



# Stabilizing Consensus Is Impossible in Lossy Iterated Immediate Snapshot Models

Stephan Felber  

Vienna University of Technology, Austria

Hugo Rincon Galeana  

Vienna University of Technology, Austria

Berlin University of Technology, Germany

---

## Abstract

A substantial portion of distributed computing research is dedicated to terminating problems like consensus and similar agreement problems. However, non-terminating problems have been intensively studied in the context of self-stabilizing distributed algorithms, where processes may start from arbitrary initial states and can tolerate arbitrary transient faults. In between lie stabilizing problems, where the processes start from a well-defined initial state, but do not need to decide irrevocably and are allowed to change their decision finitely often until a stable decision is eventually reached.

Stabilizing consensus has been studied within the context of synchronous message adversaries. In particular, Charron-Bost and Moran showed that a necessary condition for stabilizing consensus is the existence of at least one process that reaches all others infinitely often (a perpetual broadcaster). However, it was left open whether this is also a sufficient condition for solving stabilizing consensus.

In this paper, we introduce the novel Delayed Lossy-Link (DLL) model, and the Lossy Iterated Immediate Snapshot Model (LIIS), for which we show stabilizing consensus to be impossible. The DLL model is introduced as a variant of the well-known Lossy-Link model, which admits silence periods of arbitrary but finite length. The LIIS model is a variant of the Iterated Immediate Snapshot (IIS) model, which admits finite length periods of at most  $f$  omission faults per layer. In particular, we show that stabilizing consensus is impossible even when  $f = 1$ .

Our results show that even in a model with very strong connectivity, namely, the Iterated Immediate Snapshot (IIS) model, a single omission fault per layer effectively disables stabilizing consensus. Furthermore, since the DLL model always has a perpetual broadcaster, the mere existence of a perpetual broadcaster, even in a crash-free setting, is not sufficient for solving stabilizing consensus, negatively answering the open question posed by Charron-Bost and Moran.

**2012 ACM Subject Classification** Computer systems organization → Fault-tolerant network topologies; Networks → Network properties; Theory of computation → Distributed algorithms; Theory of computation → Randomness, geometry and discrete structures

**Keywords and phrases** distributed systems, dynamic networks, dynamic graphs, message adversaries, stabilizing consensus, asynchronous message passing

**Digital Object Identifier** 10.4230/LIPIcs.OPODIS.2024.18

**Funding** This work was supported by the DMAC project (P32431).

*Hugo Rincon Galeana*: supported by German Research Foundation (DFG), Schwerpunktprogramm (SPP 2378), project ReNO, 2023-2027, and the financial support of Najla Amira Ochoa Leonor.

**Acknowledgements** We would like to thank Ulrich Schmid and Kyrill Winkler for their helpful remarks, which undoubtedly improved this paper.

## 1 Introduction

Agreement tasks, and in particular consensus, have always been a focal point of distributed computing research, not only because of their practical applicability, but also because consensus tasks characterize very precisely the limits of a distributed computing model.



© Stephan Felber and Hugo Rincon Galeana;  
licensed under Creative Commons License CC-BY 4.0

28th International Conference on Principles of Distributed Systems (OPODIS 2024).

Editors: Silvia Bonomi, Letterio Galletta, Etienne Rivière, and Valerio Schiavoni; Article No. 18; pp. 18:1–18:16

Leibniz International Proceedings in Informatics



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

Understanding such limitations not only facilitates the quick assessment of problem solvability in a given model, but also sheds light on the impact that certain properties, such as synchrony, impose on the system. For instance, the celebrated FLP impossibility result [15] reveals the devastating effect of both asynchrony and process crashes on deterministic consensus solvability. In a similar vein, it has been shown that process crashes impair general task solvability [21] even in asynchronous shared memory systems, and that byzantine faults [22] and/or message loss [25, 26, 11, 29, 30] impair consensus solvability even in the context of synchronous message passing systems.

Although distributed systems research has focused extensively on *terminating* tasks, there exist interesting applications and problems that are inherently *non-terminating*. Apart from the specific class of self-stabilizing distributed algorithms [12, 14], which can even tolerate massive transient faults, asymptotic consensus [4, 5], stabilizing consensus [2, 9, 27], and approximate consensus [8, 13] are examples of non-terminating distributed tasks: processes are not required to terminate after having computed some final value, but rather eventually converge to some stable configuration. Such tasks are not only of theoretical interest, but are also essential for implementing practical distributed problems such as clock synchronization [23, 28] and sensor fusion [3].

In this paper, we focus on stabilizing consensus, which can be viewed as the non-terminating variant of consensus. As in the case of consensus, all processes start with their own input values and must *eventually* agree on a common value. The fundamental difference to terminating consensus is that every process may start arbitrarily late, and does not need to decide on a value irrevocably and exactly once, but can change its decision value *finitely often*. Nevertheless, all processes must eventually stabilize on the *same* decision value.

It should be noted that stabilizing consensus is also loosely related to asymptotic consensus [16], in the sense that, in neither problem, processes are required to terminate. A fundamental difference, however, is that asymptotic consensus allows processes to decide from a real-valued range, in contrast to the validity condition for stabilizing consensus, which only allows decisions from the pool of input values. Stabilizing consensus is a stronger problem, since any protocol that solves stabilizing consensus also solves asymptotic consensus. Indeed, we provide a simple novel model introduced below (the DLL), where stabilizing consensus is impossible (see Theorem 3.10), yet asymptotic consensus can be solved. In this light, asymptotic consensus is more closely related to approximate consensus [8], while stabilizing consensus is more closely related to consensus.

## 1.1 Related work

We study stabilizing consensus in the *Synchronous Message Passing* setting where processes communicate through uni-directional links in a round by round fashion and an *adversary* suppresses certain links in every round. More specifically, the communication in every round is modeled as a directed graph sequence, which is under the control of a *message adversary* [1]. Charron-Bost and Moran [9] provided a class of algorithms in this setting, called the *MinMax Algorithms*, solving stabilizing consensus under any message adversary that generates graph sequences adhering to two constraints: (i) there exists a process (a perpetual broadcaster) that is able to reach all other process (possibly via multiple hops) infinitely often, and (ii) the broadcasting time (in terms of rounds) for doing so is bounded by an unknown constant. The authors showed that (i) is necessary for solving stabilizing consensus, but leave the question open whether or not (i) is also sufficient.

In this paper, we extend the Lossy-Link model, introduced by Santoro and Widmayer [25], where two processes are connected by a pair of links that may drop at most one (directional) message per round, to the novel *Delayed Lossy-Link* (DLL) model, where both messages may

be lost but (some) messages are guaranteed to be transmitted infinitely often. Whereas stabilizing consensus can be solved in the Lossy-Link model, we show that stabilizing consensus is impossible in the DLL model, thus negatively answering Charron-Bost and Moran’s open question in Theorem 4.8. Our proof vaguely resembles the bivalency proof used for showing the impossibility of terminating consensus in the Lossy-Link model [25], in the sense that we construct a forever *conflicted* run, i.e., a run where processes always decide on invalid output configurations, by extending conflicted prefixes. Whereas a bivalent prefix leads to more than one decision value eventually, a conflicted prefix already yields an output configuration that violates agreement. A bivalency argument is not enough to prove stabilizing consensus impossible, however.

Afek and Gafni [1] showed that asynchronous wait-free shared memory models such as the *Iterated Immediate Snapshot* (IIS) model can also be represented by a synchronous message adversary. We extend the message adversary formulation of the IIS model by adding up to  $f$  read-omission faults per round to the snapshot operations. We call this new model the *Lossy Iterated Immediate Snapshot model* LIIS( $f$ ). In particular, in the case of only 2 processes, LIIS(1) corresponds to the DLL model. The LIIS( $f$ ) model is loosely related to the  $d$ -solo models studied by Herlihy, Rajsbaum, Raynal and Stainer [20], as the presence of read omissions may lead to  $d$ -solo rounds, i.e., rounds where up to  $d$  processes do not receive any information from the rest of the processes.

## 1.2 Contributions and paper organization

In Section 2, we provide the preliminary definitions and the overall framework that are essential for our results. In Section 3, we present the Lossy-Link model, and introduce the Delayed Lossy-Link model for which we prove stabilizing consensus to be impossible, namely Theorem 3.10. This is our first main result and constitutes the foundation for the rest of the paper. In Section 4, we present the IIS model and our novel LIIS( $f$ ) model, and show that stabilizing consensus is not solvable in LIIS( $f$ ) for any  $f \geq 1$ , via Theorem 4.8. Finally, in Section 5 we discuss the implications of our result and provide some interesting directions for future work.

## 2 The System Model

We consider a finite set of process  $\Pi = \{p_1, \dots, p_n\}$  that communicate through message passing via directional point-to-point links in lock-step synchronous rounds. We assume that rounds are communication-closed, i.e., messages sent at a round  $r$  will also arrive at round  $r$  or will not be delivered. Note that in contrast to [9] we assume synchronous process starts (i.e., starting round  $s_p = 1$  for all  $p$ ) as this special case already facilitates our impossibility result. Obviously, our result then carries over to the more general model presented in the original paper.

Each process  $p_i$  is given an initial input value  $x_i$  from a finite set of inputs  $I$ . We say that an initial global configuration is an  $n$ -tuple  $(p_i, x_i)_{i=1}^n$  of process-input pairs. We denote the set of possible *input configurations* as  $\mathcal{I}$ . Similarly, an *output configuration* is an  $n$ -tuple  $(p_i, y_i)_{i=1}^n$ . We denote the set of possible output configurations by  $\mathcal{O}$ .

We model distributed problems as input/output relations called *tasks* that map a set of input configurations  $\mathcal{I}$  to a set of valid output configurations. Formally, a task is a triple  $T = \langle \mathcal{I}, \mathcal{O}, \Delta \rangle$  where  $\Delta : \mathcal{I} \rightarrow 2^{\mathcal{O}}$  is a validity map that determines a subset of valid output configurations from a single input configuration. Tasks are widely used in the literature to capture problems in distributed setting, see for example [24, 17, 18].

The communication is governed by a message adversary that determines which processes are able to communicate at a given round. We use a directed graph  $G_r$  on  $\Pi$ , which we call the *communication graph*, to represent the message exchange at round  $r$ . We denote by  $In_G(p_i)$  the in-neighborhood of  $p_i$  in graph  $G$ , which always includes  $p_i$ . Note that  $In_{G_r}(p_i)$  represents the set of processes from which  $p_i$  receives a message at round  $r$ . We represent the communication in an execution by a communication graph sequence  $\sigma = (G_i)_{i=1}^\infty$ , called a *communication pattern*, where  $\sigma(k) = G_k$ . We denote by  $G \cdot \sigma$  the concatenation of  $G$  with  $\sigma$  and by  $G^k$  the sequence of  $G$  repeated  $k$  times. If  $\sigma = (G_i)_{i=1}^\infty$  is a communication pattern, we say that  $(G_i)_{i=1}^k$  is the  $k$ -*prefix* of  $\sigma$ , denoted by  $\sigma|_1^k$ . Conversely,  $\sigma|_k^\infty = (G_i)_{i=k}^\infty$  denotes the suffix starting from round  $k$ . We define  $\mu = (H_i)_{i=1}^\infty$  to be a sub-sequence of  $\sigma = (G_i)_{i=1}^\infty$ , denoted by  $\mu \sqsubseteq \sigma$ , iff there is a strictly increasing sequence  $s : \mathbb{N} \rightarrow \mathbb{N}$  such that  $H_i = G_{s(i)}$ . Note that the sub-sequence relation  $\sqsubseteq$  is a partial order on sequences with the same domain and co-domain. A *message adversary* is just a set of communication patterns  $\mathcal{M}$ .

Since our focus is on solvability, we assume for convenience that the processes execute a *full-information flooding protocol*. That is, at any given round  $r$ , each process will send a copy of its local state, also called *view* to all possible recipients. For terminating protocols and terminating tasks, we consider strong termination, i.e., a process stops its execution whenever a decision is made by the protocol. Thus, processes will stop sending messages and updating their views once the protocol has output a decision value. In contrast, for stabilizing tasks and protocols, we consider that the protocol never terminates and processes keep communicating and updating their views forever.

We define the *view* of a process  $p_i$  at the end of a round  $r$  under  $\sigma = (G_i)_{i=1}^\infty$  and initial configuration  $\mathcal{C}^0$  as:  $view_{\mathcal{C}^0:\sigma}(p_i, r) := \langle V(\mathcal{C}^0 : \sigma, p_i, r), p_i, r \rangle$ , where  $V(\mathcal{C}^0 : \sigma, p_i, r) = \{view_{\mathcal{C}^0:\sigma}(p_j, r-1) \mid p_j \in In_{\sigma_{r-1}}(p_i)\}$ . We define  $view_{\mathcal{C}^0:\sigma}(p_i, 0) := \langle \{x_i\}, p_i, 0 \rangle$ , where  $x_i$  is the input value of process  $p_i$  in  $\mathcal{C}^0$ .

We also employ a *subview* relation, denoted by  $view_{\mathcal{C}^0:\sigma}(p_i, r) \prec view_{\mathcal{C}^0:\sigma}(p_j, s)$ , which can be defined inductively:  $view_{\mathcal{C}^0:\sigma}(p_j, s) \prec view_{\mathcal{C}^0:\sigma}(p_i, r)$  iff  $s = r-1$  and  $view_{\mathcal{C}^0:\sigma}(p_j, s) \in V(\mathcal{C}^0 : \sigma, p_i, r)$ ; or  $view_{\mathcal{C}^0:\sigma}(p_j, s) \prec view_{\mathcal{C}^0:\sigma}(p_k, r-1)$  for some  $view_{\mathcal{C}^0:\sigma}(p_k, r-1) \in V(\mathcal{C}^0 : \sigma, p_i, r)$ . Intuitively, this subview relation captures all the previous local states that a given process  $p_i$  is aware of at a given round  $r$ .

We denote by  $inputs(view_{\mathcal{C}^0:\sigma}(p_i, r)) := \{x_j \mid \langle \{x_j\}, p_j, 0 \rangle \prec view_{\mathcal{C}^0:\sigma}(p_i, r)\}$  the set of inputs  $p_i$  has heard of by round  $r$ . Whenever  $\mathcal{C}^0$  is clear from the context or not relevant, we will omit it in favor of notational simplicity.

We define the kernel  $Ker(\sigma)$  of a sequence  $\sigma$  as the set of processes that reach all other processes infinitely often, i.e., processes in  $Ker(\sigma)$  are able to disseminate their local state at any point in the run. Formally,  $p \in Ker(\sigma)$  if and only if  $\forall r > 0 : \exists r' > r : \forall q \in \Pi : view_\sigma(p, r) \prec view_\sigma(q, r')$ .

We define the *global configuration* at the end of round  $r$  under a communication pattern  $\sigma$  as an  $n$ -tuple  $\mathcal{C}_\sigma^r := (view_\sigma(p_i, r))_{i=1}^n$ .

As we only consider deterministic protocols, global configurations are fully determined by the input configuration, and the communication pattern  $\sigma$ . Therefore, in the context of this paper, we will consider that a run consists of the input configuration and communication pattern. Furthermore, whenever it is clear from the context, or the input configuration is not relevant, we will use communication pattern and run interchangeably.

We use views to determine indistinguishability between global configurations. More precisely, a global configuration  $\mathcal{C}_{\alpha:\sigma}^r$  is *indistinguishable* from a global configuration  $\mathcal{C}_{\beta:\mu}^r$  for a process  $p_i \in \Pi$  iff  $view_{\alpha:\sigma}(p_i, r) = view_{\beta:\mu}(p_i, r)$ , where  $\alpha, \beta \in \mathcal{I}$  are input configurations, and  $\sigma, \mu \in \mathcal{M}$  are communication patterns.

Note that any protocol can be split into a full-information flooding part that generates views and sends messages, and a decision map that produces an output value. If the output values form a valid configuration, then we say that the protocol solves a problem.

Formally, we define a *decision map* of a protocol  $\mathcal{P}$  as a function that maps process views generated by the full information flooding protocol to output values. We will denote a decision map as  $\delta_{\mathcal{P}} : Views(\mathcal{M}) \rightarrow O$ , where  $Views(\mathcal{M})$  is the set of possible views induced by the message adversary  $\mathcal{M}$  under the protocol  $\mathcal{P}$ , and  $O$  is the set of possible output values. Note that for terminating protocols, an additional undecided value  $\perp$  may be used for representing that the protocol is not yet ready to decide. However, since we are focusing only on stabilizing protocols, we require that a decision is made at every view.

We define the *stabilization* property for tasks, which can be thought of as a relaxation of termination. Whereas termination requires a process to irrevocably choose its output value once and possibly even finalize its execution, stabilization allows the protocol to run indefinitely and correct its output finitely often.

Nevertheless processes must eventually *stabilize* on a common constant output value. We call the round on which a process  $p_i$  stabilizes, i.e., never changes its output again, the *stabilization round*  $s_i$ . It should be noted that stabilization may be reached agnostically. This means that there might be protocols where processes are guaranteed to stabilize, yet they are possibly never aware that they have already reached a stable state.

We model stabilizing problems as tasks, with the following definition of solvability:

► **Definition 2.1** (Stabilizing Task Solvability). *A protocol  $\mathcal{P}$  solves a stabilizing task  $T = \langle \mathcal{I}, \mathcal{O}, \Delta \rangle$  iff for any run  $\mathcal{C}^0 : \sigma$ , there exists an output configuration  $(p_i, y_i)_{i=1}^n \in \Delta(\mathcal{C}^0)$  and a stabilization round  $s_i$  for every process  $p_i$ , such that:*

$$\forall r \geq s_i : \delta_{\mathcal{P}}(\text{view}_{\mathcal{C}^0, \sigma}(p_i, r)) = y_i \text{ where } y_i \text{ denotes the output in } (p_i, y_i)_{i=1}^n$$

A protocol that solves a stabilizing task is called a *stabilizing protocol*.

## 2.1 Consensus and stabilizing consensus

Consensus is traditionally defined via the following three properties.

1. **Termination:** Any process  $p_i$  eventually irrevocably decides on an output value  $y_i$  at a round  $f_i$ .
2. **Validity:** If a process  $p_i$  decides  $y_i$ , then  $y_i$  was the input value of some process  $p_j$ , i.e.  $y_i = x_j$ .
3. **Agreement:** If process  $p_i$  decides on an output value  $y_i$ , and process  $p_j$  decides  $y_j$ , then  $y_i = y_j$ .

Note that we assume that  $I$  is finite and that any combination of input configurations  $\mathcal{C}^0 \in I^{\Pi}$  is possible. Therefore, the above conditions do not define a single unique task, but rather a family of tasks. For instance, binary consensus (where  $|I| = 2$ ), multi-valued consensus (where  $|I| > 2$ ) correspond to different tasks, while both are instances of consensus.

Stabilizing consensus is defined by replacing the termination property **1.** by the weaker condition **1'**. Since stabilizing protocols are allowed to decide at each round, the agreement condition **3.** also needs to be adjusted to **3'**.

- 1'. **Stabilization:** Any process  $p_i$  eventually stabilizes on some valid output value  $y_i$  after some round  $s_i$ .
- 3'. **Stable Agreement:** If process  $p_i$  stabilizes on an output value  $y_i$ , and another process  $p_j$  stabilizes on  $y_j$ , then  $y_i = y_j$ .

Note that a task does not consider any termination or stabilization condition. Therefore, the stabilizing consensus problem, as well as terminating consensus, is determined by a consensus task and a stabilizing (respectively terminating) condition.

► **Definition 2.2** (Consensus Task). *We define the consensus task with respect to an input configuration set  $\mathcal{I}$  as  $T_{con(\mathcal{I})} := \langle \mathcal{I}, \mathcal{O}_{con(\mathcal{I})}, \Delta_{con(\mathcal{I})} \rangle$ .*

*Possible output configurations are vectors satisfying agreement  $\mathcal{O}_{con(\mathcal{I})} := \{(p_i, y_i)_{i=1}^n \mid y_i = v \in I(\mathcal{I})\}$ , where  $I(\mathcal{I})$  denotes the set of input values induced by the input configurations  $\mathcal{I}$ , and the task map preserves validity  $\Delta_{con(\mathcal{I})}((p_i, x_i)_{i=1}^n) := \{(p_i, y)_{i=1}^n \mid y = x_j \text{ for } j \in \{1, \dots, n\}\}$ .*

### 3 The Lossy-Link, and Delayed Lossy-Link Message Adversaries

We will first introduce two particular message adversaries defined for only two processes, namely *Lossy-Link* (denoted by LL) and *Delayed Lossy-Link* (denoted by DLL). Lossy-Link was introduced by Santoro and Widmayer [25] and revisited in [26, 11], where it was shown that consensus is impossible even if at most a single message may be lost in every round. We will show that stabilizing consensus is solvable in the LL but impossible in the DLL.

Lossy-Link consists of 2 processes that communicate through a bi-directional link that may lose at most one message per round. For readability purposes, throughout this section we will denote the set of processes as  $\Pi = \{\circ, \bullet\}$ .

► **Definition 3.1** (Lossy-Link message adversary). *We define the LL as<sup>1</sup>:*

$$LL := \{\circ \rightarrow \bullet, \circ \leftarrow \bullet, \circ \leftrightarrow \bullet\}^\omega.$$

Although the Lossy-Link message adversary prohibits solving terminating consensus [25], it admits a simple stabilizing protocol that solves stabilizing consensus as it satisfies both that any LL pattern has a non empty kernel, and any LL pattern has a trivially bounded broadcast time. Hence, we provide a 2-process instance of the MinMax Algorithm, introduced in [9], for stabilizing consensus. Since we adapt it to our full-information and flooding model, it suffices to provide a decision map  $\delta_{MinMax} : Views(LL) \rightarrow O$ ,

$$\delta_{MinMax}(view_\sigma(p_i, r)) := \max_{x \in V(view_\sigma(p_i, r))} \{\min(inputs(x))\}$$

We omit a proof (for a complete formal treatment see [9]), but provide a sketch. The case where both have identical inputs is trivial, so assume w.l.o.g. that  $x_\circ < x_\bullet$ . We distinguish three cases, either (a)  $\bullet$  never hears from  $\circ$ , or conversely (b)  $\circ$  never hears from  $\bullet$ , or (c) both hear from each other. In (c) both eventually know both values after some round  $r$ . They then always choose  $x_\circ$ , as its the smallest value they know and as its the smallest value in any message they will receive, starting from round  $r + 1$ .

Case (a) is only possible in  $(\circ \leftarrow \bullet)^\omega$ , so  $\circ$  always hears from and chooses  $x_\bullet$ , as its the maximum heard of in the last round ( $\bullet$  also chooses  $x_\bullet$  by validity). Similarly, case (b) is only possible in  $(\circ \rightarrow \bullet)^\omega$ , where after round 1 the minimum  $\bullet$  has heard of is  $x_\circ$ , and  $\bullet$  will never receive a larger value, thus always chooses  $x_\circ$  ( $\circ$  chooses the same by validity).

Now consider a different setting, where the communication link between  $\circ$  and  $\bullet$  is also allowed to drop both messages, but only for at most  $k$  consecutive rounds for some fixed  $k$ . We call this model the *Bounded-Delay Lossy-Link*, abbreviated as *BDLL( $k$ )*. Following our previous notation,  $\circ \cdots \bullet$  denotes the silent graph.

<sup>1</sup> Throughout this paper,  $\omega$  denotes the first infinite ordinal, it is very convenient for expressing regular infinite sequences in a compact way.

► **Definition 3.2** (Bounded-Delay Lossy-Link message adversary). *We define the  $k$ -Bounded-Delay Lossy-Link message adversary, denoted by  $BDLL(k)$ , as:*

$$BDLL(k) := \left( \bigcup_{i=0}^k \{ \circ \cdots \bullet^i \} \cdot \{ \circ \rightarrow \bullet, \circ \leftarrow \bullet, \circ \leftrightarrow \bullet \} \right)^\omega.$$

*More generally, we define the Bounded-Delay Lossy-Link message adversary, as*

$$BDLL := \bigcup_{k \in \mathbb{N}} BDLL(k).$$

Note that, as some communication is always guaranteed,  $BDLL(k)$  has a non-empty kernel for any  $k$ , satisfying (i) and trivially (ii). This also applies to the  $BDLL$ , where any sequence is an instance of some  $BDLL(k)$ .

Nevertheless, just  $BDLL(1)$  already breaks  $\delta_{MinMax}$ , as it is not guaranteed anymore that the maximum in the system reaches all other processes in every round: Consider the graph sequence  $(\circ \leftarrow \bullet \cdot \circ \cdots \bullet)^\omega$  with  $x_\circ < x_\bullet$ , where  $\circ$  always alternates its decision value: it will choose  $x_\circ$  if the last communication graph was  $\circ \cdots \bullet$ , and  $x_\bullet$  if it was  $\circ \leftarrow \bullet$ , as it only considers messages from the last round. In order to circumvent this alternation, we could adapt  $\delta_{MinMax}$  to look back two rounds instead of only looking at the last one. In fact, we can generalize this idea, by adapting  $\delta_{MinMax}$  to consider the past  $k$  rounds and thus making it suitable for  $BDLL(k)$ :  $\delta_{MinMax}^k(view_\sigma(p_i, r)) := \max_{x \in V^k(view_\sigma(p_i, r))} \{\min(\text{inputs}(x))\}$ , where  $V^k(view_\sigma(p_i, r)) := \bigcup_{j=0}^{\min(k, r-1)} V(view_\sigma(p_i, r-j))$  is the set of views  $p_i$  has received in the last  $k$  rounds (when  $k < r$ ).

A similar correctness argument as the one for  $\delta_{MinMax}$  can be used to show that  $\delta_{MinMax}^k$  solves stabilizing consensus in  $BDLL(k)$ . However, any  $\delta_{MinMax}^k$  will fail in  $BDLL(k+1)$ . Charron-Bost and Moran cleverly fixed this in [9] and provided an algorithm capable of solving stabilizing consensus in  $BDLL(k)$  for *any* fixed  $k$ , i.e., in  $BDLL$ . Intuitively their safe MinMax algorithm works as follows: Any sequence  $\sigma \in BDLL$  is also a member of  $BDLL(k)$ , for some  $k$ . Of course, one cannot choose some  $l$  that is always larger than any  $k$ , but, as the sequence  $\sigma$  is chosen beforehand and  $k$  therefore fixed, one can gradually *increase*  $l$  such that it eventually surpasses any fixed  $k$ . And since  $\delta_{MinMax}^l$  also correctly runs for any  $k < l$ , once  $l$  surpassed  $k$  it will run correctly forever. Note that we will formalize this in Section 4 for  $n$  processes.

Nevertheless, we will prove that *no stabilizing protocol* can exist if there is *no fixed bound*  $k$  on the length of consecutive silence periods. We call this model the *Delayed Lossy-Link message adversary*, which can be seen as the limit of the  $BDLL$ .

► **Definition 3.3** (Delayed Lossy-Link Message Adversary). *We define the Delayed Lossy-Link Message Adversary, denoted as  $DLL$  as:*

$$DLL := (\circ \cdots \bullet^* \cdot \{ \circ \rightarrow \bullet, \circ \leftarrow \bullet, \circ \leftrightarrow \bullet \})^\omega,$$

$$\text{where } \circ \cdots \bullet^* \text{ is the Kleene star of } \circ \cdots \bullet, \text{ i.e., } \circ \cdots \bullet^* = \bigcup_{i \in \mathbb{N} \cup \{0\}} \{ \circ \cdots \bullet^i \}$$

Note that the  $DLL$  again has a non-empty kernel, as some communication always happens. It hence satisfies (i), does not satisfy (ii), as there is no bound on how many consecutive silent rounds may happen.

Also note that any communication pattern  $\sigma \in DLL$  is essentially a sequence in the  $LL$  with arbitrarily long but finite periods of no communication in-between. Consequently, any communication pattern  $\sigma \in DLL$  has a unique corresponding silence-free communication pattern  $\mu \in LL$  where  $\mu \sqsubseteq \sigma$ , which we call the *silence-free core* of  $\sigma$ . Reciprocally, we say that  $\sigma$  is a delayed pattern of  $\mu$ .

### 3.1 Stabilizing consensus is impossible in the DLL

In this sub section, we prove one of our main results, namely the impossibility of stabilizing consensus in the DLL.

The proof strategy is the following: we assume by contradiction that there is a stabilizing protocol that can solve consensus in the DLL. Then we show that it is possible to construct a run that infinitely often decides for a conflicting configuration, i.e., with different decision values for  $\circ$  and  $\bullet$ . The crucial argument for constructing a conflicting run is Theorem 3.9, which proves that any stabilizing protocol in the DLL can be translated into a patient protocol. A patient protocol allows processes to change their decision value only in rounds where they receive a message from another process.

We start by defining conflicted prefixes, which are somewhat similar (but different) to bivalent prefixes: whereas a bivalent prefix is a prefix that may eventually lead to mutually incompatible output configurations, a conflicted prefix is one that currently outputs an invalid configuration. For instance, the communication pattern  $\circ \rightarrow \bullet^\omega \in LL$  under  $\delta_{MinMax}$  with  $\times_\circ > \times_\bullet$  yields a bivalent prefix for any round  $r$ . Nevertheless,  $\delta_{MinMax}$  never outputs a conflicted prefix after round 1 in this setting.

► **Definition 3.4** (Conflicted prefix). *Let  $\mathcal{P}$  be a stabilizing protocol,  $T = \langle \mathcal{I}, \mathcal{O}, \Delta \rangle$  a task, and  $\sigma = (G_i)_{i \in \mathbb{N}}$  a communication pattern. We say that a prefix  $\sigma^k$  is conflicted with input configuration  $\mathcal{C}^0 = (p_i, \times_i)_{i=1}^n$ , iff  $\delta_{\mathcal{P}}(\mathcal{C}_\sigma^r) \notin \Delta(\mathcal{C}^0)$ .*

Respectively, we say that a communication pattern  $\sigma$  is *conflicted infinitely often* iff there is a strictly increasing sequence  $s : \mathbb{N} \rightarrow \mathbb{N}$  such that  $\sigma|_1^{s(i)}$ , the  $s(i)$ -prefix of  $\sigma$ , is conflicted. Clearly, a stabilizing protocol solves a task under a message adversary  $\mathcal{M}$  iff there are no communication patterns in  $\mathcal{M}$  that are conflicted infinitely often.

Note that the possibility of arbitrary long silence periods adds a new layer of uncertainty, particularly in the context of stabilizing protocols. In fact, this removes any possibility of acquiring any useful information through silence. For instance, in LL, even if process  $\circ$  does not receive a message in round  $r$ , it is at least sure that  $\bullet$  will receive a message in round  $r$ . The effect of this uncertainty on stabilizing protocols is reflected by the following definitions, Theorem 3.6, and Theorem 3.9.

Let  $\pi$  be a prefix of length  $k$ . For notational simplicity, we will use  $\delta_{p_i}(\pi)$  to denote  $\delta_{\mathcal{P}}(\text{view}_\pi(p_i, k))$ , whenever the protocol  $\mathcal{P}$  is either unique or not immediately relevant to the context; similarly  $\delta(\mathcal{C}^0 : \pi)$  denotes the output configuration  $(p_i, \delta_{p_i}(\mathcal{C}^0 : \pi))_{i=1}^n$ . Furthermore, if  $\sigma$  is an infinite communication pattern and  $\mathcal{P}$  is a stabilizing protocol, we denote by  $\bar{\delta}_{p_i}(\mathcal{C}^0 : \sigma)$  the stable decision by  $p_i$  under  $\sigma$  with initial configuration  $\mathcal{C}^0$ .

► **Definition 3.5** (Patient prefix). *Let  $\mathcal{P}$  be a stabilizing protocol, and  $T = \langle \mathcal{I}, \mathcal{O}, \Delta \rangle$  a stabilizing task in the DLL model. We say that a prefix  $\pi$  is a patient prefix for input configuration  $\mathcal{C}^0$ , iff for every  $r > 0$ ;  $\delta(\mathcal{C}^0 : \pi) = \delta(\mathcal{C}^0 : \pi \cdot \circ \cdots \bullet^r)$ .*

► **Lemma 3.6** (Patience Lemma). *Let  $\mathcal{P}$  be a stabilizing protocol for a task  $T = \langle \mathcal{I}, \mathcal{O}, \Delta \rangle$ ,  $\mathcal{C}^0 \in \mathcal{I}$ , and  $\pi$  any prefix in the DLL model. There exists  $k \geq 0$  such that  $\mathcal{C}^0 : \pi \cdot \circ \cdots \bullet^k$  is a patient prefix.*

**Proof.** Note that  $\pi \cdot \circ \rightarrow \bullet^\omega$  is an admissible execution in the DLL message adversary. Since  $\mathcal{P}$  is a stabilizing algorithm,  $\circ$  must eventually stabilize on a decision  $\bar{\delta}_\circ(\mathcal{C}^0 : \pi \cdot \circ \rightarrow \bullet^\omega)$ . Hence there is some minimal  $r_1 \geq 0$  such that for any  $r > r_1$ ,  $\delta_\circ(\mathcal{C}^0 : \pi \cdot \circ \rightarrow \bullet^r) = \bar{\delta}_\circ(\mathcal{C}^0 : \pi \cdot \circ \rightarrow \bullet^\omega)$ .

Reciprocally, there is some minimal  $r_2 \geq 0$  such that for any  $r > r_2$ ,  $\delta_\bullet(\mathcal{C}^0 : \pi \cdot \circ \leftarrow \bullet^r) = \bar{\delta}_\bullet(\mathcal{C}^0 : \pi \cdot \circ \leftarrow \bullet^\omega)$ .



Note that  $\mathcal{C}^0 : \pi \cdot \circ \rightarrow \bullet^{r_1}$  is indistinguishable from  $\mathcal{C}^0 : \pi \cdot \circ \dots \bullet^{r_1}$  to  $\circ$ . Symmetrically  $\mathcal{C}^0 : \pi \cdot \circ \leftarrow \bullet^{r_2}$  is indistinguishable from  $\mathcal{C}^0 : \pi \cdot \circ \dots \bullet^{r_2}$  to  $\bullet$ .

Let  $k = \max\{r_1, r_2\}$ , and consider  $\pi' = \mathcal{C}^0 : \pi \cdot \circ \dots \bullet^k$ . For any  $k'$ ,  $\pi' \cdot \circ \dots \bullet^{k'}$  is indistinguishable from  $\mathcal{C}^0 : \pi \cdot \circ \rightarrow \bullet^{k+k'}$  to  $\circ$ , and  $\pi' \cdot \circ \dots \bullet^{k'}$  is indistinguishable from  $\mathcal{C}^0 : \pi \cdot \circ \leftarrow \bullet^{k+k'}$  to  $\bullet$ .

Moreover, since  $k + k' \geq k$ , it follows that  $\bar{\delta} \circ (\mathcal{C}^0 : \pi \cdot \circ \rightarrow \bullet^\omega) = \delta \circ (\pi' \cdot \circ \dots \bullet^{k'}) = \delta \circ (\pi')$ , and  $\bar{\delta} \bullet (\mathcal{C}^0 : \pi \cdot \circ \leftarrow \bullet^\omega) = \delta \bullet (\pi' \cdot \circ \dots \bullet^{k'}) = \delta \bullet (\pi')$ . Consequently,  $\pi' = \mathcal{C}^0 : \pi \cdot \circ \dots \bullet^k$  is indeed a patient prefix.  $\blacktriangleleft$

Note that Theorem 3.6 allows us to define the *patience* of a prefix  $\pi$ , for a particular stabilizing protocol  $\mathcal{P}$  in the DLL and a particular input configuration  $\mathcal{C}^0$ .

► **Definition 3.7** (Patience). *Let  $\mathcal{P}$  be a stabilizing protocol,  $\mathcal{C}^0 \in \mathcal{I}$  be an input configuration, and  $\pi$  a prefix in the DLL. We say that  $k \in \mathbb{N}$  is the patience of  $\pi$  for  $\mathcal{P}$  and  $\mathcal{C}^0$ , if  $k$  is minimal such that  $\mathcal{C}^0 : \pi \cdot \circ \dots \bullet^k$  is a patient prefix.*

► **Definition 3.8** (Patient Protocol). *Let  $\mathcal{P}$  be a stabilizing protocol in the DLL. We say that  $\mathcal{P}$  is patient if the patience of any prefix  $\pi$  is 0.*

► **Lemma 3.9** (Patient Protocol Reduction). *Let  $\mathcal{P}$  be a protocol that solves a stabilizing task  $T = \langle \mathcal{I}, \mathcal{O}, \Delta \rangle$  in the DLL. Then there exists a patient protocol  $\mathcal{P}'$  that solves  $T$ .*

**Proof.** Consider a communication pattern  $\sigma = (G_i)_{i=1}^\infty \in \text{DLL}$ . We define inductively a prefix sequence  $\pi_i$  in the following way:  $\pi_1 := G_1 \cdot \circ \dots \bullet^{\alpha_1}$ ;  $\pi_{k+1} := \pi_k \cdot G_{k+1} \cdot \circ \dots \bullet^{\alpha_{k+1}}$ , where  $\alpha_1 := \text{patience}(G_1)$ ;  $\alpha_{k+1} := \text{patience}(\pi_k \cdot G_{k+1})$ . We define a communication pattern  $\sigma' := (H_i)_{i=1}^\infty$ , as the limit of  $\pi_i$ .

Consider a protocol  $\mathcal{P}'$  with decision map  $\delta'$  defined by:  $\delta'_x(\sigma|_1^k) = \delta_x(\pi_k)$ . We will show that  $\mathcal{P}'$  is a patient stabilizing protocol that solves  $T$ .

By construction, each  $\pi_i$  is a patient prefix of  $\sigma'$ , hence  $\mathcal{P}'$  is patient. Note that  $\sigma$  is an infinite sub-sequence of  $\sigma'$  that does not include an infinite silence suffix  $\circ \dots \bullet^\omega$ ,  $\sigma'$  does not include an infinite silent suffix  $\circ \dots \bullet^\omega$  either. Thus,  $\sigma' \in \text{DLL}$ .

Since  $\mathcal{P}$  is a stabilizing protocol that solves  $T$  in DLL, in particular,  $\delta$  provides a stabilizing solution for communication pattern  $\sigma'$ . Since  $\delta(\mathcal{C}^0 : \sigma') = \delta'(\mathcal{C}^0 : \sigma)$ , it follows that  $\mathcal{P}'$  is a stabilizing protocol that solves  $T$  for communication pattern  $\sigma$ .  $\blacktriangleleft$

► **Theorem 3.10** (Stabilizing consensus Impossibility in DLL). *Let  $\mathcal{P}$  be an arbitrary stabilizing protocol in DLL. There exists a valid communication pattern  $\sigma \in \text{DLL}$  such that  $\mathcal{P}$  does not solve stabilizing consensus for  $\sigma$ .*

**Proof.** Let us assume for a contradiction that there is a stabilizing protocol  $\mathcal{P}$  that solves stabilizing consensus in DLL. By Theorem 3.9, there exists a patient protocol  $\mathcal{P}'$  that solves stabilizing consensus in DLL. We will provide an inductive construction of a run that is conflicted. More precisely, we define a sequence of conflicted prefixes  $\pi_i$  such that  $\pi_i \subset \pi_{i+1}$ , and  $\lim_{i \rightarrow \infty} \pi_i = \sigma \in \text{DLL}$ .

**Base case (Empty prefix  $\varepsilon$ ).** For the base case, consider an initial input configuration  $\mathcal{C}^0$  with different values  $x_\circ \neq x_\bullet$ . Due to the validity condition, each process must decide on its own value. Since  $x_\circ \neq x_\bullet$ ,  $\mathcal{C}^0 : \varepsilon$  is a conflicted prefix.

**Induction step (Prefix  $\pi_k$ ).** Assume that  $\mathcal{C}^0 : \pi_k$  is a conflicted prefix, i.e.,  $\delta \circ (\mathcal{C}^0 : \pi_k) \neq \delta \bullet (\mathcal{C}^0 : \pi_k)$ .

We assert that at least one of  $\mathcal{C}^0 : \pi_k \cdot \circ \rightarrow \bullet$ ,  $\mathcal{C}^0 : \pi_k \cdot \circ \leftarrow \bullet$ ,  $\mathcal{C}^0 : \pi_k \cdot \circ \leftarrow \bullet$  is conflicted. Assuming that none is conflicted, we derive a contradiction (as illustrated in Figure 1).

## 18:10 Stabilizing Consensus Is Impossible in Lossy Iterated Immediate Snapshot Models

Recall that  $\mathcal{P}'$  is patient and  $\mathcal{C}^0 : \pi_k \cdot \circ \rightarrow \bullet$  is indistinguishable from  $\mathcal{C}^0 : \pi_k \cdot \circ \cdots \bullet$  for  $\circ \cdot$ . Therefore,  $\delta_\circ(\mathcal{C}^0 : \pi_k \cdot \circ \rightarrow \bullet) = \delta_\circ(\mathcal{C}^0 : \pi_k \cdot \circ \cdots \bullet) = \delta_\circ(\mathcal{C}^0 : \pi_k)$ .

Since we assumed that  $\mathcal{C}^0 : \pi_k \cdot \circ \rightarrow \bullet$  is not conflicted, then

$$\delta_\circ(\mathcal{C}^0 : \pi_k \cdot \circ \rightarrow \bullet) = \delta_\bullet(\mathcal{C}^0 : \pi_k \cdot \circ \rightarrow \bullet).$$

Note that  $\mathcal{C}^0 : \pi_k \cdot \circ \rightarrow \bullet$  is indistinguishable from  $\mathcal{C}^0 : \pi_k \cdot \circ \leftrightarrow \bullet$  for  $\bullet \cdot$ , and therefore

$$\delta_\bullet(\mathcal{C}^0 : \pi_k \cdot \circ \rightarrow \bullet) = \delta_\bullet(\mathcal{C}^0 : \pi_k \cdot \circ \leftrightarrow \bullet).$$

Similarly, since  $\mathcal{C}^0 : \pi_k \cdot \circ \leftrightarrow \bullet$  is not conflicted by assumption,

$$\delta_\bullet(\mathcal{C}^0 : \pi_k \cdot \circ \leftrightarrow \bullet) = \delta_\circ(\mathcal{C}^0 : \pi_k \cdot \circ \leftrightarrow \bullet).$$

Note that  $\mathcal{C}^0 : \pi_k \cdot \circ \leftrightarrow \bullet$  is indistinguishable from  $\mathcal{C}^0 : \pi_k \cdot \circ \leftarrow \bullet$  for  $\circ$  and thus,

$$\delta_\circ(\mathcal{C}^0 : \pi_k \cdot \circ \leftrightarrow \bullet) = \delta_\circ(\mathcal{C}^0 : \pi_k \cdot \circ \leftarrow \bullet).$$

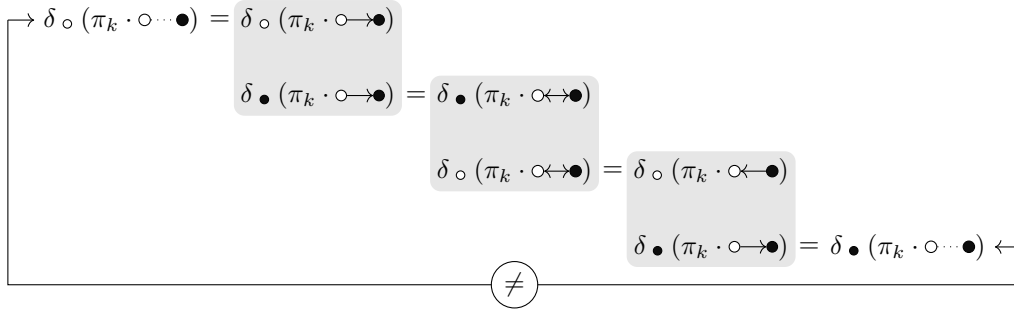
Since  $\mathcal{C}^0 : \pi_k \cdot \circ \leftarrow \bullet$  is not conflicted by assumption,  $\delta_\circ(\mathcal{C}^0 : \pi_k \cdot \circ \leftarrow \bullet) = \delta_\bullet(\mathcal{C}^0 : \pi_k \cdot \circ \leftarrow \bullet)$ .

Finally, since  $\mathcal{C}^0 : \pi_k \cdot \circ \leftarrow \bullet$  is indistinguishable from  $\mathcal{C}^0 : \pi_k \cdot \circ \cdots \bullet$  for  $\bullet$  and  $\mathcal{P}'$  is patient,

$$\delta_\bullet(\mathcal{C}^0 : \pi_k \cdot \circ \leftarrow \bullet) = \delta_\bullet(\mathcal{C}^0 : \pi_k \cdot \circ \cdots \bullet) = \delta_\bullet(\mathcal{C}^0 : \pi_k).$$

Therefore  $\delta_\circ(\mathcal{C}^0 : \pi_k) = \delta_\bullet(\mathcal{C}^0 : \pi_k)$ , which contradicts the induction hypothesis, namely that  $\mathcal{C}^0 : \pi_k$  is conflicted. Thus, either one of  $\mathcal{C}^0 : \pi \cdot \circ \rightarrow \bullet$ ,  $\mathcal{C}^0 : \pi \cdot \circ \leftrightarrow \bullet$ , or  $\mathcal{C}^0 : \pi \cdot \circ \leftarrow \bullet$  is a conflicted prefix.

Choosing  $G_{k+1} \in \{\circ \rightarrow \bullet, \circ \leftrightarrow \bullet, \circ \leftarrow \bullet\}$  such that  $\mathcal{C}^0 : \pi_k \cdot G_{k+1}$  is conflicted, we define  $\pi_{k+1} := \pi_k \cdot G_{k+1}$ . Setting  $\sigma = \lim_{k \rightarrow \infty} \pi_k$ , note that  $\sigma \in \text{DLL}$  by construction and each  $\pi_k$  is a conflicted prefix: Therefore  $\sigma$  is indeed conflicted infinitely often, and  $\mathcal{P}'$  does not solve stabilizing consensus.



■ **Figure 1** An illustration of the induction step in Theorem 3.10. Decision values grouped by a grey rectangle correspond to decision values of the same prefix, i.e., they should match otherwise we have a conflicted prefix. Decision values connected via an equality sign are necessarily identical because the process cannot distinguish the two prefixes. This creates a chain of equalities that is broken by our induction hypothesis, namely that  $\pi_k \cdot \circ \cdots \bullet$  is conflicted, implying one of the grey boxes, and consequently a longer prefix, must be conflicted as well. ◀

## 4 The Iterated Immediate Snapshot and the Lossy Iterated Immediate Snapshot Models

Previously, we introduced the Lossy-Link and Delayed Lossy-Link message adversary as a starting point for our stabilizing consensus impossibility result. This is for two main reasons: first, the simplicity of the Lossy-Link highlights the source of the impossibility result without adding distracting details to the proofs, and second, the DLL impossibility translates transparently to other message adversaries that generalize DLL.

One such message adversary is the *Iterated Immediate Snapshot* message adversary, which is computationally equivalent to the asynchronous Iterated-Immediate Snapshot model defined as an  $n$ -process asynchronous shared memory model where processes communicate

via atomic snapshots. The first approach towards this model was introduced by Borowski and Gafni [6], who presented the *Immediate-atomic-snapshot model* as a generalization of the classical FLP model [15]. They also re-used it as a convenient model for characterizing wait-free asynchronous shared memory systems [7].

In the IIS model, one considers a set  $\Pi = \{p_i\}_{i=1}^n$  of processes that communicate through a shared memory object  $M$ . Each process may write into its own register  $M[i]$  and is able to read the whole memory  $M$  in an atomic snapshot operation. A process  $p_i$  that executes an atomic snapshot writes into its register  $M[i]$  and reads all of  $M$  instantaneously. Since snapshot operations executed by different processes may occur concurrently, it is assumed that every process in a concurrent set is capable of reading each others output. Although atomic snapshot objects appear to be very strong, it has been shown that they are equivalent to single-reader single-writer wait-free shared registers [19].

Perhaps one of its most useful features of the IIS model is that it allows us to express an asynchronous shared memory model in terms of a synchronous message adversary [1], an asynchronous snapshot by all processes, can hence be represented by a communication graph. If a process  $p$  and a process  $q$  execute a snapshot operation, there are 3 possibilities, either  $p$  executed its snapshot before  $q$ ,  $q$  before  $p$ , or both snapshots are concurrent. If  $p$  executes its snapshot before  $q$ , then  $p$  cannot read what  $q$  wrote, but  $q$  will be able to read the value from  $p$ . Reciprocally, if  $q$  executes its snapshot before  $p$ ,  $q$  will see  $p$ 's value but not vice versa. If  $p$  and  $q$  have concurrent snapshots, then both are able to read each other's value. Since we consider time to be linearly ordered, and snapshots occur within the same time frame (round), then this means that for any pair of processes  $p$  and  $q$ , either  $p$  reads  $q$  or  $q$  reads  $p$ , or both. When translated to communication graphs, this implies that each communication graph in the IIS model is *semi-complete*<sup>2</sup>. Since the time at which each process executed its snapshot determines completely the memory registers that each process is able to read, every communication graph in the IIS model is also *transitive*<sup>3</sup>.

The converse can also be shown, such that for every communication pattern  $\sigma$  of semi-complete and transitive graphs, there is an IIS schedule that matches  $\sigma$ . Thus, for the rest of this paper, we will simply consider that IIS :=  $\mathcal{G}^\omega$ , where  $\mathcal{G}$  is the set of semi-complete and transitive graphs on  $\Pi$ . Note carefully that LL satisfies both transitivity and semi-completeness.

## 4.1 The Lossy Iterated Immediate Snapshot Model

Having introduced the IIS model as an  $n$ -process generalization of the Lossy-Link model, we will now add read-omission faults to the immediate snapshot operations. This means that, whenever processes execute a snapshot in a round, there may be at most  $f$  total read-omission failures of other processes' registers.

Note that the empty graph  $\circ \cdots \bullet \in \text{DLL}$ , can be viewed as a read-omission in either  $\circ \rightarrow \bullet$  or  $\circ \leftarrow \bullet$ . However, for  $n \geq 3$ , faulty snapshots are slightly more complex than simply considering empty graphs. Consequently, we define the set of *lossy snapshot graphs*,  $\Phi(f)$  as the set of directed graphs from  $\mathcal{G}$  (consisting of semi-complete and transitive directed graphs on  $\Pi$ ) with at least one missing edge and at most  $f \leq n(n-1)$  missing edges. More precisely,  $\Phi(f) := \{G \setminus F \mid G \in \mathcal{G}, F \subseteq E(G), 1 \leq |F| \leq f\}$ .

<sup>2</sup> A directed graph is semi-complete if for any pair of vertices  $v, w$ , either  $(v, w) \in E(G)$  or  $(w, v) \in E(G)$ .

<sup>3</sup> A directed graph  $G$  is transitive iff  $(v, w)(w, z) \in E(G) \Rightarrow (v, z) \in E(G)$ .

## 18:12 Stabilizing Consensus Is Impossible in Lossy Iterated Immediate Snapshot Models

Similarly to Section 3, we start by limiting the number of consecutive iterations where a snapshot might be faulty. We denote the *Bounded Lossy Iterated Immediate Snapshot* model by  $\text{BLIIS}(f, k)$ , where  $f \leq n(n-1)$  is the maximum number of omission faults per iteration, and  $k$  is the maximum number of consecutive rounds an immediate snapshot might be faulty. We then relax this restriction by requiring that correct snapshot iterations must only happen infinitely often (similar to the DLL). We call this model the *Lossy Iterated Immediate Snapshot* model, denoted by  $\text{LIIS}(f)$ . Note that for 2 processes,  $\text{BLIIS}_{\{\circ, \bullet\}}(1, k)$  corresponds to  $\text{BDLL}(k)$ , while  $\text{LIIS}_{\{\circ, \bullet\}}(1)$  corresponds to  $\text{DLL}$ .

These message adversaries can be formally defined as follows:

► **Definition 4.1** (IIS, BLIIS, LIIS Message Adversaries). *Let  $\mathcal{G}$  be the set of semi-complete and transitive graphs on  $\Pi$ , and  $\Phi(f)$  the set of faulty snapshot graphs for  $f \leq n(n-1)$ .*

- $\text{IIS} := \mathcal{G}^\omega$
- $\text{BLIIS}(f, k) := (\{\Phi(f)^i\}_{i=0}^k \cdot \text{IIS})^\omega$
- $\text{BLIIS}(f) := \bigcup_{k \in \mathbb{N}} \text{BLIIS}(f, k)$
- $\text{LIIS}(f) := (\Phi(f)^* \cdot \text{IIS})^\omega$ , where  $\Phi(f)^*$  is the Kleene star of  $\{\Phi(f)\}$ .

First, we prove that all presented message adversaries have a non-empty kernel, and hence satisfy (i).

► **Lemma 4.2.** *IIS,  $\text{BLIIS}(f, k)$ ,  $\text{BLIIS}(f)$  and  $\text{LIIS}(f)$  all have a non empty kernel, i.e., for any  $\sigma$ , we have  $\text{Ker}(\sigma) \neq \emptyset$ .*

**Proof.** As every communication graph in  $\sigma \in \text{IIS}$  is semi-complete and transitive,  $\sigma(r)$  contains at least one process that directly reaches all other processes<sup>4</sup>. As there are only finitely many processes, the pigeonhole principle guarantees that in the infinite sequence  $\sigma$  at least one process reaches all others infinitely often, implying  $\text{Ker}(\sigma) \neq \emptyset$ .

Now assume  $\sigma = (G_i)_{i=1}^\infty$  to be an admissible run in either  $\text{BLIIS}(f, k)$ ,  $\text{BLIIS}(f)$  or  $\text{LIIS}(f)$ . By construction, in any of these models, there is an infinite number of instances of graphs  $G \in \text{IIS}$  in  $\sigma$ , i.e.,  $(G_{s(i)})_{i=1}^\infty \sqsubseteq \sigma$  and  $(G_{s(i)})_{i=1}^\infty \in \text{IIS}$ . But this is a sub-sequence with a non-empty kernel as proven above, therefore  $\text{Ker}(\sigma) \neq \emptyset$  holds. ◀

We will now argue that  $\text{IIS}$ ,  $\text{BLIIS}(f, k)$  and  $\text{BLIIS}(f)$  allow solving stabilizing consensus. We start with the following technical lemma:

► **Lemma 4.3.** *Let  $\sigma$  be any communication pattern. There exists a round  $s_i$  such that no process learns any new input values anymore. Formally,  $\forall p_i \in \Pi : \forall r > s_i : \text{inputs}(\text{view}_\sigma(p_i, r)) = \text{inputs}(\text{view}_\sigma(p_i, s_i))$ . We call this set  $p_i$ 's stable value set.*

**Proof.** As the set of processes  $\Pi$  is finite and input values are never forgotten in a run, the set  $\text{inputs}(\text{view}_\sigma(p_i, r))$  can be increased only finitely often. ◀

► **Lemma 4.4.** *Let  $\sigma$  be any communication pattern. After round  $s_i$ ,  $p_i$ 's stable value set contains the stable value set of all processes in  $\text{Ker}(\sigma)$ . Formally,  $\forall p_i \in \text{Ker}(\sigma)$ ,  $\forall p_j \in \Pi : \text{inputs}(\text{views}_\sigma(p_j, s_j)) \subseteq \text{inputs}(\text{views}_\sigma(p_i, s_i))$ .*

<sup>4</sup> It is a well known result that transitive tournaments are equivalent to strict total order graphs, and hence finite transitive tournaments always contain a dominating vertex. Thus, any finite semi-complete and transitive graph also includes a dominating vertex. See for example [10] for more results on tournaments.

**Proof.** By contradiction. Assuming the contrary implies that  $p_i$  never receives a message from  $p_j \in Ker(\sigma)$  that reveals some input value known to  $p_j$ . This contradicts the assumption that  $p_j$  is in the kernel. ◀

► **Corollary 4.5.** *For any communication pattern  $\sigma$  with  $Ker(\sigma) \neq \emptyset$ , processes in the kernel have identical stable value sets.*

► **Theorem 4.6.** *For any  $\sigma \in IIS \cup BLIIS(f, k) \cup BLIIS(f)$ ,*

$$\delta_{MinMax}^{r/2}(view_{\sigma}(p_i, r)) := \max_{x \in V^{r/2}(\mathcal{C}^0: \sigma, p_i, r)} \{\min(\mathbf{inputs}(x))\}$$

*solves stabilizing consensus in  $\sigma$ . Note that the number of past rounds considered (i.e.,  $\frac{r}{2}$ ) now depends on the current round  $r$ .*

**Proof.** We set  $s_g = \max\{s_1, \dots, s_n\}$ , i.e., the round after which all processes have arrived at their stable value sets. From Theorem 4.4, it follows that for all rounds  $r > s_g$ ,  $\forall p_j \in Ker(\sigma) : \forall p_i \in \Pi : \min(\mathbf{inputs}(view_{\sigma}(p_j, r))) \geq \min(\mathbf{inputs}(view_{\sigma}(p_i, r)))$ . From Theorem 4.5, we infer that processes in the kernel have identical minima.

- Assume  $\sigma \in IIS$ , and  $r - 1 > s_g$ . As any  $\sigma(r)$  is semi-complete and transitive, all processes not in the kernel receive the view of a process  $p_j \in Ker(\sigma)$ . By the above reasoning,  $p_j$ 's minimum is larger than all other minima, therefore the decision function  $\delta_{MinMax}^{r/2}$  chooses correctly from round  $s_g + 1$  on.
- Assume  $\sigma \in BLIIS(f, k)$ , and  $r - k > s_g$ . As  $\sigma(r)$  is semi-complete and transitive for at least one graph among the ones for  $r \in \{r - k, \dots, r\}$ , we can repeat the previous reasoning but applied to views received in the previous  $k$  rounds,  $V^k(\mathcal{C}^0 : \sigma, p_i, r)$ . Eventually, i.e., after round  $s_j + k$ , we are sure that the interval  $r - k$  to  $r$  contains a view with the stable value set of a process  $p_i$  in the kernel. Note that as  $IIS \subseteq BLIIS(f, k)$ , this also solves stabilizing consensus on the IIS.
- Assume  $\sigma \in BLIIS(f)$ , and  $r > s_g$ . As  $\sigma$  does not have a predefined bound  $k$ , we cannot resort to the previous reasoning. However, we know there *is* a bound, as  $\sigma \in BLIIS(f)$  implies  $\sigma \in BLIIS(f, k)$  for some  $k$ . As any  $\delta^k$  solves stabilizing consensus also on any  $\sigma \in BLIIS(f, l)$  for  $l < k$  by inclusion, we can dynamically consider the previous  $\frac{r}{2}$  rounds. This ensures that (a) we eventually surpass the fixed bound  $k$ , and thus eventually always consider the view of a process in the kernel, and (b) eventually  $r - \frac{r}{2}$  also exceeds  $s_j$  and we only consider stable value sets of  $p_j \in Ker(\sigma)$ . Together with the fact that,  $p_j$ 's stable value set contains the largest minimum among all processes, we conclude that  $\delta^{r/2}$  solves stabilizing consensus on  $\sigma \in IIS \cup BLIIS(f, k) \cup BLIIS(f)$ . ◀

As the main result of this section, we will now prove that stabilizing consensus is impossible in LIIS(1), i.e., in the presence of just one faulty register per round finitely consecutive rounds. We extend the proof of Theorem 3.10 from DLL to LIIS(1) by identifying four admissible graphs in LIIS(1) that “replicate” the DLL impossibility result in the LIIS model.

Note that  $IIS = LIIS(0) \subseteq LIIS(1) \subseteq \dots \subseteq LIIS(f) \subseteq LIIS(f + 1)$ , thus if stabilizing consensus is impossible in LIIS(1), then it is also impossible in LIIS( $f$ ) for any  $f \geq 1$ .

For the remainder of this paper, we will consider a system with at least  $n \geq 2$  processes, and focus on two distinct processes  $p_1$  and  $p_2$ . For convenience, we assume that  $\circ \mapsto p_1$ , i.e., as the choice of index is free, we rename  $\circ$  in DLL to  $p_1$  in the context of the LIIS(1), and symmetrically  $\bullet \mapsto p_2$ . We use the following convenient notations:  $\tilde{\Pi} := \Pi \setminus \{p_1, p_2\}$ ;  $K_{\tilde{\Pi}}$  is the complete graph on  $\tilde{\Pi}$ ;  $G_1 \oplus G_2$  represents the graph defined by  $V(G_1 \oplus G_2) := V(G_1) \cup V(G_2)$ , and  $E(G_1 \oplus G_2) := E(G_1) \cup E(G_2) \cup V(G_1) \times V(G_2)$ , i.e., adding all possible edges from  $G_1$  to  $G_2$  but not the other way round, and also extend it to runs  $\sigma \oplus K_{\tilde{\Pi}} := (G_i \oplus K_{\tilde{\Pi}})_{i=1}^{\infty}$ .

We define the following graphs:  $G_{p_1 \rightarrow p_2} := (p_1 \rightarrow p_2) \oplus K_{\bar{\Pi}}$ ;  $G_{p_1 \leftrightarrow p_2} := (p_1 \leftrightarrow p_2) \oplus K_{\bar{\Pi}}$ ;  $G_{p_1 \leftarrow p_2} := (p_1 \leftarrow p_2) \oplus K_{\bar{\Pi}}$ ; and  $G_{p_1 \dots p_2} := (p_1 \dots p_2) \oplus K_{\bar{\Pi}}$ . Note that  $G_{p_1 \rightarrow p_2}$ ,  $G_{p_1 \leftrightarrow p_2}$  and  $G_{p_1 \leftarrow p_2}$  are semi-complete and transitive and therefore snapshot graphs in the IIS. Likewise,  $G_{p_1 \dots p_2} \in \Phi(1)$ , since it can be obtained by removing  $(p_1, p_2)$  from  $G_{p_1 \rightarrow p_2}$  or by removing  $(p_2, p_1)$  from  $G_{p_1 \leftarrow p_2}$ .

We show that any sequence in the DLL can be extended to a sequence in the LIIS(1) where  $p_1$  and  $p_2$  have identical views, implying that a protocol solving stabilizing consensus in the LIIS(1) necessarily solves it in the DLL, deriving a contradiction.

► **Lemma 4.7.** *Any  $\sigma \in \text{DLL}$  can be extended to a  $\sigma' \in \text{LIIS}(1)$ , s.t.,  $\text{Views}(\sigma) = \text{Views}_{\{p_1, p_2\}}(\sigma')$ .*

**Proof.** Let  $\sigma \in \text{DLL}$  and consider for  $\sigma' = \sigma \oplus K_{\bar{\Pi}}$ . By construction we know that  $\sigma'(r) \in \{G_{p_1 \rightarrow p_2}, G_{p_1 \leftrightarrow p_2}, G_{p_1 \leftarrow p_2}, G_{p_1 \dots p_2}\}$  and as  $\sigma$  does not contain an infinite silent suffix (i.e.,  $\sigma|_k^\infty \neq (p_1 \dots p_2)^\omega$  for all  $k > 0$ ),  $\sigma'$  also has no infinite suffix  $(G_{p_1 \dots p_2})^\omega$  where no communication between  $p_1$  and  $p_2$  happens. Thus,  $\sigma'$  is an admissible LIIS(1) sequence.

By construction, in  $\sigma'$ ,  $p_1$  and  $p_2$  only receive messages from each other, moreover, they only do so in rounds where they also receive a message in  $\sigma$ . Therefore their views are identical  $\text{Views}(\sigma) = \text{Views}_{\{p_1, p_2\}}(\sigma')$ . ◀

► **Theorem 4.8** (Stabilizing consensus Impossibility in LIIS(1)). *Let  $\mathcal{P}$  be an arbitrary stabilizing protocol in the LIIS(1) model.  $\mathcal{P}$  does not solve stabilizing consensus.*

**Proof.** Assume for a contradiction that  $\mathcal{P}$  solves stabilizing consensus in the LIIS(1) model. Take any  $\sigma \in \text{DLL}$ , from Theorem 4.7, it follows that there exists a  $\sigma' \in \text{LIIS}(1)$  with identical views for  $p_1, p_2$ . The decision map  $\delta_{\mathcal{P}}$  therefore solves stabilizing consensus on  $\sigma \in \text{DLL}$  directly contradicting Theorem 3.10. ◀

## 5 Conclusion

In this paper we provided the first stabilizing consensus impossibility result where the existence a non-empty kernel is guaranteed (i) but the number or rounds where everybody hears from a member in the kernel (ii) is not . While it has been shown that a non-empty kernel is necessary for solving stabilizing consensus in [9], we prove that it is not sufficient.

At the core of the stabilizing consensus impossibility lies the fact that padding the communication with arbitrarily long but finite silence periods greatly impairs the decision power of stabilizing protocols, essentially forcing them to fix a value during silence periods. This limitation enables us to find a valid prefix that has conflicting decision values, and thus prevents stabilization. This result highlights the importance of communication liveness within a system, since any bounded variants of the DLL model are capable of solving stabilizing consensus.

Furthermore, we extended this impossibility result to the Iterated Immediate Snapshot model, where the possibility of a single read-faulty snapshot makes stabilizing consensus impossible. This result sheds new light on the impact of omission faults on consensus, even when the termination condition is relaxed to stabilization, and the communication assumptions are as strong as a shared-memory model.

---

**References**

---

- 1 Yehuda Afek and Eli Gafni. Asynchrony from synchrony. In Davide Frey, Michel Raynal, Saswati Sarkar, RudrapatnaK. Shyamasundar, and Prasun Sinha, editors, *14th International Conference on Distributed Computing and Networking (ICDCN)*, volume 7730 of *LNCS*, pages 225–239. Springer Berlin Heidelberg, 2013. doi:10.1007/978-3-642-35668-1\_16.
- 2 Dana Angluin, Michael J. Fischer, and Hong Jiang. Stabilizing consensus in mobile networks. In Phillip B. Gibbons, Tarek Abdelzaher, James Aspnes, and Ramesh Rao, editors, *Distributed Computing in Sensor Systems*, pages 37–50, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. doi:10.1007/11776178\_3.
- 3 Jón Atli Benediktsson and Philip H. Swain. Consensus theoretic classification methods. *IEEE Trans. Syst. Man Cybern.*, 22(4):688–704, 1992. doi:10.1109/21.156582.
- 4 Dimitri P. Bertsekas and John N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, Belmont, MA, 1989.
- 5 V.D. Blondel, J.M. Hendrickx, A. Olshevsky, and J.N. Tsitsiklis. Convergence in multiagent coordination, consensus, and flocking. In *Proceedings of the 44th IEEE Conference on Decision and Control*, pages 2996–3000, 2005. doi:10.1109/CDC.2005.1582620.
- 6 Elizabeth Borowsky and Eli Gafni. Generalized flip impossibility result for t-resilient asynchronous computations. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '93, pages 91–100, New York, NY, USA, 1993. Association for Computing Machinery. doi:10.1145/167088.167119.
- 7 Elizabeth Borowsky and Eli Gafni. A simple algorithmically reasoned characterization of wait-free computation. In *Proceedings of the sixteenth annual ACM symposium on Principles of distributed computing (PODC)*, pages 189–198, 1997.
- 8 Bernadette Charron-Bost, Matthias Függer, and Thomas Nowak. Approximate consensus in highly dynamic networks: The role of averaging algorithms. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 9135 of *Lecture Notes in Computer Science*, pages 528–539. Springer Berlin Heidelberg, 2015. doi:10.1007/978-3-662-47666-6\_42.
- 9 Bernadette Charron-Bost and Shlomo Moran. Minmax algorithms for stabilizing consensus. *Distributed Comput.*, 34(3):195–206, 2021. doi:10.1007/s00446-021-00392-9.
- 10 Gary Chartrand, Linda Lesniak, and Ping Zhang. *Graphs & Digraphs*. Chapman & Hall/CRC, 6th edition, 2015.
- 11 Étienne Coulouma, Emmanuel Godard, and Joseph G. Peters. A characterization of oblivious message adversaries for which consensus is solvable. *Theoretical Computer Science*, 584:80–90, 2015. doi:10.1016/j.tcs.2015.01.024.
- 12 Edsger W. Dijkstra. Self-stabilizing systems in spite of distributed control. *Communications of the ACM*, 17(11):643–644, 1974. doi:10.1145/361179.361202.
- 13 Danny Dolev, Nancy Lynch, Shlomit S. Pinter, Eugene W. Stark, and William E. Weihl. Reaching approximate agreement in the presence of faults. *J. ACM*, 33(3):499–516, July 1986. doi:10.1145/5925.5931.
- 14 Shlomi Dolev. *Self-Stabilization*. MIT Press, 2000.
- 15 Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, April 1985. doi:10.1145/3149.214121.
- 16 Matthias Függer, Thomas Nowak, and Manfred Schwarz. Tight bounds for asymptotic and approximate consensus. *J. ACM*, 68(6), October 2021. doi:10.1145/3485242.
- 17 Eli Gafni. Distributed computing: a Glimmer of a theory. In *Algorithms and theory of computation handbook: general concepts and techniques*, page 29. Chapman & Hall/CRC, 2 edition, February 2010.

- 18 Hugo Rincon Galeana, Sergio Rajsbaum, and Ulrich Schmid. Continuous Tasks and the Asynchronous Computability Theorem. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, volume 215 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 73:1–73:27, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2022.73.
- 19 Maurice Herlihy, Dmitry Kozlov, and Sergio Rajsbaum, editors. *Distributed Computing Through Combinatorial Topology*. Morgan Kaufmann, Boston, 2014. doi:10.1016/B978-0-12-404578-1.00017-6.
- 20 Maurice Herlihy, Sergio Rajsbaum, Michel Raynal, and Julien Stainer. From wait-free to arbitrary concurrent solo executions in colorless distributed computing. *Theoretical Computer Science*, 683:1–21, 2017. doi:10.1016/j.tcs.2017.04.007.
- 21 Maurice Herlihy and Nir Shavit. The asynchronous computability theorem for t-resilient tasks. *Conference Proceedings of the Annual ACM Symposium on Theory of Computing*, August 1994. doi:10.1145/167088.167125.
- 22 Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, July 1982. doi:10.1145/357172.357176.
- 23 Jennifer Lundelius and Nancy Lynch. A new fault-tolerant algorithm for clock synchronization. In *Proc. 3rd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 75–88, August 1984. doi:10.1145/800222.806738.
- 24 Michel Raynal and Julien Stainer. Synchrony weakened by message adversaries vs asynchrony restricted by failure detectors. In *ACM Symposium on Principles of Distributed Computing (PODC)*, pages 166–175, Montréal, Québec, Canada, 2013. ACM Press. doi:10.1145/2484239.2484249.
- 25 Nicola Santoro and Peter Widmayer. Time is not a healer. In B. Monien and R. Cori, editors, *6th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 304–313, Berlin, Heidelberg, 1989. Springer Berlin Heidelberg. doi:10.1007/BFB0028994.
- 26 Ulrich Schmid, Bettina Weiss, and Idit Keidar. Impossibility results and lower bounds for consensus under link failures. *SIAM Journal on Computing*, 38(5):1912–1951, 2009. doi:10.1137/S009753970443999X.
- 27 Manfred Schwarz and Ulrich Schmid. Round-oblivious stabilizing consensus in dynamic networks. In Colette Johnen, Elad Michael Schiller, and Stefan Schmid, editors, *Stabilization, Safety, and Security of Distributed Systems - 23rd International Symposium, SSS 2021, Virtual Event, November 17-20, 2021, Proceedings*, volume 13046 of *Lecture Notes in Computer Science*, pages 154–172. Springer, 2021. doi:10.1007/978-3-030-91081-5\_11.
- 28 Josef Widder and Ulrich Schmid. Booting clock synchronization in partially synchronous systems with hybrid process and link failures. *Distributed Computing*, 20(2):115–140, 2007. doi:10.1007/s00446-007-0026-0.
- 29 Kyrill Winkler, Ami Paz, Hugo Rincon Galeana, Stefan Schmid, and Ulrich Schmid. The Time Complexity of Consensus Under Oblivious Message Adversaries. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, volume 251 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 100:1–100:28, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2023.100.
- 30 Kyrill Winkler, Ulrich Schmid, and Yoram Moses. A Characterization of Consensus Solvability for Closed Message Adversaries. In *23rd International Conference on Principles of Distributed Systems (OPODIS 2019)*. Schloss Dagstuhl – Leibniz Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.OPODIS.2019.17.