

# Fast and Robust Information Spreading in the Noisy $\mathcal{PULL}$ Model

Extended Abstract

Niccolò D’Archivio

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

Emanuele Natale

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

Amos Korman

Department of Computer Science, University of Haifa  
Haifa, Israel

Robin Vacus

Sorbonne Université, CNRS, LIP6  
Paris, France

## ABSTRACT

Efficient information spreading in stochastic multi-agent systems is a core challenge when communication is noisy, bandwidth-limited, and agents lack global coordination. Yet biological systems—such as ant colonies and fish schools—routinely overcome these constraints: a small number of informed individuals can reliably guide large, uncoordinated populations using minimal, noisy signals. Motivated by these observations, we investigate how reliable information dissemination can be achieved in bio-inspired stochastic settings with limited communication and no global control.

We analyze the noisy  $\mathcal{PULL}(h)$  model, covering a general setting that spans from rumor spreading to majority consensus: a subset of source agents hold initial preferences, and the goal is to converge to the majority preference. Agents passively observe noisy messages from  $h$  randomly sampled peers per round. Prior work shows that convergence requires  $\Omega(n/h)$  rounds even under favorable conditions. We ask: how far can one push simplicity—no synchronization and minimal message size—without compromising convergence speed?

We present a quasi self-stabilizing protocol using only 2-bit messages that converges from arbitrary initial states despite severe noise and asynchrony. It achieves optimal convergence time  $O((n/h) \log n)$  with high probability, and  $O(\log n)$  time in the fully connected case  $h = n$ . A key subroutine is an even simpler 1-bit protocol assuming simultaneous start, based on a natural two-phase “listen-then-amplify” mechanism reminiscent of biological strategies.

Together, our results connect biologically inspired heuristics with provable guarantees for robust, efficient information dissemination in highly unreliable and uncoordinated systems.

## KEYWORDS

Information Spreading; Noisy Communication; Zealot Consensus; Self-Stabilization; Natural Algorithms

## ACM Reference Format:

Niccolò D’Archivio, Amos Korman, Emanuele Natale, and Robin Vacus. 2026. Fast and Robust Information Spreading in the Noisy  $\mathcal{PULL}$  Model: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 3 pages. <https://doi.org/10.65109/FFHM1865>

## 1 INTRODUCTION

Designing artificial multi-agent systems is challenging when agents must achieve distributed coordination and consensus under local, asynchronous, and noisy communication constraints. Yet, nature offers inspiration: for example, in ant colonies and fish schools, a few informed individuals can lead an entire uncoordinated group using minimal, noisy signals [8, 15, 16, 19]. This motivates the study of simple, fault-tolerant algorithms for information spreading in settings with unreliable communication and severe bandwidth limits [6, 13].

In the spirit of *opinion dynamics* [4, 9, 10, 14, 18], we consider a population of  $n$  agents, each holding a binary opinion (0 or 1). We assume that a subset of agents, called *sources*, hold initial preferences. We define the correct opinion as the one supported by the majority of sources. The goal is for the entire population to converge on this correct opinion. This formulation generalizes the standard rumor spreading problem (where sources agree on the correct opinion) to include majority consensus (where all agents are sources). We refer to this broader problem as *majority rumor spreading* for simplicity.

*Noise in communication.* To capture random, unpredictable, and noisy communication, we adopt the noisy  $\mathcal{PULL}(h)$  model [6, 7, 12, 13]. In this model, time proceeds in synchronous rounds, and in each round every agent pulls messages from  $h$  randomly chosen peers (neighbors) in the network. Each message is drawn from a fixed alphabet  $\Sigma$ . Crucially, the channel is noisy: the received symbol may be corrupted randomly. We assume the noise is biased, i.e. it retains the original message with probability slightly greater than  $1/|\Sigma|$ . Despite this slight reliability, without further algorithmic design the information can easily get lost in the noise. This raises a central question: How can the agents quickly and reliably agree on the correct opinion under such extreme noise, using small  $|\Sigma|$ ?

One difficulty that arises in the noisy  $\mathcal{PULL}(h)$  model is that sources cannot safely signal the fact that they are sources using



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems ([www.ifaamas.org](http://www.ifaamas.org)). <https://doi.org/10.65109/FFHM1865>

a designated bit, since this bit itself would be noisy. As a consequence, the majority of messages perceived as originating from a source would in fact be artifacts of noise, rather than genuine source messages. Conversely, without the ability to identify sources, non-sources must somehow extract information from multiple samples, most of which are unreliable.

Notably, prior work established a fundamental time barrier: any protocol requires on the order of  $\Omega(n/h)$  rounds to spread the information with constant success probability [6]. In other words, if each agent communicates with only a constant number of others per round ( $h = O(1)$ ), it takes linear time  $\Omega(n)$  in the population size to reliably propagate a rumor; larger  $h$  may potentially linearly speed this up, but no algorithm can circumvent this bound. Our goal is to approach this optimal convergence speed using the smallest alphabet  $\Sigma$ , and to do so under minimal assumptions about coordination among agents.

*Quasi self-stabilization.* To model the lack of global coordination, we consider a variant of the self-stabilization framework [1]. A self-stabilizing protocol ensures convergence from any initial state—potentially set by an adversary—without assuming synchronized clocks or a common start time. However, the traditional model assumes complete adversarial control over all initial internal states, which may be unnecessarily strong for modeling real systems. We therefore introduce a relaxed version – quasi self-stabilization – that allows reliable access to certain information, such as the size of the population  $n$ , the distribution of the noise, or whether an agent is a source.

When agents start the execution at different times, several technical challenges arise. Without synchronous counters, it is difficult to implement phase-based strategies. This asynchrony complicates both correctness and efficiency guarantees: protocols must ensure convergence despite arbitrary misalignment of agents’ clocks, while simultaneously tolerating noisy communication.

Previous studies have explored self-stabilizing rumor spreading in the  $\mathcal{PULL}$  model, focusing on minimizing message size [2, 3, 5, 17], but without addressing noisy communication. Conversely, Boczkowski et al. [6] analyzed the noisy  $\mathcal{PULL}(h)$  model and showed that with any constant-size alphabet, convergence still requires  $\tilde{\Omega}(n/h)$  rounds—even assuming synchronized start. In this work, we address both dimensions of unreliability: noisy messages and arbitrary, uncoordinated initial configurations.

*Main contribution.* We answer the above challenge by designing protocols for noisy majority rumor spreading, emphasizing a tiny message alphabet and robustness to communication noise and/or arbitrary start conditions.

Our main result is a quasi self-stabilizing majority rumor spreading algorithm, called SSF, that uses a 4-symbol alphabet (only 2 bits of information per message) and operates correctly from arbitrary initial states. It achieves convergence in  $O\left(\frac{n}{h} \log n\right)$  rounds w.h.p. In the extreme well-connected case of  $h = n$ , it achieves  $O(\log n)$  rounds, an exponential improvement over naive gossip. For general  $h$ , the upper bound matches the  $\Omega(n/h)$  lower bound up to a logarithmic factor, which could be attributed to high probability guarantee of SSF, versus the constant probability guarantee of the lower bound.

To our knowledge, SSF is the first protocol to solve noisy majority rumor spreading with such a small message alphabet in a quasi self-stabilizing manner, advancing the state of the art in fault-tolerant consensus in multi-agent systems subject to strong noise.

## 1.1 Overview of Key Algorithmic Ideas

SSF builds on a simple 1-bit protocol, SF, which assumes a simultaneous start and achieves fast majority rumor spreading through a two-phase process. SSF extends this method to the fully self-stabilizing setting by embedding lightweight synchronization directly into the protocol, eliminating the need for a global clock.

In Phase 1, agents collect  $h$  samples per round to form a weak opinion—a guess correct with probability slightly above  $1/2$ . This is achieved by having sources continuously broadcast their bit, while non-sources contribute minimal interference. Weak opinions are computed using only first-hand information—direct observations of other agents—to maintain independence and avoid bias from relayed noise.

In Phase 2, agents share their weak opinions and update by majority rule, amplifying the bias to achieve consensus with high probability. This boosting step resembles mechanisms used in prior works on majority dynamics under noise [12, 13].

For such a procedure to work, it is desirable that all weak-opinions are mutually independent, so that the majority of weak-opinions would be correct with high probability. However, having non-sources relay information between themselves can cause dependencies. Therefore, in our algorithms, weak-opinions are computed in a way that is oblivious of the information obtained (and the randomness involved) by other non-sources. Hence, in a sense, the weak-opinions are based on “first-hand” information, that is, information that comes directly from the sources, despite the fact that sources are not easily detectable.

The self-stabilizing protocol SSF implicitly cycles through these phases and uses the additional symbols (beyond binary opinions) to coordinate phase transitions and recover from inconsistent states. Instead of relying on an external synchronization subroutine, as in [5], which may be difficult to implement in the presence of noise, SSF integrates coordination directly into the message structure. The first bit indicates whether the sender is a source (1) or not (0), while the second bit conveys the weak opinion. Of course, upon observation, these bits may be flipped due to noise. Agents compute weak opinions by applying a majority vote over the second bits of messages marked as coming from sources, and update their actual opinion, which is kept internally, via a majority over all received second bits, regardless of the value of their first bit. This design allows SSF to converge in near-optimal time, despite both noise and lack of coordination, using only 2-bit messages.

A full version of this work with complete proofs is available at [11].

**Acknowledgements.** This work has been supported the French government National Research Agency (ANR) through the UCA JEDI (ANR-15-IDEX-01) and EUR DS4H (ANR-17-EURE-004), through the 3IA Cote d’Azur Investments in the project with the reference number ANR-23-IACL-0001, and by the AID INRIA-DGA project n°2023000872 “BioSwarm. A.K was supported by the Israel Science Foundation (<https://www.isf.org.il>) Grant 1574/24.

## REFERENCES

- [1] Karine Altisen, Stéphane Devismes, Swan Dubois, and Franck Petit. 2022. *Introduction to distributed self-stabilizing algorithms*. Springer Nature.
- [2] Paul Bastide, George Giakkoupis, and Hayk Saribekyan. 2021. Self-Stabilizing Clock Synchronization with 1-bit Messages. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, Alexandria, Virginia, U.S., 2154–2173.
- [3] Luca Becchetti, Andrea Clementi, Amos Korman, Francesco Pasquale, Luca Trevisan, and Robin Vacus. 2023. On the Role of Memory in Robust Opinion Dynamics. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI '23)*. Macao, P.R.China, 29–37. <https://doi.org/10.24963/ijcai.2023/4>
- [4] Luca Becchetti, Andrea Clementi, and Emanuele Natale. 2020. Consensus dynamics: An overview. *ACM SIGACT News* 51, 1 (2020), 58–104.
- [5] Lucas Boczkowski, Amos Korman, and Emanuele Natale. 2017. Minimizing Message Size in Stochastic Communication Patterns: Fast Self-Stabilizing Protocols with 3 Bits. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) (Proceedings)*. Society for Industrial and Applied Mathematics, 2540–2559. <https://doi.org/10.1137/1.9781611974782.168>
- [6] Lucas Boczkowski, Emanuele Natale, Ofer Feinerman, and Amos Korman. 2018. Limits on reliable information flows through stochastic populations. *PLoS computational biology* 14, 6 (2018), e1006195.
- [7] Andrea Clementi, Luciano Gualà, Emanuele Natale, Francesco Pasquale, Giacomo Scornavacca, and Luca Trevisan. 2020. Consensus vs Broadcast, with and without Noise. In *ITCS 2020 - 11th Annual Innovations in Theoretical Computer Science (11th Innovations in Theoretical Computer Science Conference)*. Seattle, United States, 42–43. <https://doi.org/10.4230/LIPcs.ITCS.2020.42>
- [8] Iain D Couzin, Jens Krause, Nigel R Franks, and Simon A Levin. 2005. Effective leadership and decision-making in animal groups on the move. *Nature* 433, 7025 (2005), 513–516.
- [9] Emilio Cruciani, Emanuele Natale, André Nusser, and Giacomo Scornavacca. 2018. Phase Transition of the 2-Choices Dynamics on Core-Periphery Networks. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (Stockholm, Sweden) (AAMAS '18)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 777–785.
- [10] Emilio Cruciani, Emanuele Natale, and Giacomo Scornavacca. 2019. Distributed community detection via metastability of the 2-choices dynamics. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence (Honolulu, Hawaii, USA) (AAAI'19/IAAI'19/EAAI'19)*. AAAI Press, Article 742, 8 pages. <https://doi.org/10.1609/aaai.v33i01.33016046>
- [11] Niccolò D'Archivio, Amos Korman, Emanuele Natale, and Robin Vacus. 2024. Fast and Robust Information Spreading in the Noisy PULL Model. arXiv:2411.02560 [cs.DC] <https://arxiv.org/abs/2411.02560>
- [12] Ofer Feinerman, Bernhard Haeupler, and Amos Korman. 2017. Breathe before speaking: efficient information dissemination despite noisy, limited and anonymous communication. *Distributed Computing* 30, 5 (2017), 339–355. <https://doi.org/10.1007/S00446-015-0249-4>
- [13] Pierre Fraigniaud and Emanuele Natale. 2019. Noisy rumor spreading and plurality consensus. *Distributed Computing* 32, 4 (2019), 257–276. <https://doi.org/10.1007/S00446-018-0335-5>
- [14] Marcelo Matheus Gaury, Anna Abramishvili, Eduardo Colli, Tiago Madeira, Fredrik Mallmann-Trenn, Vinicius Franco Vasconcelos, and David Kohan Marzagao. 2025. Voter Model Meets Rumour Spreading: A Study of Consensus Protocols on Graphs with Agnostic Nodes. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (Detroit, MI, USA) (AAMAS '25)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 848–857.
- [15] Aviram Gelblum, Itai Pinkoviezky, Ehud Fonio, Abhijit Ghosh, Nir Gov, and Ofer Feinerman. 2015. Ant groups optimally amplify the effect of transiently informed individuals. *Nature communications* 6, 1 (2015), 7729.
- [16] Aviram Gelblum, Itai Pinkoviezky, Ehud Fonio, Nir S Gov, and Ofer Feinerman. 2016. Emergent oscillations assist obstacle negotiation during ant cooperative transport. *Proceedings of the National Academy of Sciences* 113, 51 (2016), 14615–14620.
- [17] Amos Korman and Robin Vacus. 2022. Early adapting to trends: Self-stabilizing information spread using passive communication. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, 235–245.
- [18] Hicham Lesfari, Frédéric Giroire, and Stéphane Pérennes. 2022. Biased Majority Opinion Dynamics: Exploiting Graph k-Domination. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, Vienna, Austria, 377–383. <https://doi.org/10.24963/ijcai.2022/54>
- [19] Nitzan Razin, Jean-Pierre Eckmann, and Ofer Feinerman. 2013. Desert ants achieve reliable recruitment across noisy interactions. *Journal of the Royal Society Interface* 10, 82 (2013), 20130079.