


# Brief Announcement: Simple Majority Consensus in Networks with Unreliable Communication

Ariel Livshits ✉ 

Technion – Israel Institute of Technology, Haifa, Israel

Yonatan Shadmi ✉ 

Technion – Israel Institute of Technology, Haifa, Israel

Ran Tamir (Averbuch) ✉ 

Technion – Israel Institute of Technology, Haifa, Israel

---

## Abstract

In this work, we consider a synchronous model of  $n$  faultless agents, with a complete communication graph and messages that are lost with some constant probability  $q \in (0, 1)$ . In this model we show that there exists a protocol, called the Simple Majority Protocol, that solves consensus in 3 communication rounds with probability of agreement converging to 1 as  $n \rightarrow \infty$ . We also prove that 3 communication rounds are necessary for the SMP to achieve consensus, with high probability.

**2012 ACM Subject Classification** Theory of computation  $\rightarrow$  Distributed algorithms

**Keywords and phrases** Majority consensus, probabilistic message loss, distributed systems

**Digital Object Identifier** 10.4230/LIPIcs.DISC.2021.59

**Related Version** Simple Majority Consensus in Networks with Unreliable Communication

*Full Version:* <https://arxiv.org/pdf/2104.04996.pdf> [6]

**Acknowledgements** The authors thank Yoram Moses for his invaluable comments and suggestions.

## 1 Introduction

In the binary consensus problem, every agent is initially assigned some binary value, referred to as the agent’s initial value. The goal of a protocol that solves consensus is to have every agent eventually decide on the same value, thus reaching agreement throughout the system. More formally, given any initial configuration of the agents, every run of a protocol that solves consensus must exhibit these three properties: *Decision* – every agent eventually decides on some value  $v \in \{0, 1\}$ , *Agreement* – if some agent has decided on  $v$ , no value other than  $v$  can be decided on by any other agent, and *Validity* – if some agent has decided on  $v$ , then  $v$  was initially assigned to some agent.

In the pursuit of developing distributed protocols for consensus, much of the literature routinely makes two powerful assumptions. The first is that communication links are reliable, and the second is that there exists an upper bound on the transmission delay of messages from one agent to another [1]. Nonetheless, communication networks are notoriously unreliable in many practical systems [2]. In fact, actual communication links may suffer from sudden crashes, resulting in messages in transit being lost forever.

In practice, the reliability and upper bound assumptions do not hold simultaneously. However, it has been proven that consensus is not solvable without assuming reliable communication [5], and without a bound on message delivery times, protocols solving consensus cannot provide any bound on when they halt. Alternatively, we may consider a model where there exists a stationary probability distribution on the event that messages are delivered by a certain deadline, but in such a model we may only provide probabilistic



© Ariel Livshits, Yonatan Shadmi, and Ran Tamir (Averbuch);  
licensed under Creative Commons License CC-BY 4.0

35th International Symposium on Distributed Computing (DISC 2021).

Editor: Seth Gilbert; Article No. 59; pp. 59:1–59:4



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

guarantees for consensus. Fortunately, the growth in scale of modern distributed systems (e.g. cryptocurrency) has provided ample incentive for work in probabilistic models, given that large numbers naturally lend themselves well to probabilistic analysis.

We consider a model of  $n$  agents connected via a complete communication graph. Namely, we assume that each agent has an active communication channel to every other agent in the system. This is a valid approximation for unstructured overlays in peer to peer networks, e.g., Freenet, Gnutella and Fast Track [4]. Also, we assume that each agent has access to a synchronized global clock. This permits time to be split into equal-length communication rounds, in which all messages between agents are sent at the beginning of a communication round, and either arrive by the end of the round or are considered lost. Additionally, we assume that the agents themselves are faultless, and that all message loss events are i.i.d. with some constant probability  $q \in (0, 1)$ .

The Simple Majority Protocol (SMP) is a round-based, majority-rule protocol that attempts to solve majority consensus, a stronger variation of consensus, in which the validity clause stipulates that if a majority of agents were initially assigned the same value, then all agents must decide on this value. The SMP can be described as follows: In each round, every agent sends its current value to all other agents. Then, it waits to receive messages from all other agents. If a majority of received messages propose the same value, then the agent adopts this value for the next round. All ties are reconciled by readopting the agent's current value. After a fixed number of rounds  $r$ , each agent decides on its currently adopted value.

In this work, we prove that the SMP solves consensus in 3 communication rounds with probability of agreement converging to 1 as  $n \rightarrow \infty$  (i.e., with high probability, denoted by *w.h.p.*). We also prove that if the relative majority (defined in this work as the excess over 50%) of the initial values of agents is of the order of at least  $\sqrt{n}$ , then the SMP solves majority consensus *w.h.p.* in at most 2 communication rounds. Additionally, we show that 3 communication rounds are not only sufficient, but also necessary for the SMP to achieve consensus, with high probability.

## 2 Main Results

We begin by defining a set of useful notations. Let  $S_n$  be a system of  $n$  agents, and let  $c(S_n)$  be the set of all  $2^n$  possible configurations of  $S_n$ . Let  $\sigma = (c_1, c_2, \dots, c_n, \dots)$  be an infinite sequence of configurations, where the  $k$ -th element of the sequence is some configuration of the system  $S_k$ , i.e.,  $\forall k \geq 1 : c_k \in c(S_k)$ . Let  $I_0(c_n)$  and  $I_1(c_n)$  be the number of zeros and ones, respectively, in the configuration  $c_n$ . The relative majority of a configuration is then defined as  $\delta(c_n) = \max\{I_0(c_n), I_1(c_n)\} - \lfloor \frac{n}{2} \rfloor$ , i.e., the excess over 50% of the majority value. We will also refer to the sequence of relative majorities  $\delta(\sigma) = (\delta(c_1), \delta(c_2), \dots, \delta(c_n), \dots)$ , in which the  $k$ -th element is the relative majority of the  $k$ -th configuration in  $\sigma$ . We say that  $\delta(\sigma) \in \Omega(\sqrt{n})$  *w.h.p.* holds for some sequence  $\sigma = (c_1, c_2, \dots, c_n, \dots)$ , if for any  $\epsilon > 0$  there exists some constant  $a > 0$  and index  $k \in \mathbb{N}$  such that  $\forall n \geq k : \mathbb{P}(\delta(c_n) \geq a \cdot \sqrt{n}) \geq 1 - \epsilon$ . Similarly,  $\delta(\sigma) \in \omega(\sqrt{n})$  *w.h.p.* holds if  $\forall n \geq k : \mathbb{P}(\delta(c_n) > a \cdot \sqrt{n}) \geq 1 - \epsilon$ . We say that  $\delta(\sigma) \in \Omega(\sqrt{n})$  (or  $\delta(\sigma) \in \omega(\sqrt{n})$ ) if this holds also for  $\epsilon = 0$ .

The SMP defines *a priori* the number of rounds  $r$  until termination. Denote the event that the SMP reaches consensus from an initial configuration  $c_n$  in  $r$  rounds by  $\mathcal{C}(c_n, r)$ . Similarly, we denote by  $\mathcal{C}^m(c_n, r)$  the event that the SMP reaches majority consensus. The probability spaces over which our results hold are  $\{(R(c_n), 2^{R(c_n)}, \mathbb{P})\}_{n \geq 1}$  where  $R(c_n)$  is the set of all runs of a system of  $n$  agents, that start in the initial configuration  $c_n$ , and  $\mathbb{P}(\cdot)$  is the natural probability measure on runs induced by the distribution on message loss, i.e. Bernoulli with parameter  $q$  (as done in, e.g., [3]).

► **Lemma 1.** *For any sequence of initial configurations  $\sigma = (c_1, c_2, \dots, c_n, \dots)$ , let  $\tilde{\sigma} = (\tilde{c}_1, \tilde{c}_2, \dots, \tilde{c}_n, \dots)$  be the next sequence of configurations reached after a single round of the SMP. Then,  $\delta(\tilde{\sigma}) \in \Omega(\sqrt{n})$  w.h.p. holds for  $\tilde{\sigma}$ .*

The lemma states that as the scale of the system approaches infinity, the relative majority of the configuration reached after a single round of the SMP is of the order of at least  $\sqrt{n}$ , regardless of the initial configuration of the system. To gain an intuitive understanding of this result, observe the worst case scenario. Specifically, consider the case where the initial configuration of the system is perfectly symmetric, i.e., the number of ones equals the number of zeros. In this case, after a single communication round, every agent receives approximately the same number of messages reporting zeros and ones. Thus, we assume that the configuration of the agents at the end of the first round is approximately an i.i.d. random vector whose components are Bernoulli random variables with parameter 0.5. The central limit theorem states that  $\frac{1}{\sqrt{n}} \cdot \delta(\tilde{c}_n)$  converges in distribution to a Gaussian random variable. This means that  $\frac{1}{\sqrt{n}} \cdot \delta(\tilde{c}_n) \in \Theta(1)$  w.h.p., and hence  $\delta(\tilde{c}_n) \in \Theta(\sqrt{n})$  w.h.p., which implies that  $\delta(\tilde{\sigma}) \in \Omega(\sqrt{n})$  w.h.p. (see Proposition 3 in [6] for a formal proof).

► **Lemma 2.** *Let  $\sigma = (c_1, c_2, \dots, c_n, \dots)$  be a sequence of initial configurations. Then:*

- *If  $\delta(\sigma) \in \Omega(\sqrt{n})$ , then  $\mathbb{P}[\mathcal{C}^m(c_n, 2)] \xrightarrow{n \rightarrow \infty} 1$*
- *If  $\delta(\sigma) \in \omega(\sqrt{n})$ , then  $\mathbb{P}[\mathcal{C}^m(c_n, 1)] \xrightarrow{n \rightarrow \infty} 1$*

The lemma implies that for initial configurations that, to begin with, have a significant relative majority to one of the values, then w.h.p. the SMP reaches majority consensus in one or two communication rounds, depending on how much larger the initial relative majority is than  $\sqrt{n}$ . However, majority consensus cannot be ensured w.h.p. for all possible initial configurations. Intuitively, if the initial relative majority is too weak (e.g., on the order of  $\log(n)$ ), then it is most likely that the random losses in the network will completely hide it. In fact, the state at the end of round 1 is probabilistically equivalent to a sequence of  $n$  fair coin tosses, and hence, the majority at the end of round 1 will switch sides with a probability of about one half (see Propositions 1 & 2 in [6] for a formal proof). Our main result is:

► **Theorem 3.** *Let  $\sigma = (c_1, c_2, \dots, c_n, \dots)$  be a sequence of initial configurations matching to a sequence of systems with  $n$  agents that grow in size as  $n \rightarrow \infty$ . Then:  $\mathbb{P}[\mathcal{C}(c_n, 3)] \xrightarrow{n \rightarrow \infty} 1$ .*

This theorem asserts that for a sequence of systems of  $n$  agents that progressively grow larger with  $n$ , the SMP will reach agreement w.h.p., after only 3 communication rounds, regardless of the initial configurations of those systems. This result seems intuitive for initial configurations that already possess a large relative majority towards some value, considering that it is very likely that the SMP will cause a small minority of agents to adopt the value of the majority. However, it is surprising that this result holds for initial configurations whose relative majority is close to 0, especially if the initial configuration is perfectly symmetric. However, Lemma 1 maintains that even in the worst case, the symmetry will break equiprobably to one of the values after a single round of the SMP, the relative majority will be of the order of at least  $\sqrt{n}$ , and then, by Lemma 2, consensus will be reached in at most 2 additional rounds (see Theorem 1 in [6] for a formal proof).

► **Theorem 4.** *Let  $\sigma = (c_1, c_2, \dots, c_n, \dots)$  be a sequence of initial configurations such that for all  $c_n \in \sigma$  it holds that  $I_0(c_n) = I_1(c_n)$ . Then:  $\mathbb{P}[\mathcal{C}(c_n, 2)] \xrightarrow{n \rightarrow \infty} 0$ .*

The theorem provides a lower bound of 2 rounds for the SMP to reach consensus. Specifically, if the initial configuration of the system is perfectly symmetric, then 3 rounds are not only sufficient, but also necessary (see Theorem 2 in [6] for a formal proof).

---

References

---

- 1 Marcos K Aguilera. Stumbling over consensus research: Misunderstandings and issues. In *Replication*, pages 59–72. Springer, 2010.
- 2 Rachid Guerraoui, Michel Hurfin, Achour Mostéfaoui, Riucarlos Oliveira, Michel Raynal, and André Schiper. Consensus in asynchronous distributed systems: A concise guided tour. In *Advances in Distributed Systems*, pages 33–47. Springer, 2000.
- 3 Joseph Y Halpern and Mark R Tuttle. Knowledge, probability, and adversaries. *Journal of the ACM (JACM)*, 40(4):917–960, 1993.
- 4 Eng Keong Lua, Jon Crowcroft, Marcelo Pias, Ravi Sharma, and Steven Lim. A survey and comparison of peer-to-peer overlay network schemes. *IEEE Communications Surveys & Tutorials*, 7(2):72–93, 2005.
- 5 Nancy A Lynch. *Distributed algorithms*. Elsevier, 1996.
- 6 Ran Tamir, Ariel Livshits, and Yonatan Shadmi. Simple majority consensus in networks with unreliable communication. *arXiv preprint*, 2021. [arXiv:2104.04996](https://arxiv.org/abs/2104.04996).