$$

Therefore,

$$(A_{t+1} - A_t \mid \mathcal{E}) = \left(\mathbf{E}(Y_{t+1} \mid Y_t) - Y_t \mid \mathcal{E} \right) = \mathbf{E}(Y_{t+1} \mid Y_t = y_t) - y_t \leq 0.$$

Since $m_t \geq y_t$, this implies

$$(A_{t+1} \mid \mathcal{E}) \leq (A_t \mid \mathcal{E}) = (Y_t - M_t \mid \mathcal{E}) = y_t - m_t \leq 0,$$

and thus,

$$(Y_{t+1} \mid \mathcal{E}) = (M_{t+1} + A_{t+1} \mid \mathcal{E}) \leq (M_{t+1} \mid \mathcal{E}),$$

which concludes the proof of Claim 7. ◁

18:10 On the Limits of Information Spread by Memory-Less Agents

Now, using Azuma's inequality and Assumptions (ii) and (iii), we establish high probability bounds on the martingale M_t .

▷ **Claim 8.** With high probability, for every $t \leq T$, $a_2 n + T < M_t < a_3 n - T$.

Proof. By construction, and since X_t and Y_t only differ by a deterministic quantity,

$$M_{t+1} - M_t = Y_{t+1} - \mathbf{E}(Y_{t+1} | Y_t) = X_{t+1} - \mathbf{E}(X_{t+1} | X_t).$$

By assumption (iii) in the statement, this implies

$$\mathbb{P}\left(|M_{t+1} - M_t| > n^{\frac{1}{2} + \frac{\epsilon}{4}}\right) \leq 2 \exp(-2n^{\frac{\epsilon}{2}}).$$

By the union bound,

$$\mathbb{P}\left(\exists s \leq t, |M_{s+1} - M_s| > n^{\frac{1}{2} + \frac{\epsilon}{4}}\right) \leq 2t \exp(-2n^{\frac{\epsilon}{2}}).$$

Let $\alpha = (a_3 - a_2)/4$, so that

$$M_0 + \alpha n = X_0 + \alpha n = \frac{a_2 + a_3}{2} n + \alpha n = a_3 n - \alpha n, \quad \text{and } M_0 - \alpha n = a_2 n + \alpha n. \quad (6)$$

By the Azuma-Hoeffding inequality applied to $\{M_t\}_{t \in \mathbb{N}}$, we then have for every $t \leq T$,

$$\begin{aligned} \mathbb{P}(|M_t - M_0| > \alpha n) &\leq 2 \exp\left(-\frac{\alpha^2 n^2}{2t n^{1 + \frac{\epsilon}{2}}}\right) + 2t \exp(-2n^{\frac{\epsilon}{2}}) \\ &\leq 2 \exp\left(-\frac{\alpha^2 n^2}{2T n^{1 + \frac{\epsilon}{2}}}\right) + 2T \exp(-2n^{\frac{\epsilon}{2}}) \quad (\text{since } t \leq T) \\ &= 2 \exp\left(-\frac{\alpha^2}{2} \cdot n^{\frac{\epsilon}{2}}\right) + 2T \exp(-2n^{\frac{\epsilon}{2}}). \quad (\text{since } T = n^{1-\epsilon}) \end{aligned}$$

By the union bound,

$$\mathbb{P}(\exists t \leq T, |M_t - M_0| > \alpha n) \leq 2T \exp\left(-\frac{\alpha^2}{2} \cdot n^{\frac{\epsilon}{2}}\right) + 2T^2 \exp(-2n^{\frac{\epsilon}{2}}) = o(n^{-2}). \quad (7)$$

Finally, note that for n large enough, $T = n^{1-\epsilon} < \alpha n$, and hence,

$$\begin{aligned} &\mathbb{P}(\exists t \leq T, M_t \notin \{a_2 n + T, \dots, a_3 n - T\}) \\ &\leq \mathbb{P}(\exists t \leq T, M_t \notin \{a_2 n + \alpha n, \dots, a_3 n - \alpha n\}) \\ &= \mathbb{P}(\exists t \leq T, |M_t - M_0| > \alpha n) \quad (\text{by Equation (6)}) \\ &= o(n^{-2}), \quad (\text{by Equation (7)}) \end{aligned}$$

which concludes the proof of Claim 8. ◁

Next, we use Claim 7 and Claim 8 and Assumption (ii) of the theorem to show that Y_t can never jump over M_t , with high probability.

▷ **Claim 9.** With high probability, for every $t \leq T$, $M_t \geq Y_t$.

Proof. We will be conditioning on the two following events:

$$\begin{aligned} \mathcal{E}_1 &:= \{\forall t \leq T, Y_t \leq a_1 n - t \implies Y_{t+1} \leq a_2 n - t\}, \\ \mathcal{E}_2 &:= \{\forall t \leq T, a_2 n + T < M_t < a_3 n - T\}. \end{aligned}$$

Note that \mathcal{E}_2 happen w.h.p. as a consequence of Claim 8. Moreover,

$$\begin{aligned}
\mathbb{P}(\mathcal{E}_1) &= 1 - \mathbb{P}\left(\bigcup_{t=0}^T \{Y_t \leq a_1 n - t\} \cap \{Y_{t+1} > a_2 n - t\}\right) \\
&= 1 - \mathbb{P}\left(\bigcup_{t=0}^T \{X_t \leq a_1 n\} \cap \{X_{t+1} > a_2 n\}\right) \\
&\geq 1 - \sum_{t=0}^T \mathbb{P}(X_t \leq a_1 n \cap X_{t+1} > a_2 n) && \text{(by the union bound)} \\
&\geq 1 - \sum_{t=0}^T \mathbb{P}(X_{t+1} > a_2 n \mid X_t \leq a_1 n) \\
&\geq 1 - T \exp(-n^{\Omega(1)}) && \text{(by assumption (ii))} \\
&\geq 1 - o(n^{-2}). && \text{(since } T = n^{1-\varepsilon}\text{)}
\end{aligned}$$

Hence, \mathcal{E}_1 also happens w.h.p., and so does $\mathcal{E}_1 \cap \mathcal{E}_2$. To conclude the proof, we will show that

$$\mathcal{E}_1 \cap \mathcal{E}_2 \implies \forall t \leq T, M_t \geq Y_t. \quad (8)$$

We will proceed by induction on t . By definition, we have $M_0 = Y_0$. Now, let $t < T$, and consider the case that \mathcal{E}_1 and \mathcal{E}_2 hold, and that $M_t \geq Y_t$.

- If $Y_t < a_1 n - t$, we have

$$Y_{t+1} \underset{(\mathcal{E}_1)}{\leq} a_2 n - t \leq a_2 n + T \underset{(\mathcal{E}_2)}{<} M_{t+1}.$$

- Otherwise, by induction hypothesis and \mathcal{E}_2 , we have $a_1 n - t \leq Y_t \leq M_t < a_3 n - T < a_3 n - t$. In this case, $M_{t+1} \geq Y_{t+1}$ follows as a consequence of Claim 7.

By induction, we deduce that Equation (8) holds, which concludes the proof of Claim 9. \triangleleft

Finally, we are ready to conclude. By Claim 9, with high probability, for every $t \leq T$, $X_t \leq M_t + t$. Therefore,

$$\inf\{t \in \mathbb{N}, X_t > a_3 n\} \geq \inf\{t \in \mathbb{N}, M_t > a_3 n - t\} \quad \text{w.h.p.}$$

Moreover, by Claim 8,

$$\inf\{t \in \mathbb{N}, M_t > a_3 n - t\} > T \quad \text{w.h.p.,}$$

which gives the desired result. \triangleleft

By symmetry, the following result can be deduced directly from Theorem 6.

► **Corollary 10.** *Let $\{X_t\}$ be a Markov chain on \mathbb{Z} and $\varepsilon > 0$. If there are $a_1 < a_2 < a_3 \in \mathbb{R}$ s.t.*

- (i) *for every $x_t \in \{[a_1 n], \dots, [a_3 n]\}$, $\mathbf{E}(X_{t+1} \mid X_t = x_t) \geq x_t - 1$,*
 - (ii) *for every $x_t > a_3 n$, $\mathbb{P}(X_{t+1} < a_2 n \mid X_t = x_t) = \exp(-n^{\Omega(1)})$,*
 - (iii) *$\mathbb{P}(|X_{t+1} - \mathbf{E}(X_{t+1} \mid X_t)| > n^{1/2+\varepsilon/4}) < 2 \exp(-2n^{\varepsilon/2})$,*
- then for $X_0 = \frac{a_1 + a_2}{2} \cdot n$ and for n large enough, we have w.h.p.*

$$\inf\{t \in \mathbb{N}, X_t \leq a_1 n\} \geq n^{1-\varepsilon}.$$

Proof. It is easy to check that if the assumptions of Corollary 10 hold for $\{X_t\}_{t \in \mathbb{N}}$ and some constants $a_1 < a_2 < a_3$, then the assumptions of Theorem 6 hold w.r.t. $\{-X_t\}_{t \in \mathbb{N}}$ and $-a_3 < -a_2 < -a_1$ respectively – and the conclusion follows. \blacktriangleleft

4 The Main Proof

4.1 The Voter Dynamics

First, we focus on the Voter dynamics (Algorithm 1) and show that it satisfies the lower bound stated in Theorem 1, which we will prove in its full generality in Section 4.2. We start by observing that for g^{voter} defined in Equation (13), and by definition in Equation (1),

$$F_n^{\text{voter}}(p) = -p + \sum_{k=0}^{\ell} \binom{\ell}{k} p^k (1-p)^{\ell-k} \cdot \frac{k}{\ell} = -p + p = 0.$$

Then, we conclude by applying the following result. Due to space constraints, its proof is deferred to Appendix D.

► **Lemma 11.** *Consider a protocol \mathcal{P} satisfying $F_n = 0$ for every n large enough. There exists a sequence of configurations $\{C_n\}$ such that for every $\epsilon > 0$, with high probability,*

$$\tau_n(\mathcal{P}, C_n) > n^{1-\epsilon}.$$

4.2 General Case

Our main result (Theorem 1) will follow as a consequence of Theorem 12 below.

► **Theorem 12.** *For every $\epsilon > 0$ and every protocol \mathcal{P} , there exists an infinite set $\mathcal{S} \subseteq \mathbb{N}$ and a sequence of configurations $\{C_n\}$ such that for every $n \in \mathcal{S}$,*

$$\mathbb{P}(\tau_n(\mathcal{P}, C_n) > n^{1-\epsilon}) \geq 1 - \frac{1}{n^{\Omega(1)}}.$$

In other words, the convergence time of \mathcal{P} restricted to \mathcal{S} is greater than $n^{1-\epsilon}$ w.h.p.

Proof. Recall the definition of F_n in Equation (1). If there is $N \in \mathbb{N}$ s.t. for every $n \geq N$, $F_n = 0$, then we can conclude by applying Lemma 11. Otherwise, there is an infinite set $\mathcal{S}_0 \subseteq \mathbb{N}$ s.t. for every $n \in \mathcal{S}_0$, $F_n \neq 0$, which we will assume from now on.

By definition in Equation (1), F_n is a polynomial of degree at most $\ell + 1$. For $n \in \mathcal{S}_0$, let d_n be the number of roots of F_n in the interval $[0, 1]$ (counted with multiplicity). By Proposition 3, $g_n^{[0]}(0) = 0$ and $g_n^{[1]}(\ell) = 1$, so $F_n(0) = F_n(1) = 0$, and thus $d_n \in \{2, \dots, \ell + 1\}$. Since d_n can only adopt finitely many values, there exists $d \in \{2, \dots, \ell + 1\}$ and an infinite set $\mathcal{S}_1 \subseteq \mathcal{S}_0$ s.t. for every $n \in \mathcal{S}_1$, $d_n = d$.


For $n \in \mathcal{S}_1$, let $0 = r_n^{(1)} \leq \dots \leq r_n^{(d)} = 1$ be the roots of F_n within the interval $[0, 1]$, with multiplicity, in increasing order. The sequence $\{(r_n^{(1)}, \dots, r_n^{(d)})\}_{n \in \mathcal{S}_1}$ is bounded in \mathbb{R}^d by definition. Hence, by the Bolzano-Weierstrass theorem, there exists a converging sub-sequence of $\{(r_n^{(1)}, \dots, r_n^{(d)})\}_{n \in \mathcal{S}_1}$; i.e., there are $0 = r_\infty^{(1)} \leq \dots \leq r_\infty^{(d)} = 1$ together with an infinite set $\mathcal{S}_2 \subseteq \mathcal{S}_1$ s.t. for every $k \in [d]$,

$$\lim_{\substack{n \rightarrow +\infty \\ n \in \mathcal{S}_2}} r_n^{(k)} = r_\infty^{(k)}. \quad (9)$$

Let $k_0 := \min\{k \in [d], r_\infty^{(k)} = 1\}$. Note that $k_0 \geq 2$ (since $r_\infty^{(1)} = 0$) and by definition, $r_\infty^{(k_0-1)} < 1 = r_\infty^{(k_0)}$. Moreover, for every $n \in \mathcal{S}_2$, F_n is non-zero and has constant sign on $(r_n^{(k_0-1)}, r_n^{(k_0)})$. Therefore, there exists an infinite set $\mathcal{S}_3 \subseteq \mathcal{S}_2$ s.t.

1. either $\forall n \in \mathcal{S}_3$, $F_n < 0$ on $(r_n^{(k_0-1)}, r_n^{(k_0)})$,
2. or $\forall n \in \mathcal{S}_3$, $F_n > 0$ on $(r_n^{(k_0-1)}, r_n^{(k_0)})$.

In the remainder of the proof, we will analyse these two cases separately. The reader is strongly encouraged to consult Figure 2 and 3 respectively for a clearer understanding of the argument.



■ **Figure 2 Illustration of the arguments for Case 1.** We consider a configuration in which the correct opinion is 1. Constant a_1 is fixed arbitrarily in the interval $(r_\infty^{(k_0-1)}, 1)$. Then, a_2 is chosen according to Proposition 4 to ensure that X_t cannot jump from below $a_1 n$ to above $a_2 n$. Finally, a_3 is set anywhere in the interval $(a_2, 1]$. By assumption, $F_n < 0$ on $[a_1, a_3]$, and we can eventually apply Theorem 6.

Case 1

Let $a_1 \in (r_\infty^{(k_0-1)}, 1)$. Let $a_2 = a_2(a_1, \ell) \in (a_1, 1)$ given by Proposition 4, s.t. for n large enough,

$$\text{for every } x_t \leq a_1 n, \quad \mathbb{P}(X_{t+1} \leq a_2 n \mid X_t = x_t) \geq 1 - \exp(-2n^{-1/2}).$$

Let $a_3 \in (a_2, 1)$. By Equation (9), for n large enough, $r_n^{(k_0-1)} < a_1$ and $r_n^{(k_0)} > a_3$. We now wish to use Theorem 6, with a_1, a_2, a_3 as we just defined. Let us check that every assumption holds:

- For every $x_t \in \{a_1 n, \dots, a_3 n\}$, we have $x_t/n \in [a_1, a_3] \subset [r_n^{(k_0-1)}, r_n^{(k_0)}]$, and so $F_n(x_t/n) < 0$ by assumption. Therefore, Equation (3) gives

$$\mathbf{E}(X_{t+1} \mid X_t = x_t) \leq x_t + n F_n\left(\frac{x_t}{n}\right) + 1 < x_t + 1, \quad (10)$$

so assumption (i) in the statement of Theorem 6 holds.

- Assumption (ii) holds by Proposition 4.
- Finally, assumption (iii) follows from Hoeffding's bound: conditioning on X_t , X_{t+1} is the sum of n Bernoulli random variables, then the result follows choosing $\delta = n^{1/2+\varepsilon/4}$.

Assuming that the source has opinion $z = 1$, we apply Theorem 6, which implies the existence of an initial configuration C_n s.t. the convergence time is bounded w.h.p.:

$$\tau_n(g, C_n) = \inf\{t \in \mathbb{N}, X_t = n\} \geq \inf\{t \in \mathbb{N}, X_t \geq a_3 n\} \geq n^{1-\varepsilon}.$$


Case 2

Let $a_1, a_2, a_3 \in (r_\infty^{(k_0-1)}, 1)$, with $r_\infty^{(k_0-1)} < a_1 < a_2 < a_3 < 1$. First, we show that by taking n large enough, we can have F_n be arbitrarily close to 0 on the interval $[r_n^{(k_0)}, 1]$.

▷ **Claim 13.** For every $\delta > 0$, for n large enough: $r_n^{(k_0)} > (1+a_3)/2$, and for every $p \in [r_n^{(k_0)}, 1]$, $F_n(p) > -\delta$.

Proof. Let $\delta > 0$ and $n \in \mathcal{S}_2$. Since F_n has bounded coefficients and degree d , and since $F_n(r_n^{(k_0)}) = F_n(1) = 0$, by Claim 17 in Appendix D we obtain the existence of C_0 s.t.

$$\text{for every } p \in [r_n^{(k_0)}, 1], \quad |F_n(p)| < C_0 \cdot (1 - r_n^{(k_0)}). \quad (11)$$



■ **Figure 3 Illustration of the arguments for Case 2.** We consider a configuration in which the correct opinion is 0. Constants a_1, a_2 and a_3 are chosen arbitrarily in the interval $(r_\infty^{(k_0-1)}, 1)$. By assumption, $F_n > 0$ on $[a_1, a_3]$. Moreover, once a_2 and a_3 are fixed, we give a lower-bound on F_n on the interval $[r_n^{(k_0)}, 1]$, by letting $r_n^{(k_0)}$ be sufficiently close to 1, in order to ensure that X_t cannot jump from above $a_3 n$ to below $a_2 n$. Eventually, we are able to apply Corollary 10.

By Equation (9) and by definition of k_0 , $r_n^{(k_0)}$ tends to 1 as n goes to $+\infty$. If n is large enough,

$$r_n^{(k_0)} > \max \left\{ 1 - \frac{\delta}{C_0}, \frac{1 + a_3}{2} \right\}. \quad (12)$$

By Equation (11), this implies that $|F_n(p)| < \delta$ on $[r_n^{(k_0)}, 1]$, which, together with Equation (12), concludes the proof of Claim 13. \triangleleft

Now, we use the previous result to establish a lower bound on $p + F_n(p)$ when $p \geq a_1$.

\triangleright **Claim 14.** For n large enough,

- for every $p \in [a_1, a_3]$, $p + F_n(p) > p$.
- for every $p \in [a_3, 1]$, $p + F_n(p) > a_3$.

Proof. For $p \in [a_1, r_n^{(k_0)})$, $F_n(p) > 0$ by assumption. Therefore:

- for every $p \in [a_1, a_3]$, $p + F_n(p) > p$.
- for every $p \in [a_3, r_n^{(k_0)})$, $p + F_n(p) > p \geq a_3$.

All is left to prove is that $p + F_n(p) > a_3$ on $[r_n^{(k_0)}, 1]$. Let $\delta = (1 - a_3)/2$, and let n be large enough for Claim 13 to hold w.r.t. δ . For every $p \in [r_n^{(k_0)}, 1]$, we have

$$\begin{aligned}
 p + F_n(p) &> r_n^{(k_0)} - \delta && \text{(by Claim 13 and definition of } p) \\
 &= a_3 + \left(\frac{1 - a_3}{2} - \delta \right) + \left(r_n^{(k_0)} - \frac{1 + a_3}{2} \right) \\
 &> a_3, && \text{(by Claim 13 and definition of } \delta)
 \end{aligned}$$

which concludes the proof of Claim 14. \triangleleft

Finally, similarly to the first case, we use Corollary 10 to conclude. Again, we start by checking that all assumptions hold. Let n be large enough for Claim 14 to hold.

- For every $x_t \in \{a_1 n, \dots, a_3 n\}$, we have $x_t/n + F_n(x_t/n) > x_t/n$ by Claim 14. Therefore, Equation (4) rewrites

$$\mathbf{E}(X_{t+1} \mid X_t = x_t) \geq x_t + n F_n \left(\frac{x_t}{n} \right) - 1 \geq x_t - 1,$$

so assumption (i) in the statement of Corollary 10 holds.

- For every $x_t \in \{a_3 n, \dots, n\}$, we have $x_t/n + F_n(x_t/n) > a_3$ by Claim 14. Therefore, Equation (4) rewrites

$$\mathbf{E}(X_{t+1} \mid X_t = x_t) \geq x_t + n F_n\left(\frac{x_t}{n}\right) - 1 \geq a_3 n - 1.$$

Therefore, by Hoeffding’s bound,

$$\begin{aligned} \mathbb{P}(X_{t+1} < a_2 n \mid X_t = x_t) &\geq \mathbb{P}\left(X_{t+1} < \mathbf{E}(X_{t+1}) - \frac{a_3 - a_2}{2} n \mid X_t = x_t\right) \\ &\leq \exp\left(-2 \left(\frac{a_3 - a_2}{2}\right)^2 n\right), \end{aligned}$$

and so assumption (ii) holds.

- Finally, assumption (iii) follows from Hoeffding’s bound, as in the first case.

Assuming that the source has opinion $z = 0$, we apply Corollary 10, which implies the existence of an initial configuration C_n s.t. the convergence time is bounded:

$$\tau_n(g, C_n) = \inf\{t \in \mathbb{N}, X_t = 0\} \geq \inf\{t \in \mathbb{N}, X_t \leq a_1 n\} \geq n^{1-\varepsilon}. \quad \blacktriangleleft$$

5 Discussion and Future Works

In this paper, we explore the minimal requirements for simultaneously reaching consensus and propagating information in a distributed system. We consider memory-less and anonymous agents, which update their opinion synchronously after observing the opinions of a few other agents sampled uniformly at random, and whose goal is to converge on the correct opinion held by a single “source” individual. In addition, we adopt the self-stabilizing framework, which in a memory-less setting, means that convergence must happen for any possible initialization of the opinions of the agents (including the source). Under this model, we show that to obtain a convergence time better than $n^{1-\varepsilon}$, the number ℓ of samples obtained by each agent in every round must necessarily tend towards infinity as n increases. Our result extends the range of values of ℓ for which the performance of the “minority” dynamics (Algorithm 2) is characterized. Our technique, which consists in translating the sample size into the degree of a well-chosen polynomial, and then inspecting its roots, is simple yet quite novel (to the best of our knowledge), and may be used to show similar results in other settings.

Our ultimate goal is to fully characterize the complexity of the bit-dissemination problem in the parallel setting and in the absence of memory, as a function of the sample size. Regarding values of ℓ allowing poly-logarithmic convergence time, there is still a large gap between our lower bound $\ell = \Omega(1)$ and the upper bound $\ell = O(\sqrt{n \log n})$ mentioned in [7]. Closing or narrowing this gap, even specifically for the minority dynamics, would be of appreciable interest in our opinion.

Another natural continuation would be to generalize our result to protocols using a constant amount of memory. If feasible, the resulting lower bound would still be compatible with the algorithm of [23], which requires $\Omega(\log \log n)$ bits of memory.

References

- 1 Yehuda Afek, Noga Alon, Omer Barad, Eran Hornstein, Naama Barkai, and Ziv Bar-Joseph. A biological solution to a fundamental distributed computing problem. *Science*, 331(6014):183–185, January 2011. doi:10.1126/science.1193210.

- 2 James Aspnes and Eric Ruppert. An introduction to population protocols. In Benoît Garbinato, Hugo Miranda, and Luís Rodrigues, editors, *Middleware for Network Eccentric and Mobile Applications*, pages 97–120. Springer, Berlin, Heidelberg, 2009. doi:10.1007/978-3-540-89707-1_5.
- 3 M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic. Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the National Academy of Sciences of the United States of America*, 105(4):1232–1237, January 2008. doi:10.1073/pnas.0711437105.
- 4 Paul Bastide, George Giakkoupis, and Hayk Saribekyan. Self-stabilizing clock synchronization with 1-bit messages. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Proceedings, pages 2154–2173. Society for Industrial and Applied Mathematics, January 2021. doi:10.1137/1.9781611976465.129.
- 5 Luca Becchetti, Andrea Clementi, Amos Korman, Francesco Pasquale, Luca Trevisan, and Robin Vacus. On the role of memory in robust opinion dynamics. In Edith Elkind, editor, *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, volume 1, pages 29–37, Macao, August 2023. International Joint Conferences on Artificial Intelligence Organization. doi:10.24963/ijcai.2023/4.
- 6 Luca Becchetti, Andrea Clementi, and Emanuele Natale. Consensus dynamics: An overview. *ACM SIGACT News*, 51(1):58–104, March 2020. doi:10.1145/3388392.3388403.
- 7 Luca Becchetti, Andrea Clementi, Francesco Pasquale, Luca Trevisan, Robin Vacus, and Isabella Ziccardi. The minority dynamics and the power of synchronicity. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Proceedings, pages 4155–4176. Society for Industrial and Applied Mathematics, Alexandria, Virginia, U.S., January 2024. doi:10.1137/1.9781611977912.144.
- 8 William Bialek, Andrea Cavagna, Irene Giardina, Thierry Mora, Edmondo Silvestri, Massimiliano Viale, and Aleksandra M. Walczak. Statistical mechanics for natural flocks of birds. *Proceedings of the National Academy of Sciences*, 109(13):4786–4791, March 2012. doi:10.1073/pnas.1118633109.
- 9 Lucas Boczkowski, Amos Korman, and Emanuele Natale. Minimizing message size in stochastic communication patterns: Fast self-stabilizing protocols with 3 bits. In *Proceedings of the 2017 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Proceedings, pages 2540–2559. Society for Industrial and Applied Mathematics, January 2017. doi:10.1137/1.9781611974782.168.
- 10 Fan Chung and Linyuan Lu. Concentration inequalities and martingale inequalities: A survey. *Internet Mathematics*, 3(1):79–127, January 2006. doi:10.1080/15427951.2006.10129115.
- 11 Andrea Clementi, Francesco d’Amore, George Giakkoupis, and Emanuele Natale. Search via parallel Lévy walks on z_2 . In *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing, PODC’21*, pages 81–91, New York, NY, USA, July 2021. Association for Computing Machinery. doi:10.1145/3465084.3467921.
- 12 Emilio Cruciani, Hlafo Alfie Mimun, Matteo Quattropiani, and Sara Rizzo. Phase transitions of the k -majority dynamics in a biased communication model. In *Proceedings of the 22nd International Conference on Distributed Computing and Networking, ICDCN ’21*, pages 146–155, New York, NY, USA, January 2021. Association for Computing Machinery. doi:10.1145/3427796.3427811.
- 13 Emilio Cruciani, Emanuele Natale, André Nusser, and Giacomo Scornavacca. Phase transition of the 2-Choices dynamics on core–periphery networks. *Distributed Computing*, 34(3):207–225, June 2021. doi:10.1007/s00446-021-00396-5.
- 14 Étienne Danchin, Luc-Alain Giraldeau, Thomas J. Valone, and Richard H. Wagner. Public information: From nosy neighbors to cultural evolution. *Science*, 305(5683):487–491, July 2004. doi:10.1126/science.1098254.

- 15 Bartłomiej Dudek and Adrian Kosowski. Universal protocols for information dissemination using emergent signals. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, pages 87–99, New York, NY, USA, June 2018. Association for Computing Machinery. doi:10.1145/3188745.3188818.
- 16 Yuval Emek and Roger Wattenhofer. Stone age distributed computing. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing*, PODC '13, pages 137–146, New York, NY, USA, July 2013. Association for Computing Machinery. doi:10.1145/2484239.2484244.
- 17 Ofer Feinerman and Amos Korman. Individual versus collective cognition in social insects. *Journal of Experimental Biology*, 220(1):73–82, January 2017. doi:10.1242/jeb.143891.
- 18 Ofer Feinerman, Amos Korman, Zvi Lotker, and Jean-Sebastien Sereni. Collaborative search on the plane without communication. In *Proceedings of the 2012 ACM Symposium on Principles of Distributed Computing*, PODC '12, pages 77–86, New York, NY, USA, July 2012. Association for Computing Machinery. doi:10.1145/2332432.2332444.
- 19 Mikaela Irene D. Fudolig and Jose Perico H. Esguerra. Analytic treatment of consensus achievement in the single-type zealotry voter model. *Physica A: Statistical Mechanics and its Applications*, 413:626–634, November 2014. doi:10.1016/j.physa.2014.07.033.
- 20 Mohsen Ghaffari, Cameron Musco, Tsvetomira Radeva, and Nancy Lynch. Distributed house-hunting in ant colonies. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing*, PODC '15, pages 57–66, New York, NY, USA, July 2015. Association for Computing Machinery. doi:10.1145/2767386.2767426.
- 21 Yehuda Hassin and David Peleg. Distributed probabilistic polling and applications to proportionate agreement. *Information and Computation*, 171(2):248–268, December 2001. doi:10.1006/inco.2001.3088.
- 22 Yael Katz, Kolbjørn Tunstrøm, Christos C. Ioannou, Cristián Huepe, and Iain D. Couzin. Inferring the structure and dynamics of interactions in schooling fish. *Proceedings of the National Academy of Sciences*, 108(46):18720–18725, November 2011. doi:10.1073/pnas.1107583108.
- 23 Amos Korman and Robin Vacus. Early adapting to trends: Self-stabilizing information spread using passive communication. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, PODC'22, pages 235–245, New York, NY, USA, July 2022. Association for Computing Machinery. doi:10.1145/3519270.3538415.
- 24 Melissa B. Miller and Bonnie L. Bassler. Quorum sensing in bacteria. *Annual Review of Microbiology*, 55(1):165–199, 2001. doi:10.1146/annurev.micro.55.1.165.
- 25 V. Moeinifar and S. Gündüç. Zealots' effect on opinion dynamics in complex networks. *Mathematical Modeling and Computing*, 8(2):pp. 203, Wed, 05/05/2021 - 17:53.
- 26 Arpan Mukhopadhyay, Ravi R. Mazumdar, and Rahul Roy. Binary opinion dynamics with biased agents and agents with different degrees of stubbornness. In *2016 28th International Teletraffic Congress (ITC 28)*, pages 261–269, Würzburg, Germany, September 2016. IEEE. doi:10.1109/ITC-28.2016.143.
- 27 V. Sood and S. Redner. Voter model on heterogeneous graphs. *Physical Review Letters*, 94(17):178701, May 2005. doi:10.1103/PhysRevLett.94.178701.
- 28 David J. T. Sumpter, Jens Krause, Richard James, Iain D. Couzin, and Ashley J. W. Ward. Consensus decision making by fish. *Current Biology*, 18(22):1773–1777, November 2008. doi:10.1016/j.cub.2008.09.064.
- 29 Anthony Trewavas. Plant intelligence: Mindless mastery. *Nature*, 415(6874):841–841, February 2002. doi:10.1038/415841a.
- 30 Tamás Vicsek and Anna Zafeiris. Collective motion. *Physics Reports*, 517(3):71–140, August 2012. doi:10.1016/j.physrep.2012.03.004.
- 31 Ercan Yildiz, Asuman Ozdaglar, Daron Acemoglu, Amin Saberi, and Anna Scaglione. Binary opinion dynamics with stubborn agents. *ACM Transactions on Economics and Computation*, 1(4):19:1–19:30, December 2013. doi:10.1145/2538508.

A Well-known Dynamics

In this section, we define two important dynamics. Algorithms 1 and 2 below only describe the behaviour of non-source agents, since the source never changes its opinion.

<p>Algorithm 1 Voter dynamics.</p> <p>Input: An opinion sample S of size ℓ.</p> <p>$X_{t+1}^{(i)} \leftarrow$ a random opinion in S</p> <p><i>(Since S is already sampled uniformly at random, this protocol yields the same dynamics regardless of the value of ℓ, and is usually defined with $\ell = 1$).</i></p>	<p>Algorithm 2 Minority dynamics [7].</p> <p>Input: An opinion sample S of size ℓ.</p> <p>if all opinions in S are equal to x then $X_{t+1}^{(i)} \leftarrow x$ else $X_{t+1}^{(i)} \leftarrow$ minority opinion in S <i>(breaking ties randomly).</i></p>
---	--

In terms of our definition, the Voter dynamics writes

$$g_n^{[0]}(k) = g_n^{[1]}(k) =: g^{\text{voter}}(k) = \frac{k}{\ell}, \quad \text{for every } k \in \{0, \dots, \ell\}. \quad (13)$$

Similarly, given that ties are broken u.a.r., the Minority dynamics is given by

$$g_n^{[0]}(k) = g_n^{[1]}(k) =: g^{\text{minority}}(k) = \begin{cases} 1 & \text{if } k = \ell \text{ or } 0 < k < \frac{\ell}{2}, \\ \frac{1}{2} & \text{if } k = \frac{\ell}{2}, \\ 0 & \text{if } k = 0 \text{ or } \frac{\ell}{2} < k < \ell. \end{cases} \quad (14)$$

B Probabilistic Tools

► **Theorem 15** (Hoeffding's bound). *Let X_1, \dots, X_n be i.i.d. random variables taking values in $\{0, 1\}$, let $X = \sum_{i=1}^n X_i$ and $\mu = \mathbf{E}(X) = n\mathbb{P}(X_1 = 1)$. Then it holds for all $\delta > 0$ that*

$$\mathbb{P}(X \leq \mu - \delta), \mathbb{P}(X \geq \mu + \delta) \leq \exp\left(-\frac{2\delta^2}{n}\right).$$

The following version of Azuma's inequality, which accounts for the possibility that the martingale makes a large jump with a small probability, appears in [10, Section 8, p.34].

► **Theorem 16** (Azuma-Hoeffding inequality). *Let $(X_t)_{t \in \mathbb{N}}$ be a martingale, and let $T \in \mathbb{N}$. If there is $p > 0$ and c_1, \dots, c_T such that*

$$\mathbb{P}(\exists t \leq T, X_t - X_{t-1} > c_t) \leq p,$$

then for every $\delta > 0$,

$$\mathbb{P}(|X_T - X_0| > \delta) \leq 2 \exp\left(-\frac{\delta^2}{2 \sum_{t=1}^T c_t^2}\right) + p.$$

C Upper Bound for the Voter Dynamics

The following proof involves only classical arguments (see, for instance, [31] and [21, Section 2.4]). We nonetheless present it here for the sake of completeness.

Proof of Theorem 2. The reader who is not already familiar with the idea is strongly encouraged to refer to Figure 4 for an illustration.

Without loss of generality, we consider the Voter dynamics with $\ell = 1$. In this case, the sample $S_t^{(i)}$ of a non-source agent $i \neq 1$ in round t is simply an element of the set $I = \{1, \dots, n\}$ of all agents, drawn uniformly at random. In the case of the source, for the sake of the argument, we let $S_t^{(1)} = 1$ for every $t \in \mathbb{N}$, i.e., we consider that the source agent applies the Voter rule but always samples itself.

Now, let us fix an horizon T . We will proceed by examining n random walks $\{W_{T-t}^{(i)}\}_{t \leq T}$, defined on the same randomness, but for which the time flows backward. Specifically, for every $i \in I$, let $W_T^{(i)} = i$, i.e., every random walk “starts” at a different position. Moreover, for $t < T$, let

$$W_t^{(i)} := S_t^{(W_{t+1}^{(i)})}.$$

In other words, if a random walk is in position j in round $t + 1$, and if Agent j samples i in round t , then the random walk “moves” to i in round t . Note that by definition, if a random walk moves to position 1 in round t , it will remain in position 1 for all remaining rounds:

$$W_t^{(i)} = 1 \implies \forall s \in \{1, \dots, t\}, W_s^{(i)} = 1. \quad (15)$$

Moreover, if the random walk indexed by i ends up in position 1 (in round 1), it implies that the agent of index $W_t^{(i)}$ holds the correct opinion in round t :

$$W_1^{(i)} = 1 \implies X_t^{(W_t^{(i)})} = z. \quad (16)$$

To show that, we can simply proceed by induction on t , using the definition of the random walks. Equation (15) and Equation (16) together implies that if the random walks indexed by i ever moves to position 1, then Agent i has the correct opinion in round T :

$$\exists t \leq T, W_t^{(i)} = 1 \implies X_T^{(i)} = z. \quad (17)$$

Therefore, for every $i \neq 1$,

$$\begin{aligned} \mathbb{P}\left(X_T^{(i)} \neq z\right) &\leq \mathbb{P}\left(\forall t \leq T, W_t^{(i)} \neq 1\right) && \text{(by Equation (17))} \\ &\leq \prod_{t=0}^{T-1} \mathbb{P}\left(S_t^{(W_{t+1}^{(i)})} \neq 1 \mid W_{t+1}^{(i)} \neq 1\right) && \text{(by the chain rule)} \\ &= \left(1 - \frac{1}{n}\right)^T. && \text{(uniform and independent samples)} \end{aligned}$$

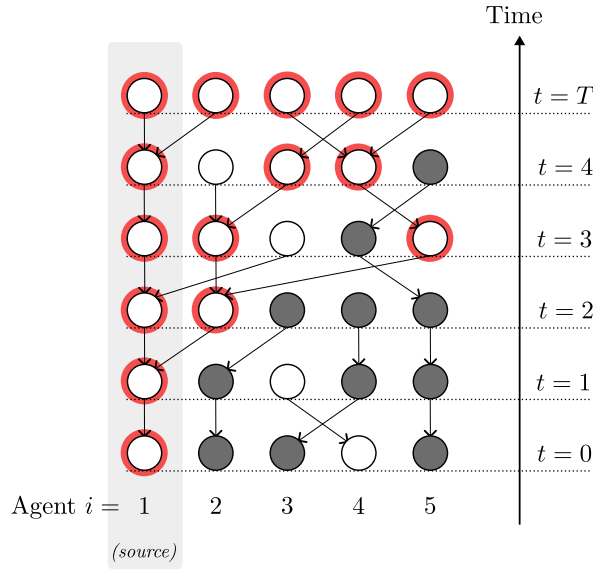
Note that this bound holds trivially for $i = 1$. Therefore, by the union bound,

$$\mathbb{P}\left(\forall i \in I, X_T^{(i)} = z\right) = 1 - \mathbb{P}\left(\exists i \in I, X_T^{(i)} \neq z\right) \geq 1 - \sum_{i \in I} \mathbb{P}\left(X_T^{(i)} \neq z\right) \geq 1 - n \left(1 - \frac{1}{n}\right)^T.$$

Taking $T = 2n \log n$, for n large enough,

$$\left(1 - \frac{1}{n}\right)^{2n \log n} = \exp(2n \log n \log(1 - 1/n)) \leq \exp(-2 \log n) = \frac{1}{n^2},$$

which concludes the proof of Theorem 2. \blacktriangleleft



■ **Figure 4** Depiction of the dual process behind the proof of Theorem 2. The color of the circle in row t , column i , corresponds to the opinion of Agent i in round t : it is black if $X_t^{(i)} = 1$, and white otherwise. An arrow is drawn from $(i, t+1)$ to (j, t) if $S_t^{(i)} = j$, i.e., if Agent i observes Agent j in round t (and thus adopts their opinion in round $t+1$). Red circles depict the locations of n coalescing random walks going backward in time, and initially present at every location. Random walks at a location $i > 1$ make a move using the same randomness as the samples, while the source acts like a sink. If all random walks have coalesced in less than T rounds, it implies that the opinion of each agent in round T comes from the source, and thus that the dynamics has reached consensus on the correct opinion.

D Missing Proofs

\triangleright **Claim 17.** For every M, d , there exists $C_0 = C_0(M, d) > 0$ s.t. for every polynomial P of degree d and coefficients bounded by M , every $a, b \in [0, 1]$ with $P(a) = P(b) = 0$, and every $x \in [a, b]$, $P(x) < C_0 \cdot (b - a)$.

Proof. Since P has degree d and coefficients bounded by M , there exists $C = C(M, d)$ s.t. $|P'(x)| < C$ on $[0, 1]$. Therefore, for every $x \in [a, (a+b)/2]$, we have

$$|P(x)| = |P(x) - P(a)| < C \cdot (x - a) < C \cdot \frac{b - a}{2}.$$

Similarly, for every $x \in [(a+b)/2, b]$, we have

$$|P(x)| = |P(b) - P(x)| < C \cdot (b - x) < C \cdot \frac{b - a}{2}.$$

Taking $C_0 = C/2$ concludes the proof of Claim 17. \triangleleft

Proof of Lemma 11. Let $\varepsilon > 0$. We want to apply Theorem 6 with $a_1 = 1/4$, $a_2 = 1/2$, $a_3 = 3/4$. If the three hypotheses hold, we have for the initial configuration $C_n := (z = 1, X_0 = \frac{a_2 n + a_3 n}{2})$,

$$\tau_n(\mathcal{P}, C_n) \geq \inf\{t \in \mathbb{N}, X_t > a_3 n\} \geq n^{1-\varepsilon}.$$

Now, let us show that the hypotheses hold:

- **Proving (i).** Since $F_n = 0$ for n large enough, we have by Proposition 5 that $\mathbf{E}(X_{t+1} | X_t = x_t) \leq x_t + 1$.
- **Proving (iii).** Conditioning on X_t , X_{t+1} is a sum of n independent Bernoulli random variables. Therefore, we can use Hoeffding's bound to obtain

$$\mathbb{P}(|X_{t+1} - \mathbf{E}(X_{t+1} | X_t)| > n^{1/2+\varepsilon/4}) < 2 \exp(-2n^{\varepsilon/2}),$$

which establishes (iii).

- **Proving (ii).** If $x_t < a_1 n$ and n is large enough, (ii) follows again from Hoeffding's bound:

$$\begin{aligned} \mathbb{P}(X_{t+1} > a_2 n | X_t = x_t) &\leq \mathbb{P}(X_{t+1} > n^{1/2+1/4} + \mathbf{E}(X_{t+1} | X_t = x_t) | X_t = x_t) \\ &< 2 \exp(-2n^{1/2}), \end{aligned}$$

which establishes (ii) and concludes the proof of Lemma 11. ◀