

Heterogeneous Paxos 2.0: the Specs

Aleksandr Karbyshev ^a and Isaac Sheff ^a

^aHeliax AG

* E-Mail: aleksandr@heliax.dev, isaac@heliax.dev

Abstract

We present *Heterogeneous Paxos 2.0* (HP2.0), an improved version of Heterogeneous Paxos consensus algorithm (HP). In a nutshell, HP2.0 simplifies the algorithm logic, reduces bandwidth usage, and enables a more efficient implementation.

HP2.0 is compatible with the requirements of HP and satisfies the same correctness properties. A formal specification of HP2.0 in TLA⁺ is available as a separate software artefact, with a formal proof of the key safety property of *Agreement* in TLAPS.

This report provides an accessible account of HP2.0, of the design space in which it exists, and of the design choices that led us to our current system.

Keywords: consensus ; distributed algorithm ; heterogeneous Paxos ;

(Received Jun 17, 2024; Version: Nov 19, 2024)

Contents

1	Introduction	2
1.1	Original Heterogeneous Paxos	2
1.2	Why Heterogeneous Paxos 2.0?	2
2	Differences From Original Heterogeneous Paxos	3
2.1	Broadcast Primitive	3
2.2	2a Messages	3
2.3	One Message In, at Most One Message Out	3
2.4	Byzantine Behaviour Detection	4
2.5	Logic changes	4
2.5.1	Buried	4
2.5.2	Well-Formed	4
2.5.3	Acceptor Algorithm	5
3	Specification	5
3.1	Network Model	5
3.2	Learner Graph	5
3.3	Protocol Message Structure	6
3.3.1	Properties of All Protocol Messages	6
3.3.2	Properties of Proposal Messages	6
3.3.3	Properties of Non-Proposal Messages	6

3.4	Definitions	7
3.5	Protocol	9
3.6	Protocol Properties	11
3.7	Mailbox Layer	11
4	Future Work	12
	References	12

1. Introduction

We present *Heterogeneous Paxos 2.0* (HP2.0): an improved, simpler, and more efficient version of Heterogeneous Paxos consensus protocol (HP) [SWvRM20]. The assumptions and guarantees of HP2.0 are similar to those of HP, but we simplify the algorithm logic and reduce communication complexity, thus improving efficiency.

A formal specification of HP2.0 in TLA⁺ [Lam02] is available in the TYRPHON repository [Ano]. The TLA⁺ spec has a formally verified proof of Agreement, a key safety property [Ano]. Here, rather than focusing on proofs, we attempt to explain the protocol in more reader-friendly terms.

This report is intended as the definitive “source of truth” for the definition of HP2.0. However, as we continue to integrate and prove more optimizations and simplifications, HP2.0 remains a work in progress.

1.1. Original Heterogeneous Paxos

This report assumes readers are familiar with the original HP technical report [SWvRM20], and is intended to explain the improvements we have made for HP2.0.

Original HP generalizes the Paxos consensus protocol [Lam98] for a setting with arbitrary quorums of acceptors, tolerating mixed crash and byzantine failures, with each learner tolerating *different* failure scenarios. The motivation, roles for participants, and trust model of HP2.0 are identical to those in HP [SWvRM20, §1–4, 8].

1.2. Why Heterogeneous Paxos 2.0?

HP2.0 improves upon HP by being both simpler and more efficient. The simplifications not only make the protocol easier to formally specify in TLA⁺, but also facilitate proving things about it and streamline its implementation. The differences between HP and HP2.0 are outlined in [Section 2](#).

HP2.0 admits a substantially more efficient implementation than HP, with no loss in power. For example, calculating which byzantine acceptors are *caught* in a message could naively require exponential time for HP. The

Byzantine behaviour detection modification of HP2.0 (Section 2.4) enables an implementation that can perform this check in linear time.

2. Differences From Original Heterogeneous Paxos

The HP2.0 consensus algorithm results from five substantial improvements to HP, which we describe in the remainder of this section.

2.1. Broadcast Primitive

HP2.0 assumes a *broadcast* primitive, which is responsible for sending messages to each acceptor and learner. For the liveness property to hold, *broadcast* must guarantee that if a message is received by one honest acceptor (i.e., an acceptor that is both safe and live), then that message is eventually received by all acceptors.

HP2.0 assumes a partially synchronous network [DLS88], where after some unknown global stabilization time (GST), all messages arrive within some (unknown but fixed) latency bound. After GST, the broadcast guarantee must deliver messages to all honest recipients within an (unknown) finite bound dependent on network latency. We do not require any particular message ordering guarantees.

In the HP algorithm, *broadcast* was implemented by having each acceptor echo well-formed messages to all other acceptors. In contrast, HP2.0 leaves the exact implementation of *broadcast* flexible, allowing for multiple possible approaches. One such implementation, as used in HP, involves all acceptors echoing all received messages to all other acceptors and learners. A more bandwidth-efficient implementation, however, would require acceptors to explicitly request any missing messages when a received message references one they do not recognize.

2.2. 2a Messages

Instead of sending a *2a*-message for each learner as in HP [SWvRM20], we send a single *2a* (non-proposal) message associated with a set of learners. Conceptually, this is similar to sending a set of *2a*-messages, but in practice, it is more efficient, both to send and to track with *refs* (Section 3.3).

2.3. One Message In, at Most One Message Out

With the broadcast primitive and the *2a*-messages, we can substantially simplify our protocol: In fact, we can remove all recursion and broadcast at most one message for each message received. This entails another minor change: instead of each actor receiving its own message in the same atomic action that it sends it (messages are broadcast, so actors receive their own messages), they receive them in some future action, just like any other message.

This change, in turn, may increase implementation efficiency by breaking up atomic actions into smaller schedulable pieces.

2.4. Byzantine Behaviour Detection

Each message specifies the previous message from the same sender (Section 3.3). This makes it much easier to detect certain kinds of Byzantine behaviour: two messages referencing the same parent form a proof of misbehaviour. In contrast, the original protocol called for comparing transitive history sets [SWvRM20]. This change makes detecting Byzantine behaviour much easier to implement without sacrificing any guarantees.

2.5. Logic changes

In some cases, we were able to simplify Heterogeneous Paxos by removing unnecessary constraints and checks. Our formal proof shows that the streamlined HP2.0 protocol retains the same safety conditions (Agreement and Validity) as the original. In each case, simplification can only improve liveness: all acceptor actions allowed without the simplification remain allowed with the simplification. Termination guarantees therefore remain unchanged.

We make three such simplifications: in the definition of Buried, the definition of Well-Formed, and the acceptor algorithm.

2.5.1. Buried

In HP2.0, to *bury* a $2a$ -message x , the acceptor must have seen a different $2a$ -message with a higher ballot and a different value. In contrast, HP required that a quorum of acceptors observe such a message [SWvRM20, Definition 28]. Upon analysing the formal safety proof for HP, we found that the quorum was only used to establish the existence of a single higher ballot $2a$ -message with a different value. We have formally proven that the quorum condition is unnecessary for burying an older $2a$ -message.

By eliminating this unnecessary constraint, HP2.0 consensus algorithm can achieve faster progress than HP in certain scenarios, as the acceptors no longer need to wait for redundant $2a$ -messages to propagate.

2.5.2. Well-Formed

We removed the “ $\text{Sig}(z) \in \text{Sig}(q(z))$ ” constraint for a $2a$ -message to be considered *well-formed* [SWvRM20, Assumption 29]. This constraint required that an acceptor must itself be part of the $1b$ -quorum from the same ballot in order to issue a $2a$ -message. However, this constraint was not necessary in both original Paxos [Lam98] and Byzantized Paxos [Lam11], and it has proven to be superfluous here as well. By removing it, acceptors can issue $2a$ -messages under strictly wider conditions: this change does not damage liveness.

2.5.3. Acceptor Algorithm

In [SWvRM20, Figure 3], acceptors could only issue $2a$ -messages in response to a $1b$ -message with the most recent ballot. HP2.0 allows acceptors to issue $2a$ -messages in response to any $1b$ -message. While a $1b$ -message from an old ballot would not enable any new $2a$ -messages, and so it is not clear *why* an acceptor would ever need to send a $2a$ -message in response to an old $1b$, allowing this action does not affect the protocol's formal guarantees: Agreement, Termination, and Validity. We therefore remove this condition to keep the formal protocol as simple as possible.

3. Specification

3.1. Network Model

We assume the *closed-world* distributed system consisting of a fixed finite set of *acceptors* \mathbb{A} , a fixed finite set of *proposers* \mathbb{P} and a fixed finite set of *learners* \mathbb{L} . We denote by $\mathcal{S} \subseteq \mathbb{A}$ a set of *safe*, non-Byzantine, acceptors that follow the protocol.

We assume that message broadcast is *reliable*, i.e., every message sent or received by a correct (i.e., safe and live) acceptor or proposer is eventually received by all live acceptors and learners¹. The delivery order of sent messages does not have to be preserved.

3.2. Learner Graph

We use the notion of *learner graph* introduced in [SWvRM20]. We recap its formal definition below.

Let \mathbb{L} be a fixed finite set of learners, and \mathbb{A} a fixed finite set of acceptors.

- Nodes of learner graph are elements $\alpha \in \mathbb{L}$.
- Each learner α is labelled with a set $Q_\alpha \subseteq 2^{\mathbb{A}}$. The elements of Q_α are quorums $q \subseteq \mathbb{A}$. We assume that each Q_α is closed upwards, i.e., for any $q' \supseteq q \in Q_\alpha$, we have $q' \in Q_\alpha$.
- For any pair $\alpha, \beta \in \mathbb{L}$, there is a graph edge labelled with a set of quorums, $\alpha-\beta \subseteq 2^{\mathbb{A}}$, called a *safe set* of α and β . We assume that any safe set $\alpha-\beta$ is closed upwards, i.e., for any $q' \supseteq q \in \alpha-\beta$, we have $q' \in \alpha-\beta$.
- The graph is undirected, i.e., $\alpha-\beta = \beta-\alpha$.

Definition 1 (Condensation). *We say that the learner graph is **condensed** if for all $\alpha, \beta, \gamma \in \mathbb{L}$*

$$\alpha-\beta \cap \beta-\gamma \subseteq \alpha-\gamma$$

¹As learners are never required to send messages, all learners are trivially live.

Definition 2 (Validity). We say that the learner graph is **valid** if for any $s \in \alpha\text{-}\beta$, $q \in Q_\alpha$, and $r \in Q_\beta$ we have $s \cap q \cap r \neq \emptyset$.

Definition 3 (Entanglement). We say that learners α and β are **entangled** if the set of safe acceptors \mathcal{S} belongs to the learners' safe set.

$$\text{Entangled}(\alpha, \beta) \stackrel{\text{def}}{=} \mathcal{S} \in \alpha\text{-}\beta$$

3.3. Protocol Message Structure

As far as message encoding “on the wire” is concerned, there are two types of messages: **proposal messages**, also called 1a-messages, and **non-proposal messages**. The encoding of these must be distinct, allowing recipients to identify the type of the received message.

3.3.1. Properties of All Protocol Messages

1. Each protocol message x carries a cryptographic signature that uniquely identifies the message signer, denoted by $\text{Sig}(x)$, with $\text{Sig}(x) \in \mathbb{A} \cup \mathbb{P}$.
2. Each protocol message has a unique hash, which other messages can use to reference it.

3.3.2. Properties of Proposal Messages

1. For each proposal message x , $\text{Sig}(x)$ must be a proposer: a participant authorized to generate new proposals.
2. Each proposal message x has a field $x.\text{val}$ containing a proposed value.
3. Each proposal message x has a field $x.\text{bal}$: a natural number specific to this proposal—its *ballot* number. We assume that ballots are value-specific: $x.\text{bal} = y.\text{bal}$ implies $x.\text{val} = y.\text{val}$. One way to implement this is to include the hash of the value in the (least significant bits of the) ballot.

3.3.3. Properties of Non-Proposal Messages

1. Each non-proposal message x carries a finite list of references to some other messages, $x.\text{refs}$. In practice, each reference is represented by a hash of the referenced message. Acceptors remember previously received messages (*known_messages*), allowing them to resolve references.
2. Each non-proposal message x carries a distinguished reference, $x.\text{prev}$, to a previous message signed by the same signer, identified by $\text{Sig}(x)$. If x is the first message sent by the sender, $x.\text{prev}$ contains a special non-reference value \perp .

3.4. Definitions

Definition 4 (1b). A non-proposal message x is called a 1b-message, denoted as $x:1b$, if its references contain a proposal message. Formally:

$$x:1b \stackrel{\text{def}}{=} \exists y \in x.\text{refs}. y:1a$$

Definition 5 (2a). Non-proposal message x is called a 2a-message, denoted as $x:2a$, if its references do not contain a proposal message. Formally:

$$x:2a \stackrel{\text{def}}{=} \forall y \in x.\text{refs}. \neg y:1a$$

We extend Sig over sets of messages to mean the set of signers of those messages:

Definition 6 (Message set signers). For any set of messages M :

$$\text{Sig}(M) \stackrel{\text{def}}{=} \{\text{Sig}(m) \mid m \in M\}$$

Definition 7 (Transitive references).

$$\text{Tran}(x) \stackrel{\text{def}}{=} \{x\} \cup \bigcup_{m \in x.\text{refs}} \text{Tran}(m)$$

Definition 8 (Transitive previous).

$$\text{PrevTran}(x) \stackrel{\text{def}}{=} \{x\} \cup \begin{cases} \text{PrevTran}(x.\text{prev}) & \text{if } x.\text{prev} \neq \perp \\ \emptyset & \text{otherwise} \end{cases}$$

Definition 9 (Message 1a).

$$\text{Get1a}(x) \stackrel{\text{def}}{=} \underset{m:1a \in \text{Tran}(x)}{\text{argmax}} m.\text{bal}$$

Definition 10 (Ballot).

$$B(x) \stackrel{\text{def}}{=} \text{Get1a}(x).\text{bal}$$

Definition 11 (Value).

$$V(x) \stackrel{\text{def}}{=} \text{Get1a}(x).\text{val}$$

Definition 12 (Caught acceptors).

$$\text{Caught}(x) \stackrel{\text{def}}{=} \text{Sig} \left(\left(\begin{array}{l} m \in \text{Tran}(x) \\ \exists m' \in \text{Tran}(x). \\ \text{Sig}(m) = \text{Sig}(m') \\ \wedge m \notin \text{PrevTran}(m') \\ \wedge m' \notin \text{PrevTran}(m) \end{array} \right) \right)$$

Definition 13 (Connected learners). For any learner α and message x :

$$\text{Con}_\alpha(x) \stackrel{\text{def}}{=} \{\beta \in \mathbb{L} \mid \exists s \in \alpha\text{-}\beta. s \cap \text{Caught}(x) = \emptyset\}$$

Definitions 14 to 18 are mutually recursive. They are sound because the recursive call is done on descendants of argument messages relative to the message reference relation.

Definition 14 (Buried 2a-messages). For any learner α , 2a-message x and message y :

$$\text{Buried}_\alpha(x, y) \stackrel{\text{def}}{=} \exists z \in \text{Tran}(y). z:2a \wedge \alpha \in \text{lrns}(z) \wedge B(z) > B(x) \wedge V(z) \neq V(x)$$

Definition 15 (Connected 2a-messages). For any learner α and message x :

$$\text{Con2as}_\alpha(x) \stackrel{\text{def}}{=} \left\{ m \in \text{Tran}(x) \left| \begin{array}{l} m:2a \\ \wedge \text{Sig}(m) = \text{Sig}(x) \\ \wedge \exists \beta \in \text{lrns}(x). \\ \beta \in \text{Con}_\alpha(x) \wedge \neg \text{Buried}_\beta(m, x) \end{array} \right. \right\}$$

Definition 16 (Freshness). For any learner α and 1b-message x :

$$\text{fresh}_\alpha(x) \stackrel{\text{def}}{=} \forall m \in \text{Con2as}_\alpha(x). V(m) = V(x)$$

Definition 17 (Quorum of message). For any learner α , 2a-message x :

$$q_\alpha(x) \stackrel{\text{def}}{=} \text{Sig}(\{m \in \text{Tran}(x) \mid m:1b \wedge \text{fresh}_\alpha(m) \wedge B(m) = B(x)\})$$

Definition 18 (2a-message learners). For any 2a-message x :

$$\text{lrns}(x) \stackrel{\text{def}}{=} \{\alpha \in \mathbb{L} \mid q_\alpha(x) \in Q_\alpha\}$$

The following *chain* property does not have a direct analogue in HP. It requires that the parent of a message be from the same sender, and included in the messages *refs* field.

Definition 19 (Chain property). For any message x :

$$\text{ChainRef}(x) \stackrel{\text{def}}{=} x.\text{prev} \neq \perp \rightarrow x.\text{prev} \in x.\text{refs} \wedge \text{Sig}(x.\text{prev}) = \text{Sig}(x)$$

Definition 20 (Decision). For any learner α and set of messages S :

$$\text{Decision}_\alpha(S) \stackrel{\text{def}}{=} \text{Sig}(S) \in Q_\alpha \wedge \forall x, y \in S. x:2a \wedge \alpha \in \text{lrns}(x) \wedge B(x) = B(y)$$

Any set of messages S such that $\text{Sig}(S) \in Q_\alpha$ is called a message quorum. If $\text{Decision}_\alpha(S)$ holds, we shall write $V(S)$ to denote $V(x)$ for some message $x \in S$.

Definition 21 (Well-formedness). *For any message x :*

$$\text{WellFormed1b}(x) \stackrel{\text{def}}{=} \forall y \in \text{Tran}(x). x \neq y \wedge y \neq \text{Get1a}(x) \rightarrow B(y) \neq B(x)$$

$$\text{WellFormed2a}(x) \stackrel{\text{def}}{=} \text{lrns}(x) \neq \emptyset$$

$$\begin{aligned} \text{WellFormed}(x) &\stackrel{\text{def}}{=} \\ &\quad \text{ChainRef}(x) \\ &\quad \wedge (x:1b \rightarrow \text{WellFormed1b}(x)) \\ &\quad \wedge (x:2a \rightarrow x.\text{refs} \neq \emptyset \wedge \text{WellFormed2a}(x)) \end{aligned}$$

3.5. Protocol

The formal specification of the protocol is formulated in TLA⁺ [Lam02] in the TYPHON repository [Ano]. The algorithm itself is formulated in PLUSCAL language [Lam09], which gets translated to TLA⁺ code.

For better readability, we present the pseudocode of learner and acceptor algorithms in Figures 1 and 2, respectively. Like TLA⁺, we specify exactly which actions are *safe* for each actor. This does not rule out unnecessary or repetitive actions. Formally, each action is a *predicate* over state changes: the action is only deemed safe if every part of the predicate holds true. Actions can include other actions.

To ensure liveness, we require *weak fairness* [Lam02]: in any execution trace, if an action is safe for an infinite sequence, it eventually occurs.

Next, we describe the semantics of some particular instructions.

Instruction Semantics.

broadcast(z) When called, it sends message z to every learner and acceptor, including the caller. Formally, this would mean that the defined state change is only allowed if the new state includes z in the network.

assume P is a synonym of PLUSCAL's "when" instruction [Lam24]. The instruction, defined for Boolean state predicate P , restricts possible executions and should be considered as a guard à la Dijkstra [Dij75]: execution of the function containing "assume P " is only possible if the predicate evaluates to true during such a putative execution.

with $x \in S$: C is adopted from the PLUSCAL counterpart [Lam24]: an element s is non-deterministically chosen from the given set S and C is executed with x being locally defined to be equal to s . The instruction "with $x = v$: C " is equivalent to "with $x \in \{v\}$: C ".

$C_1 \parallel \dots \parallel C_k$ means a parallel execution of C_i , $i = 1, \dots, k$.

```

1 Acceptor::init():
2   known_messages = {}
3   recent_messages = {}
4   prev_message =  $\perp$ 
5
6 Acceptor::process_1a( $m$ ):
7   assume  $m:1a$ 
8   with  $z = \text{msg}(\text{prev} = \text{prev\_message}, \text{refs} = \text{recent\_messages} \cup \{m\})$ :
9     if WellFormed1b( $z$ ):
10      recent_messages = { $z$ }
11      prev_message =  $z$ 
12      broadcast( $z$ )
13
14 Acceptor::process_1b( $m$ ):
15   assume  $m:1b$ 
16   with  $z = \text{msg}(\text{prev} = \text{prev\_message}, \text{refs} = \text{recent\_messages} \cup \{m\})$ :
17     if WellFormed2a( $z$ ):
18      recent_messages = { $z$ }
19      prev_message =  $z$ 
20      broadcast( $z$ )
21     else:
22      recent_messages  $\cup = \{m\}$ 
23
24 Acceptor::process_2a( $m$ ):
25   assume  $m:2a$ 
26   recent_messages  $\cup = \{m\}$ 
27
28 Acceptor::receive( $m$ ):
29   assume  $m \notin \text{known\_messages}$ 
30   assume  $\forall r \in m.\text{refs}. r \in \text{known\_messages}$ 
31   known_messages  $\cup = \{m\}$ 
32
33 Acceptor::process_message( $m$ ):
34   assume WellFormed( $m$ )
35   receive( $m$ )
36   process_1a( $m$ ) || process_1b( $m$ ) || process_2a( $m$ )

```

Figure 1. Heterogeneous Paxos 2.0 acceptor specification.

```

1 Learner::init():
2   known_messages = {}
3   decision = ⊥
4
5 Learner::receive(m):
6   assume m ∉ known_messages
7   assume ∀r ∈ m.refs. r ∈ known_messages
8   known_messages ∪= {m}
9
10 Learner::process_message(m):
11   assume WellFormed(m)
12   receive(m)
13
14 Learner::decide():
15   with s ⊆ known_messages:
16     assume Decisionself(s)
17     decision = V(s)

```

Figure 2. Heterogeneous Paxos 2.0 learner specification.

3.6. Protocol Properties

Theorem 22 (Validity). *For any learner α , and set of messages s :*

$$Decision_{\alpha}(s) \implies \exists x. x: 1a \wedge V(s) = V(x)$$

Proof. Directly follows from the definitions of B , V and $Decision$. □

Theorem 23 (Safety). *Let the learner graph of the network be valid and condensed. Let $\alpha, \beta \in \mathbb{L}$ be learners such that $Entangled(\alpha, \beta)$. For any protocol execution and reachable network state, if $Decision_{\alpha}(s_{\alpha})$ and $Decision_{\beta}(s_{\beta})$ hold in that state, for some message quorums s_{α} and s_{β} , then $V(s_{\alpha}) = V(s_{\beta})$.*

Proof. See TLA⁺ protocol formalization [Ano]. □

3.7. Mailbox Layer

The predicates `Acceptor::receive()` and `Learner::receive()` (Figures 1 and 2) define *causal message delivery*, i.e., the message can be received only once and only if all its direct references have already been received. In practice, this logic can be separated out into a special *mailbox* component. This component can be used by both local acceptor and learner, avoiding code duplication and allowing for more efficient message processing.

A mailbox implementation could be integrated into the broadcast implementation (Section 2.1): one mailbox might request missing messages from another when it receives a message m , but has not yet received m 's causal dependencies.

4. Future Work

Due to the recursive structure of protocol messages, a naïve implementation of the presented algorithm—utilizing caching of learner values for $2a$ -messages—would result in message processing complexity that is polynomial in l and n , where l represents the number of learners in the network and n is the total number of messages received so far.

In ongoing work, we are developing an efficient implementation of HP2.0 using additional data structures achieving a linear time complexity of message processing, $\mathcal{O}(l \cdot n)$. We also plan to formally prove that this implementation correctly realizes the proposed algorithm.

References

- Ano. Anoma. Typhon project github repository. <https://github.com/anoma/typhon/releases/tag/hpaxos-2.0-art-v3>. (cit. on pp. 2, 9, and 11.)
- Dij75. Edsger W. Dijkstra. Guarded commands, nondeterminacy and formal derivation of programs. *Commun. ACM*, 18(8):453–457, 1975. (cit. on p. 9.)
- DLS88. Cynthia Dwork, Nancy A. Lynch, and Larry J. Stockmeyer. Consensus in the presence of partial synchrony. *J. ACM*, 35(2):288–323, 1988. (cit. on p. 3.)
- Lam98. Leslie Lamport. The part-time parliament. *ACM Trans. Comput. Syst.*, 16(2):133–169, 1998. (cit. on pp. 2 and 4.)
- Lam02. Leslie Lamport. *Specifying Systems, The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley, 2002. (cit. on pp. 2 and 9.)
- Lam09. Leslie Lamport. The pluscal algorithm language. In Martin Leucker and Carroll Morgan, editors, *Theoretical Aspects of Computing - ICTAC 2009, 6th International Colloquium, Kuala Lumpur, Malaysia, August 16-20, 2009. Proceedings*, volume 5684 of *Lecture Notes in Computer Science*, pages 36–60. Springer, 2009. (cit. on p. 9.)
- Lam11. Leslie Lamport. Byzantizing paxos by refinement. In David Peleg, editor, *Distributed Computing - 25th International Symposium, DISC 2011, Rome, Italy, September 20-22, 2011. Proceedings*, volume 6950 of *Lecture Notes in Computer Science*, pages 211–224. Springer, 2011. (cit. on p. 4.)
- Lam24. Leslie Lamport. *A PlusCal User's Manual. C-Syntax Version 1.8*, 2024. (cit. on p. 9.)
- SWvRM20. Isaac Sheff, Xinwen Wang, Robbert van Renesse, and Andrew C. Myers. Heterogeneous paxos: Technical report, 2020. (cit. on pp. 2, 3, 4, and 5.)