

# Extending Classic Paxos for High-performance Read-Modify-Write Registers

Vasilis Gavrielatos, Antonios Katsarakis, Vijay Nagarajan  
The University of Edinburgh  
FirstName.LastName@ed.ac.uk

## 1 Introduction

In this paper we pose the problem of executing Read-Modify-Writes (RMWs) over replicated data, assuming that machines can crash and that processing and networking delays can be unbounded. This is also known as solving *asynchronous consensus*. Only that instead of assuming that machines must reach agreement on the next value of a “register”, we assume that machines agree on the next RMW to be executed over a replicated key-value pair that resides within an in-memory Key-Value Store (KVS). We also pose two additional requirements.

Firstly, we are looking for a solution that does not sacrifice availability when a machine fails. This renders the majority of consensus protocols unsuitable for our purpose. For instance, consider leader-based [6, 8, 9] consensus protocol. When the leader fails the rest of the machines must block waiting for a timeout to expire before performing leader election. Such protocols sacrifice availability for the duration of that timeout.

Secondly, we want to deploy the system in the datacenter over modern hardware. This means that the protocol must scale across the many cores of a modern server and must be able to utilize the high bandwidth of modern network interfaces. This requirement rules out the last standing protocol, EPaxos [7], that can solve asynchronous consensus without availability penalties. This is because EPaxos is so complex that is not clear how it can be deployed in a scalable fashion across the many threads of each machine. Absent that thread-scalability, it is impossible to leverage the high bandwidth offered by modern networks.

The requirements of 1) solving asynchronous consensus, 2) avoiding any availability penalty and 3) leveraging modern hardware, all point to Classic Paxos [5] (from now on *CP* or *Paxos*), which solves asynchronous consensus without blocking on a machine crash or delay while it can be deployed on a per-key basis without requiring any synchronization across threads. However there is a problem. For each key, CP can decide only a single value once. However, we must be able to perform many RMWs on the same key. Therefore, to use CP, we must first extend it, such that it can run repeatedly for each key, while preserving important invariants such that each RMW executes exactly once. Finally, we must make sure to avoid livelocks, which are a known issue of CP.

In this work we provide a detailed specification of how we extended and implemented Classic Paxos to execute Read-Modify-Writes in Kite [2]. In addition, we also specify how we implemented All-aboard Paxos [3] over Paxos and how we use carstamps [1], to also add ABD reads and writes, to accelerate the common case, where RMWs are not needed. Our specification targets a Key-Value-Store that is deployed within the datacenter, is replicated across 3 to 7 machines and supports reads, writes and RMWs.

First, to allow CP to run repeatedly we use the abstraction of a log per key, but – crucially – without actually implementing a log per key. Specifically, every time we commit an RMW for a key, we conceptually move to the next log entry, and we denote so through a single counter that remembers the current entry. The log abstraction allows us to execute CP repeatedly across the different entries, without actually modifying CP. The modularity of this approach is crucial as it allows us to rely on the well-established correctness guarantees of CP.

Second, once we have established the log abstraction, we must ensure that an RMW that has been committed in log entry  $X$  can never be committed again in log entry  $Y$ ,  $Y > X$ . We achieve this by tagging RMWs with unique identifiers (*rmw-ids*), and remembering the latest *rmw-id* committed from each session (bounded storage). In addition, we add commit messages to ensure that each replica knows when an RMW has been committed, and can thus stop it from getting proposed again in the future.

Finally, recall that our goal is to create a multi-threaded system that can stress the many-core servers and high-bandwidth NICs found in today’s datacenter. To achieve this we need to allow for thousands of RMWs to execute concurrently. However such concurrency can be problematic with CP as it can trigger livelocks. We avoid this scenario by implementing a back-off mechanism that for each key, allows each machine to have at most one active RMW. Furthermore, if machine M1 sees that machine M2 is trying to perform an RMW, then M1 will back-off, giving M2 a chance to complete its RMW unencumbered.

Beyond extending CP, we also provide the specification of an optimization called All-Aboard Paxos, which we implemented over our CP implementation. All-aboard is an optimization sketched in Howard’s thesis [3] as an application of the Flexible Paxos [4] theorem. To the best of our knowledge we are the first to fully specify and measure an All-aboard implementation. Finally, to provide a system that

offers a read, write, RMW API, we implement ABD reads and writes using carstamps [1] over our CP implementation.

The rest of this paper is organized as follows. Section 2 briefly describes how CP works (without our extensions). Section 3 describes our execution model and the basic data structures that underpin our implementation. Section 4 specifies our implementation of CP following the lifetime of a request. Section 5 and Section 6 explain how back-off and helping are implemented. Section 7 provides a series of correctness arguments and informal proofs for our extensions. Section 8 provides the rationale behind an array of our design choices. Section 9 presents the theory and implementation of All-aboard Paxos. Section 10 discusses how we added ABD writes to our implementation using carstamps [1]. Finally, Section 11 describes how we added ABD reads.

## 2 Classic Paxos Algorithm

Here we will give a very brief overview of Classic Paxos through Table 1. We will not explain why it works, but rather just how it works. There are two types of messages: proposes and accepts. A machine first issues proposes, if it gathers a majority of acks it then issues accepts. If accepts are acked by a majority of machines then the command is said to have been decided (i.e., committed).

Proposes and accepts carry logical timestamps (i.e., Tses, discussed in 3.1). Table 1 shows the case where machine M1 sends messages and M2 receives them. The leftmost column shows four potential messages that M1 can send, and the top row shows the four possible states that M2 can be in. We assume two timestamps  $H$  and  $L$ , where  $H > L$ . The rest of the cells denote the answer that M2 sends to M1 based on M1’s message and M2’s state. Each cell is colour-coded as blue, red or green, allowing us to group similar cases. Below we discuss each of the colours, noting that the crux of the protocol is found in the red cells.

**Green cells.** The green cells are when an accept message of M1 finds M2 having only received a propose with the same TS. This is the most straight-forward case where there has not been another intervening command, and M2 will always accept and reply with an ack.

**Blue cells.** Blue cells capture obscure corner cases, such as propose-L finds that M2 has already seen a propose-L. This is obscure because TS should be unique and we should expect that no TS is used twice for proposes. These can occur for example when we send the message twice, because we suspect the first message was lost. These cells are not of particular interest to understanding CP.

**Red cells.** Finally, red cells shows how CP handles Conflicts between commands with different Timestamps. The crux of this protocol lays here. Note the (simplified) rules.

1. Propose-H blocks propose-L and accept-L. I.e. proposes block proposes and accepts with lower Tses

	M2: Already seen propose-L	M2: Already seen accept-L	M2: Already seen propose-H	M2: Already seen accept-H
M1: Sends propose-L	Nack- restart	Nack- restart	Nack- restart	Nack- restart
M1: Sends accept-L	Ack	Ack	Nack- restart	Nack- restart
M1: Sends propose-H	Ack	Nack - Help	Nack- restart	Nack- restart
M1: Sends accept-H	Ack	Ack	Ack	Ack

Table 1: A simplified version of the core CP algorithm

2. Accept-H blocks propose-L and accept-L. I.e. accepts block proposes and accepts with lower Tses
3. Propose-L cannot block accept-H or Propose-H. I.e. proposes with lower TS are discarded
4. Accept-L cannot block accept-H, but it does block propose-H forcing it to help it. I.e. accepts with lower Tses cannot block accepts, but they do block proposes, and force them to help them

Note that our rules assume that proposes and accepts will all be seen by a majority of nodes and thus will intersect. Later we will relax this assumption (see §9).

## 3 System model and data structures

There are a finite number of machines (aka servers or nodes). Typical numbers are 3 to 7. Each machine runs a number of worker software threads (*workers*). Typical numbers are 20 to 30. Each worker runs a number of sessions. Typically 40 to 80. Each session maps to an external client. There is one FIFO per session. Client requests are inserted in the FIFO. Workers execute the requests of each session in order. Requests from different sessions can run concurrently. We will assume that all requests are read-modify-writes (RMWs).

Therefore, at any given moment the worker can be working concurrently on one RMW for each session. With 20 workers, 40 sessions and 5 machines, at any given moment there are 4000 RMWs running concurrently.

Each machine has the same Key-Value Store (KVS) in its main memory. For each RMW a Paxos instance is running. RMWs to different *KV-pairs* do not have any relation whatsoever. RMWs to the same *KV-pair* are said to *conflict*.

### 3.1 Data-structures and preliminaries

We start by establishing the data structures that will be used by the implementation.

**Message Types.** Paxos requires 2 broadcast rounds, a *propose* broadcast followed by an *accept* broadcast. Finally, if the accept is successful, a *commit* message will be broadcast.

### 3.1.1 The *KV-pair* and its metadata

***KV-pair***. For each key in the KVS, the KVS stores a data structure that contains the key, the value and some metadata. We will refer to this structure as *KV-pair*. The *KV-pair* is where RMWs to the same key from different machines/workers/sessions meet and synchronize. The metadata of the *KV-pair* will describe the state of the executing RMW (if any) at any given moment as perceived by the machine that stores the *KV-pair*.

***KV-pair* fields**. Below is a list of all the fields of *KV-pair*. The rest of this subsection is devoted to explaining these fields.

1. key
2. value
3. accepted-value
4. state
5. log-no
6. last-committed-log-no
7. proposed-TS
8. accepted-TS
9. rmw-id
10. last-committed-rmw-id

Crucially, every piece of metadata is absolutely necessary as we will see throughout this paper.

**State**. Each *KV-pair* stores a *state* variable. The state can be *Invalid*, *Proposed* or *Accepted*. The *state* variable refers to the last not-yet-committed *log-no*. *Invalid* means that the *KV-pair* knows of no ongoing RMW attempt. On receiving a propose (locally or from a different machine) the state will go to *Proposed*. On receiving an accept the state will go to *Accepted*. As we explain the protocol we will demonstrate when such state transitions are legal.

**Logical Timestamps—*TSes***. Paxos makes heavy use of logical timestamps (aka Lamport Clocks). A timestamp is a tuple of a version and a machine-id. To compare two timestamps, we compare their versions, using the machine-ids as tie breakers. We will refer to logical timestamps simply as *TSes*.

***KV-pair* Timestamps**. As part of the *KV-pair*'s metadata there are two *TSes*. A *proposed-TS*, which remembers the highest propose seen and an *accepted-TS*, which remembers the highest TS that has been accepted. The *accepted-TS* is used only in one case: when the *KV-pair* is in *Accepted* state and a propose with a lower TS than the *accepted-TS* is received.

**Log number - *log-no***. One of *KV-pair*'s metadata fields is the *log-no*, which is a counter for the number of RMWs that have committed for that particular *KV-pair*. Despite its name, there is no actual log – or need for one – in the implementation<sup>1</sup>. However, it is useful to think of a log for each *KV-pair*; for

each of the log's slots, Paxos must run to decide the winner that gets to commit its RMW.

Therefore, when a *KV-pair* stores a *log-no* = 10, that means that the key has already been successfully RMWed 9 times (at least), and we are currently working on the tenth. We do not actually keep a log, because we need not remember any of the values committed in slots 1 - 8. The fields: *state*, *proposed-TS*, *accepted-TS*, *accepted-value* and *rmw-id*, all refer to the *log-no*, i.e., to the log number we are currently working on. If the state is *Invalid*, these fields are meaningless. The *KV-pair* also stores a *last-committed-log-no*, that refers to the most recent *log-no* that has been committed (that it knows of). The *value* always refers to the value committed in *last-committed-log-no*.

**Helping**. In Paxos it is possible for machine M1 to “help” an RMW of a different machine M2. The help is necessitated by the uncertainty of the asynchronous environment: M1 cannot always know whether M2 has failed or not, and it cannot always know if an RMW is committed or not (because the RMW may have been committed by a majority that M1 cannot reach in its entirety).

**RMW-ids**. As a result of helping, it is possible that machine M2 attempts to commit its RMW, even though machine M1 helped that RMW in the past. Therefore, if we are not careful it may be the case that M1 commits an RMW in *log-no* = 10 and then machine M2 completes the exact same RMW in *log-no* = 11. This violates correctness. To ensure that an RMW is committed exactly once, we use *rmw-ids*. An *rmw-id* is an 8-Byte value, the LSBs of which contain the global-session-id (it could also be a tuple of <id, global-session-id>). Each session in the system has its own global-session-id and thus each RMW is granted a unique *rmw-id*. This allows us to identify different RMWs.

**Registering *rmw-id***. Each machine keeps the latest RMW that it knows has been committed by each session globally. Therefore, with 5 machines, 20 workers, 40 sessions, each machine holds an array with 800 fields. Each field denotes the most recent *rmw-id* of every other session. Note that if we know that the *rmw-id* <10, 253>, from session 253, has been committed, then it must be that the RMW with *rmw-id* <9, 253> has also been committed. Therefore, we need only remember the latest *rmw-id* committed by each session. When a worker learns about a committed-id, it immediately registers it, i.e. it updates the array with committed RMWs accordingly.

As a result when a machine tries to commit an already committed RMW, other machines are able to detect it and reply to it that the RMW has already been committed.

**Necessary *rmw-ids* in the *KV-pair***. The *KV-pair* stores: 1) an *rmw-id*, referring to the the RMW currently being worked on the *log-no* and 2) a *last-committed-rmw-id* which refers to the last committed RMW (in *last-committed-log-no*).

<sup>1</sup>Interestingly Lamport also observes that there is no need to keep an actual log for a KVS in his original paper [5].

**Accepted-value.** In order to facilitate help, we must know the result of the RMW we are helping, this is why the *KV-pair* has the field *accepted-value*, to remember the value that the highest accepted RMW wants to commit. The *accepted-value* field is only valid, if the key is in *Accepted* state and it only refers to the *log-no*. Note that alternatively, we could only be storing the identity of the RMW operation. For example, if the RMW is a Fetch-and-Increment, we can store its opcode instead of the value it creates, and calculate the value on demand (although the value is probably no more than 8 bytes in this scenario). However, this implementation assumes that the most commonly RMWs are Compare-and-Swaps, where remembering the whole identity of the RMW would be twice as expensive as storing its result (it has a compare-value and an exchange-value!).

### 3.1.2 Local-entries

**Local-entries.** Each worker has one *Local-entry* pre-allocated for each session. The *Local-entry* holds all the necessary state for the RMW. It remembers the exact operation the RMW wants to achieve and to which key, whether the RMW has communicated with the local *KV-pair* or other machines, what types of responses it has received and so on. Note the contrast between a *KV-pair* and a *Local-entry*: the *KV-pair* is shared among all worker threads and stores the state of the RMW that is currently in the front stage executing; *Local-entries* are thread-local and store the state of every RMW, including that of RMWs that are currently sidelined, in a back-off state waiting to get access to the *KV-pair*.

The *Local-entry* has a state variable, describing the state of the RMW: the state can be *Invalid*, *Proposed*, *Accepted*, *Needs-KV-pair*, *Retry-with-higher-TS*, *Bcast-commits*, *Bcast-commits-from-help* and *Committed*. We will discuss these states as we go along. Here is a list of field members of the *Local-entry*.

1. *key*
2. *state*
3. *accepted-value*
4. *accepted-log-no*
5. *rmw-id*
6. *back-off-counter*
7. *helping-flag*
8. *helping-local-entry*

**Unique Identifier – lid.** When receiving a reply to a message we need to know to which RMW it corresponds. But because the same RMW may have triggered multiple broadcasts of the same type, we tag each broadcast with a unique identifier called *lid*. The *Local-entry* remembers the *lid* used for the most recent broadcast (propose or accept). Replies to that broadcast always include the *lid*. On receiving a reply, the worker searches if there is an active *Local-entry* with that *lid*. If not it disregards the message. Otherwise, the worker passes the reply to the *Local-entry*. As an optimization, we use the

least significant bits of the *lid*, to store the session-id, such that we can immediately locate the *Local-entry*, a reply refers to, without the need for a (linear) search.

This is not merely a networking implementation detail. A unique identifier is necessary because it allows us to discard replies to older messages, e.g., a reply for the same session, same *rmw-id* and *log-no*, may refer to an older propose attempt.

**Jargon.** We will use the term local/locally to refer to the issuing machine of an RMW, or the machine that stores the KVS. For example, an RMW is said to be locally accepted, if it has gotten the *KV-pair* that is stored in the issuing machine's KVS in *Accepted* state.

### 3.1.3 Execution Model

The worker operates in a while(true) loop. In every iteration it: 1) polls for remote messages and takes action based on them by altering local state and enqueueing replies, 2) inspects all active *Local-entries* (i.e. not in *Invalid* state), and it takes action based on their state, 3) it sends enqueued messages (either broadcasts or unicasts which are always replies to broadcasts) and 4) probes the client FIFOs to pull requests for sessions that are not blocked. Therefore Paxos protocol actions are triggered, only in response to a received message or when inspecting the *Local-entry* of an RMW.

## 4 The protocol spec

Here we will see the lifetime of an RMW step-by-step. The description includes a number of forward pointers explaining additional cases.

### 4.1 Grabbing the local KV-pair

When the worker pulls a new RMW for a session, it immediately takes it to the local KVS. If the *KV-pair* is in *Invalid* state, then the *KV-pair* is “grabbed” by the RMW, transitioning it in *Proposed* state and essentially stopping other RMWs from the same machine, from attempting RMWs. If the *KV-pair* is already grabbed, then the RMW must *back-off* by transitioning its *Local-entry* to a state called *Needs-KV-pair*.

Note that if the *KV-pair* can be in *Invalid* state then the RMW on the previous *log-no* has been committed:  $KV-pair.last-committed-log-no + 1 == KV-pair.log-no$ . (It is possible that an RMW has worked on the current *log-no* before, but it gave up and reverted it to *Invalid* state– discussed in §8.1).

Let us assume that the RMW successfully grabs the *KV-pair*. (Section 5 discusses the *Needs-KV-pair* state as part of the back-off mechanism.) The RMW fills its corresponding *Local-entry*, with information about its current state: the key to be RMWed, its *rmw-id*, the *log-no* it works on, and the TS it will use for the propose message. The TS.version can be any arbitrary number. We use the number 3 (this will become important when discussing the All-aboard optimization

in Section 9). The TS.machine-id is the machine-id of the proposer.

**Propose.** The RMW must broadcast propose messages to the rest of the machines. A propose message includes the key, the TS of the propose, the *log-no* and the *rmw-id*.

## 4.2 On receiving a propose

The receiver will go to its local KVS to examine the *KV-pair* and sent back one of the following replies: *Rmw-id-committed*, *Log-too-low*, *Log-too-high*, *Seen-higher-prop*, *Seen-higher-acc*, *Seen-lower-acc*, or *Ack*. Note that all replies that are not an *Ack* are essentially a nack. The replies always include an opcode. Below we discuss the following: when each reply is triggered, what is its additional payload and what happens to the *KV-pair* in the machine that sends the reply.

- *Rmw-id-committed*: if the *rmw-id* of the propose message has been registered. The reply message does not include additional payload. The *KV-pair* is not altered. (Section 8.1 elaborates on this.)
- *Log-too-low*: if the *log-no* of the propose message is smaller than the current working *log-no* of the *KV-pair*. This means the proposer does not know of the latest committed RMW, and thus attempts to commit on an already used *log-no*. The reply message includes the last committed RMW (*last-committed-log-no*, *last-committed-rmw-id* and value). The *KV-pair* is not altered.
- *Log-too-high*: if the *log-no* of the propose message is higher than the current working *log-no* of the *KV-pair*. This means the proposer knows of a committed RMW, that the receiver does not know of. The reply message does not include additional payload. The *KV-pair* is not altered.
- *Seen-higher-prop*: if the *KV-pair* is in *Proposed* state but with a higher or equal *proposed-TS* than the propose's TS. The reply message includes the *proposed-TS* of the *KV-pair*. The *KV-pair* is not altered.
- *Seen-higher-acc*: if the *KV-pair* is in *Accepted* state but with a higher or equal *proposed-TS* than the propose's TS. The reply message includes the *proposed-TS* of the *KV-pair*. The *KV-pair* is not altered. (This is identical as above.) (Note: the *accepted-TS* is not inspected at all, because if the *proposed-TS* is higher than the propose's TS, the proposer has to retry either way.)
- *Seen-lower-acc*: if the *KV-pair* is in *Accepted* state and its *accepted-TS* is lower than the propose's TS. Crucially, the *KV-pair* remains in *Accepted* state, but its *proposed-TS* is advanced to the propose's TS, if it is smaller. The reply message includes the details of the accepted rmw, i.e., the *accepted-TS*, the *rmw-id*, and the *accepted-value*.
- *Ack*: if the *KV-pair* is currently in *Invalid* state, or the *KV-pair* is in *Proposed* state with a lower *proposed-TS* than the propose's TS. The *KV-pair* goes to *Proposed* state

and updates its proposed-TS. The reply message does not include any payload.

## 4.3 On receiving a propose-reply

The replies are steered to local entries (as explained in §3.1.2). The worker periodically polls active *Local-entries* and takes action when the *Local-entry* signals that a majority (or more) of replies has been gathered or one reply of type *Rmw-id-committed/Log-too-low/Seen-higher-prop/Seen-higher-acc* has been received. With 5 machines we wait for 2 replies, since we already have a reply –typically an *Ack*, but not always–from the local machine. Then the worker acts on the replies as follows.

Note that when processing the replies, there is a question of priority; e.g., if there is 1 seen-higher-acc, 1 seen-lower-acc and 2 acks, what should we do? The following discussion takes that into account.

- *Rmw-id-committed*: if any *Rmw-id-committed* is received, then the proposer commits its RMW in the local *KV-pair* (using the *Local-entry* fields *accepted-value* and *accepted-log-no*). The *Local-entry* transitions to *Bcast-commits* state so that commits are broadcast (again using the *Local-entry* fields *accepted-value* and *accepted-log-no*). Section 7.2.2 explains why using *accepted-value* is correct. Section 8.1 shows an optimization that allows us to avoid broadcast commits when not necessary.
- *Log-too-low*: if any *Log-too-low* is received, the RMW it contains is committed in the local *KV-pair*. If none of the above replies have been received, then the *Local-entry* transitions to *Needs-KV-pair* state, so that it tries to grab the entry again, but in a later *log-no*. Section 8.2 elaborates on this.
- *Seen-higher-prop/Seen-higher-acc*: if a *Seen-higher-prop* or *Seen-higher-acc* is received and none of the above is received then the *Local-entry* transitions to *Retry-with-higher-TS* state. The RMW then will retry to broadcast proposes using a TS that is higher than all TSes included in the received replies. Section 8.4 describes in detail the different cases when retrying.
- *Ack*: if a majority of *Ack* replies has been gathered and none of the above is received, then the RMW will attempt to accept locally, by transitioning the local *KV-pair* to *Accepted* state. Section 8.5 describes in detail the different cases when accepting locally. If it is successful, it will then calculate the *accepted-value*, store it in the *Local-entry* and the *KV-pair*, store the *log-no* in the *Local-entry* as *accepted-log-no* and broadcast accepts. The accept messages will include, the *accepted-value*, the *log-no*, the *rmw-id* and the same TS as the propose.
- *Seen-lower-acc*: If a *Seen-lower-acc* reply has been gathered and none of the above conditions have been triggered, then we will attempt to help the received RMW. Section 8.5 describes in detail how helped RMWs are accepted locally.

After accepting locally, accepts for the helped RMW are broadcast.

- *Log-too-high*: If a *Log-too-high* reply has been received and none of the above has occurred (e.g., with 5 machines we have received 1 *Log-too-high* and the 2 *Acks*), then we retry the RMW, transitioning the *Local-entry* to *Retry-with-higher-TS-state*. Retrying is discussed at length in Section 8.4. Notably, if this occurs repeatedly, we will time-out, broadcasting commits from the immediately previous *log-no*. This is discussed in detail in Section 8.7 .

**Priority.** The priority is due to performance and not correctness. If the RMW has been committed, we should not waste our time with other replies. If we are working on an already committed *log-no*, the same. If our TS is too small, we should not bother checking acks, we will probably end-up loosing anyway. Similarly, if acks have been gathered, we should not try to help a lower-TS accept, as we are sure to have blocked it. Finally, we should only care about the *Log-too-high* reply, only if none of the above is triggered.

#### 4.4 Accept

After managing to accept locally, accept messages are broadcast to remote machines. Note that when accepting locally the value-to-be-written and value-to-be-read for the RMW are calculated. The value-to-be-written is called *accepted-value*. The *accepted-value*, must be stored in both the *KV-pair*'s metadata in the KVS and the *Local-entry*. The *KV-pair* needs the *accepted-value* to facilitate help, by responding to proposes with higher TS with the *Seen-lower-acc* reply, that includes the *accepted-value*. The *Local-entry* needs the *accepted-value*, both to create accept messages that will be broadcast, but also in case it gets helped by another machine, to know what value must be read. Help is discussed in Section 6, reading the correct value is discussed in Section 7.2.2. The accept messages will include, the *accepted-value*, the *log-no*, the *rmw-id* and the same TS as the propose.

#### 4.5 On receiving an accept

The receiver will go to its local KVS to examine the *KV-pair* and send back one of the following replies: *Rmw-id-committed*, *Log-too-low*, *Log-too-high*, *Seen-higher-prop*, *Seen-higher-acc* or *Ack*. Note that all replies that are not an *Ack* are essentially a nack. The replies always include an opcode. Here is when each reply is triggered, what it is its additional payload and what happens to the *KV-pair* in the receiving side. A lot of the replies are identical to the ones described for the propose, in that case we will simply denote so.

- *Rmw-id-committed*: <identical to propose>
- *Log-too-low*: <identical to propose>
- *Log-too-high*: <identical to propose>
- *Seen-higher-prop*: if the *KV-pair* is in *Proposed* state but with a higher (but not equal, this is the difference with

propose-replies) *proposed-TS* than the accepts's TS. The reply message includes the *proposed-TS* of the *KV-pair*. The *KV-pair* is not altered.

- *Seen-higher-acc*: if the *KV-pair* is in *Accepted* state but with a higher (but not equal, this is the difference with propose-replies) *proposed-TS* than the accept's TS. The reply message includes the *proposed-TS* of the *KV-pair*. The *KV-pair* is not altered. (This is identical as above.)
- *Ack*: if the *KV-pair* is currently in *Invalid* state, or the *KV-pair* is in *Proposed* state with a lower or equal *proposed-TS* than the accept's TS, or the *KV-pair* is in *Accepted* state with a lower or equal *proposed-TS* than the accept's TS. The *KV-pair* goes to *Accepted* state and updates its *proposed-TS*, *accepted-TS* and *accepted-value*. The reply message does not include any payload.

#### 4.6 On receiving an accept-reply

Identically to the propose replies, accept replies are steered to *Local-entries* (as explained in §3.1.2). The worker periodically polls active *Local-entries* and takes action when the *Local-entry* signals that a majority (or more) of replies has been gathered, or if one reply of type *Rmw-id-committed/Log-too-low* has been received. With 5 machines we wait for 2 replies, since we already have a reply – always an *Ack*, because we must have accepted locally– from the local machine. Then the worker acts on the replies as follows. (Note that the list below also denotes the priority with which replies are handled).

- *Rmw-id-committed*: <identical to propose>
- *Log-too-low*: <identical to propose>
- *Ack*: if a majority of *Ack* replies has been gathered and none of the above is received, then the *Local-entry* transitions to *Bcast-commits* state, so that it broadcasts commits the next time it is inspected.
- *Seen-higher-prop/Seen-higher-acc*: <identical to propose>.
- *Log-too-high*: If a *Log-too-high* reply has been received and none of the above has occurred then we retry the RMW, transitioning the *Local-entry* to *Retry-with-higher-TS-state*. Retrying is discussed at length in Section 8.4. This is different than proposes, as repeated *Log-too-high* replies do not lead to a commit. This is discussed at length in 8.7.

**Helping.** If the accept was helping (the *Local-entry* has a *helping-flag* to let us know), then if any one nack is received (i.e., one of *Rmw-id-committed*, *Log-too-low*, *Seen-higher-prop*, *Seen-higher-acc*, *Log-too-high*), then we stop helping (lowering the *helping-flag*) and transition the *Local-entry* to the *Needs-KV-pair* state, where it will attempt to “grab” the *KV-pair* to perform its own RMW.

Otherwise, if a majority of *Acks* are gathered, we broadcast a commit message by transitioning the *Local-entry* to *Bcast-commits-from-help* state (so that it knows to broadcast the value getting helped).

**Explanation of priority.** Firstly, if the RMW has already been committed or the *log-no* has been used, then there is no reason into inspecting other replies. Again as before there is a question of priority. The Paxos invariant is that if a majority of machines ack an accept, then that command is committed. Therefore, if a majority of *Acks* is received, we can prioritize it. Otherwise if any of *Seen-higher-prop*/*Seen-higher-accl*/*Log-too-high* has been received, the worker will create a TS higher than all received and switch the *Local-entry* to the state *Retry-with-higher-TS*.

Note that it is possible that a majority of machines have acked our accept, but we don't see that (i.e. because we only inspected 2 out of 4 replies and one of the two was a *Seen-higher-prop* reply). One of two things can happen in this case: we will get helped by a remote machine, or we will end-up "helping ourselves" (help is discussed in §6).

#### 4.7 Commits

On inspecting the *Local-entry* and finding it in state *Bcast-commits* or *Bcast-commits-from-help*, commit messages are broadcast.

**Why commits.** In principle, an RMW is committed if it has been accepted by a majority of machines. Still commit messages are necessary. The reason is that it is impossible to know whether an RMW that is accepted, is also committed, unless you help it yourself. Therefore, without commits before beginning an RMW you would have to first help any previously accepted RMW, even if that has happened a long time ago. In addition, commits ensure that an RMW can not be committed twice (e.g., helped in *log-no* = 5, and then committed by the originator again in *log-no* = 7, is possible without commit messages). Finally, they allow us to reply to the client with the guarantee, that if they read they will see the committed value, without having to run Paxos. In Section 7, we show how commits help us ensure the correctness invariants for moving across *log-nos*.

Commits include all of the RMW details (its *rmw-id*, *log-no* and *value*). This allows us to ensure that even machines that have not seen the previous accept messages can commit the RMW. However machines that have acked this RMW already have this data. We discuss the implementation of an optimization that leverage this observation in Section 8.6.

After broadcasting commits, *Local-entry* is transitioned to the *Committed* state. Commits are always acked; upon gathering a majority of acks, we commit the RMW locally and reply to the client that the RMW is completed including the value to be read. Receiving a commit message always results in an unconditional commit of the RMW. When committing, we must register the new RMW, if not already registered, we must update the *value*, *last-committed-log-no* and *last-committed-rmw-id* fields of the *KV-pair* if no later *log-no* has been committed and we must transition the *KV-pair* in *Invalid* state, if it currently works on the *log-no* to be committed.

#### 4.8 A Note on the accepted-TS

You may be tempted to optimize away the *accepted-TS* of the *KV-pair*, after all it is only used once. Do not. It is an integral part of the Paxos algorithm and critical for correctness. Without the *accepted-TS* the proposers have no idea what they should help, and the receivers of proposes do not know what to tell proposers they should help.

In a nutshell, the *proposed-TS* is there to block lower proposes and accepts. The *accepted-TS* is there to tell proposes what they should help.

### 5 The back-off mechanism

We saw that if an RMW does not find the *KV-pair* in *Invalid* state, it does not grab it. Rather its *Local-entry* transitions to the *Needs-KV-pair* state. The *Local-entry* also stores the RMW that is currently holding the *KV-pair* and its state.

When the worker inspects the *Local-entry*, if it finds it in *Needs-KV-pair* state, it attempts to grab the *KV-pair*. The *KV-pair* can be grabbed if it's in *Invalid* state. When the *KV-pair* cannot be grabbed, the *Local-entry* notes the details of the *KV-pair*, its state, its current *proposed-TS* etc. Therefore, when the RMW fails repeatedly to grab the *KV-pair*, it takes note of whether the RMW that holds the *KV-pair* has made any progress. If nothing has changed since the the last time that the RMW attempted to grab the *KV-pair*, then a counter that is stored in the *Local-entry* is incremented. The counter measures the time in which the *KV-pair* has not changed. When the counter reaches a threshold (pre-determined e.g., at compile-time), then the RMW will attempt to steal or help the entry. Otherwise, if there is progress since last inspected, then the counter goes to zero.

Essentially, this is a back-off mechanism. Recall that we can have thousands of RMW running concurrently. If a few of them are on the same *KV-pair*, then there is no chance of progress. Instead, we limit the conflicting RMWs to one per machine, in the worst case. This is because a *KV-pair* can be grabbed by a remote machine.

If the counter reaches the threshold, that typically means that the *KV-pair* is held by a remote machine, which has failed. In this instance, if the *KV-pair* is in *Proposed* state, an RMW can simply "steal" the *KV-pair*, by using a higher *proposed-TS*. The RMW will trigger propose broadcasts, and will transition its *Local-entry* to *Proposed* state, whereby it will wait for propose replies.

However, if the *KV-pair* is in *Accepted* state, we can *not* steal it. We must help it.

### 6 Help

If a propose finds a *KV-pair* in *Accepted* state with a lower *accepted-TS*, then help must be triggered. The reason is that, it is possible that the propose reached some machines before the *Accept*. Therefore, the *Accept* is blocking the propose in some machines, and the *Propose* is also blocking the *Accept*

in some machines. The proposer cannot know for sure, because it only inspects a majority of responses and therefore it pessimistically breaks the deadlock, by helping the Accept.

A propose helps the accept, by broadcasting accepts with its own TS. This is crucial. The proposer will use its own TS, which is the big one. Why does then a *Seen-lower-acc* reply has the *accepted-TS* inside it, if it's not going to be used? Because the propose must help the highest *accepted-TS*. Again this is a matter of correctness.

Help can occur when a *Seen-lower-acc* reply is received to a propose, or when the *Local-entry* times out on waiting to grab the *KV-pair* locally, and the *KV-pair* is in *Accepted* state.

**Helping after a *Seen-lower-acc* reply.** Assume we broadcast a propose for an RMW – let's call it l-RMW – and we receive a reply *Seen-lower-acc* from a different RMW – let's call it h-RMW. We then attempt to help h-RMW. However, after we have finished helping h-RMW, we must go back to executing our own l-RMW. To help us, the *Local-entry* has the following fields: a flag called *helping-flag*, that is raised to denote that the *Local-entry* is currently helping h-RMW, and a *helping-local-entry*, a data structure similar to the *Local-entry*, which will be used to store information about the h-RMW, ensuring that no information about l-RMW is overwritten.

From that point we move as before, attempting to accept h-RMW locally and then commit it. However, if we are not able to accept locally, or if we receive a *Seen-higher-prop* or *Seen-higher-acc* reply, then the help is cancelled: we take down the *helping-flag* and we transition the *Local-entry* to *Needs-KV-pair* state, so that it starts over.

If we manage to receive a majority of acks for h-RMW, that means it is committed. We then transition the *Local-entry* to *Bcast-commits-from-help* state, denoting that commits must be broadcast, but they must use the *helping-local-entry* for the correct value. Finally, when a majority of commit acks have been gathered, we commit to the local KVS, the *helping-flag* goes down and the *Local-entry* transitions to the *Needs-KV-pair* state. Notably, if we were helping ourselves, we should detect it and simply free the session.

**Help after waiting locally.** In the case, where we timeout on attempting to grab a *KV-pair* locally, which is in *Accepted* state, the *KV-pair* cannot be stolen (like we would do if it were in *Proposed* state). The reason is that the *Accepted* RMW –let's called it h-RMW – may be accepted by 3 out of 5 machines and thus committed. It is possible that the issuer had broadcast commits and then died, and thus it is possible that out of the 4 remaining machines, one received the commit and has been reading the RMW, two machines have no idea about the h-RMW and our local machine has *Accepted* but not received the commit. Therefore, if we were to simply steal the *KV-pair*, and put it in *Proposed* state, we would open the door for other RMWs to be committed in the *log-no*, where an RMW is already committed.

Therefore, we must act as if we had sent a propose message to the local KVS and it has responded with a *Seen-lower-acc* reply. The *KV-pair* remains in *Accepted* state, but its *proposed-TS* is advanced to the TS of our propose (for the propose we use a bigger TS than the *KV-pair's proposed-TS*). We then broadcast proposes to remote machines. From that point we proceed as we would normally. Specifically, if we can gather a majority of acks from the remote machines, then we can get the local *KV-pair Accepted* but for our RMW. Otherwise, if we must help, then we proceed to help as described above.

**Crucial Take-away.** It is crucial to remember that a *KV-pair* can never transition from *Accepted* to *Proposed* (and remain in the same *log-no*). It can only transition from *Accepted* to *Accepted* with a higher *accepted-TS*.

## 7 Correctness

We have added several extensions to Paxos to make it run repeatedly (i.e., across *log-nos*) for a given key.

It is challenging to argue that the protocol is both correct and efficient. We will attempt to do this in the following manner. First we will list a few invariants. We will argue (or informally prove) that the protocol enforces said invariants. We will show how additions to the protocol help in doing that. Then, we will link the invariants with high-level correctness requirements: an RMW commits exactly-once, the issuer of an RMW reads the correct value. We will try to provide arguments that the invariants are both sufficient and necessary.

### 7.1 Invariants

- If machine M works (proposes or accepts) on  $log-no = X$ , then it must be that all previous *log-nos* have been committed (*inv-1*)
- If machine M works (proposes or accepts) on  $log-no = X$ , then it must be that M knows what has been committed in  $log-no = X - 1$  (*inv-2*)
- It can never be that an RMW gets accepted locally for  $log-no = X$ , if it has already been committed in a  $log-no = Y$ , where  $Y < X$  (*inv-3*)

Roughly, we will enforce *inv-1* by ensuring that work on  $log-no = X$  can start only after  $log-no = X - 1$  has been committed. *Inv-2* and *Inv-3* will be enforced by nacking with *Log-too-high* all proposes and accepts that refer to  $log-no = X$  when  $log-no = X - 1$  is not known to be committed.

#### 7.1.1 Inv-1

*Inv-1* mandates that a machine M can never work on  $log-no = X$ , unless all smaller *log-nos* have been committed.

This is proved trivially by the following statement. For machine M to work on  $log-no = X$ , it must be that at some point one machine grabbed its *KV-pair* for  $log-no = X$ , with an *Invalid KV-pair* and *last-committed-log-no = X - 1*. There is no other way for work to start on  $log-no = X$ .



### 7.1.2 Inv-2

Inv-2 mandates that machine M cannot work (i.e. propose or accept) on  $log-no = X$ , unless it knows what was committed in  $log-no = X - 1$ . Note that this is subtly different than inv-1, because it mandates that M itself must know what was committed on the immediately previous  $log-no$ .

For an RMW to work on a  $log-no = X$ , it must be that it has grabbed the *KV-pair*. There are two cases on how an RMW can grab a *KV-pair*.

**1st case.** The RMW grabs the *KV-pair* it found in *Invalid* state. The invariant is trivially enforced here: the entry cannot be in *Invalid* state, unless it knows that the previous  $log-no$  has been committed.

**2nd case.** The second case is that the RMW attempts to steal or help an entry after the back-off timeout expires. However, in the first place, the *KV-pair* can only transition to *Proposed* or *Accepted* state iff the immediately previous  $log-no$  is committed, i.e  $log-no = last-committed-log-no - 1$ . This is because remote proposes/accepts are nacked with a *Log-too-high* reply if the previous  $log-no$  than the one they want has not been committed. Same goes for local proposes/accepts.

### 7.1.3 Inv-3

Inv-3 mandates that it can never be that an RMW gets accepted locally for  $log-no = X$ , if it has already been committed in a  $log-no = Y$ , where  $Y < X$  (inv-3).

**Proof.** Assume machine M proposes in  $log-no = Z$ , an RMW it previously locally accepted in  $log-no = X$ . Also assume that the RMW get committed in  $log-no = X$ , because some other machine helps it. It suffices to show that M will receive at least one *Rmw-id-committed* reply to its propose message.

Firstly, it must be that  $X < Z - 1$ , because M, will only attempt to propose on  $log-no = Z$  if it knows what has been committed in  $log-no = Z - 1$  (from inv-2), and therefore it has registered the RMW committed in  $log-no = Z - 1$ . Furthermore, we know that the RMW has already been committed – by help – in  $log-no = X$ , before M proposes for  $log-no = Z$ . This is because the  $log-no = Z - 1$ , can only be committed if the  $log-no = X$  is already committed (from inv-1).

Therefore, before M can issue proposes for  $log-no = Z$ , it must be that a majority of machines have already acked an accept for every  $log-no \leq Z - 1$ . Because machines only ack accepts, if they have committed the previous  $log-no$  (or else the respond with a *Log-too-high*), we can infer that for each RMW committed in  $log-nos$  smaller than Z, there is a majority of machines that have committed the RMW. Therefore, when M issues proposes for  $log-no = Z$ , a majority of machines must have committed and registered the same RMW in  $log-no = X$  and thus M must received at least one *Rmw-id-committed* reply.

## 7.2 Guarantees

Let's now see how enforcing the above invariants results into higher-level guarantees. Inv-1 is just there to simplify proving inv-2 and inv-3. Firstly we will discuss why each of inv-2 and inv-3 are necessary. Then we will see a counter-example of what can happen without them.

### 7.2.1 Necessity of inv-2

Assume we ack accepts, without knowing the previously committed RMW. For example, in a 5-machine deployment, assume machines M1, M2, M3, M4 ack an accept from M5 for  $log-no = 10$ , but they all have a *last-committed-log-no* = 2. This can happen if M5 has committed all the intermediate *log-nos*. Then assume that M5 dies. M1 steals the *KV-pair* to help in  $log-no = 10$ . But M1 cannot work on  $log-no = 10$ , because it does not know what has been committed in  $log-no = 9$ , and therefore, if it fails to help, it will have to do its own RMW, but it wont know what value to use as the previously committed. Therefore, M1 has to start from the *last-committed-log-no*+ 1 (i.e., 3) and work its way up. However, if it sends a propose for  $log-no = 3$ , everyone else will answer with a *Log-too-low* nack, but they will only include the RMW committed on  $log-no = 2$ !

**Alternative solution.** It is possible, in the above example to work directly on  $log-no = 9$ , which we know has been committed (from inv-1) and therefore if M1 issues proposes for  $log-no = 9$ , it should be able to track down the highest accepted-TS, so that it can commit it directly without bothering with accepts. However, the other machines will still respond with *Log-too-low* nacks to a propose for  $log-no = 9$  (because they have accepted M5's RMW to  $log-no = 10$ ). So we then should have to take a lot of extra care (a lot of added metadata and complexity) to have the machines remember the last accepted value/ts for both the current  $log-no$  and the previous one, so that they can answer to proposes for both  $log-no = 10$  and 9.

Note this alternative would also violate inv-3, as no machine will know what has been committed in  $log-nos$  3 through 8. So it is possible that it is one of their own RMWs, and they can end up repeating it, breaking exactly-once semantics. We could solve this, by having accept messages including a *last-registered-rmw-id* field that would allow a majority of machines to register committed rmws, even if it has not actually committed them itself. However, that would incur a fixed overhead in all accept messages.

Overall this alternative seems very complicated and with potentially high overhead.

### 7.2.2 Necessity of inv-3

**Exactly-once semantics.** Inv-3 enforces exactly-once semantics as it ensures that an RMW can never get committed twice in two different logs, by ensuring that it cannot even get accepted.

**The RMW creates correct value.** The RMW creates its value when it gets accepted locally. For the RMW to create the correct value, it must be that it knows the value committed in the previous log-no. Inv-2 guarantees exactly that.

Note that the additional step, we take to enforce inv-2 was to force the RMW to work on *KV-pair.last-committed-log-no* instead of on *KV-pair.log-no*, when stealing/helping a stuck RMW after a back-off timeout. If we didn't do that then, it would be possible for an RMW to get accepted locally, without knowing the value committed in the previous log. Then, creating an *accepted-value* based on the current value of *KV-pair* (i.e., *KV-pair.value*) would be wrong.

**Ensuring that the RMW reads the correct value.** An RMW that gets retried, may be informed that it has already been committed. However, remote machines cannot be reasonably expected to say in which *log-no* the RMW has been committed or what value is it supposed to read. To do that, we would require unbounded storage. Remote machines simply check their registered *rmw-ids* (bounded storage) and reply back simply that the RMW has been committed.

What value is then a RMW to read, when it learns that it has been committed? To solve this we leverage a critical insight: for an RMW to be helped it must be that it got locally accepted. This is because, only accepted RMWs can be helped, and an RMW must get accepted locally before it can be accepted by another machine.

We can then combine the insight with the inv-3 (i.e., it can never be that an RMW gets accepted locally for  $log-no = X$ , if it has already been committed in a  $log-no = Y$ , where  $Y < X$ ). Assuming that inv-2 is enforced, then it must be that an RMW can always return the *accepted-value* from its *Local-entry*, which gets calculated every time the RMW is locally accepted. Plainly, if the RMW got locally accepted in  $log-no = X$ , if it later finds out that it has been committed, then it must be that it has been committed in  $log-no = X$ , and thus it can read the *Local-entry's accepted-value*.

### 7.2.3 Counter-example for the invariants

Let's look at an example of what bad can happen without the invariants. Assume 5 machines: M1, M2, M3, M4 and M5:

1. M1 accepts RMW-1 locally in  $log-no = 1$  and then fails to get accept majority
2. But, M2 helps M1's RMW-1, with an accept from majority of M2, M3, M4.
3. M2 broadcasts commits for RMW-1
4. M3 sees the RMW-1 commit and immediately proposes, accepts and commits RMW-3 in  $log-no = 2$ , with propose and accept majorities of M1, M3, M4
5. M1 sees the committed RMW-3 and then goes on to the next  $log-no = 3$  to retry its RMW-1
6. M1's propose for RMW-1 in  $log-no = 3$  gets a majority of acks from M1, M4 and M5
7. M1 accepts locally RMW-1 for  $log-no = 3$

From this point on M1 will either find out that its RMW has already been committed, in which case it will not know what value to read (even if it could remember previous accepted-values in its local-entry, it wouldn't know which one got helped!) or M1 will manage to commit RMW-1 in  $log-no = 3$ , breaking the exactly-once semantics.

The problem here is created because M4 acks M1's propose. M4 has not received the commit yet from M2 which helped RMW-1, and thus has not registered RMW-1, yet. With the *Log-too-high* nacks to proposes/accepts, M4 can not ack M3's attempt to commit on  $log-no = 2$ , before it commits and registers the RMW on  $log-no = 1$ .

## 8 Why and how

In this section we are tying loose ends, discussing the why and the how of several implementation choices an optimizations.

### 8.1 Why: *Rmw-id-committed*

Assume M1 sends a a propose/accept for an RMW that has already been committed; M2 replies back with an *Rmw-id-committed* message that contains the entire last committed *rmw*.

Firstly the *Rmw-id-committed* reply is absolutely essential, to ensure exactly-once semantics, seeing as RMWs can be helped.

When M1 receives the *Rmw-id-committed* rep it attempts to commit its RMW locally, using the last-accepted-value and the accepted-log-no from its *Local-entry*.

In addition, M1 needs to ensure that, before it reports completion of the RMW to the client, a majority of machines have committed it. To do that it may need to broadcast commits from the RMW. However, if M2 has already committed a later  $log-no = X$ , than the one M1 is shooting for (e.g.,  $log-no = Z$ ), that means that the RMW has been already committed in a  $log-no \leq Z$ , and thus it is guaranteed to have been committed in a majority of machines (inv-1 in §7.1). M2 includes that information in the opcode, to help M1 avoid unnecessary commits. Therefore, two different opcodes are used for *Rmw-id-committed* rep, as a performance optimization.

**Optimization.** We are implementing one more optimization. If the *accepted-log-no* of the *Local-entry* is  $X$  and its  $log-no = Z$ , and if  $Z > X$ , this means that the RMW has been committed (from a helper) in an older  $log-no$ . There is then the potential problem, that the RMW has now grabbed the *KV-pair* for the fresh  $log-no = Z$ , which it did not released when it committed itself on the older  $log-no = X$ . That may delay other local RMWs that are attempting to grab the *KV-pair*, but are waiting. To alleviate this, if we detect this case (*Local-entry.accepted-log-no < Local-entry.log-no*, after receiving an *Rmw-id-committed*reply), we inspect the *KV-pair* and if it's still held in *Proposed* state we revert it to *Invalid* state. (The *KV-pair* cannot be in *Accepted* for that RMW in this case, because that would mean that an RMW was accepted locally,

but then learnt it was committed in a previous *log-no*, which is impossible – from inv-3 § 7.1.)

Note something very important: this optimization means that a *KV-pair* can transition to *Invalid* state without advancing its *last-committed-log-no*.

It's worth noting, that optimizations are not really important for *Rmw-id-committed* replies as they are triggered only once every 5k to 50k committed RMWs.

## 8.2 Why: Log-too-low

The *Log-too-low* means that the RMW was attempting to commit in a *log-no* that has been used by a different RMW. Therefore, the RMW must go and work in a bigger *log-no*. As a result the TSes used so far are meaningless because they refer to a specific *log-no*. The RMW must start fresh. Also it must grab the *KV-pair* from scratch, because the *log-no* it used to have is gone as we committed the RMW included in the *Log-too-low* reply.

## 8.3 Why: Seen-lower-acc

Firstly, why is it okay to handle *Seen-lower-acc* replies with higher priority than *Log-too-high* ones? Because, if anyone has accepted a value locally, then it must be that majority of machines already knows the committed RMWs in previous logs.

There is a case where a *Seen-lower-acc* reply can be turned into an *Ack* purely as a performance optimization. Specifically, on receiving a propose, if the same *rmw-id* is in *Accepted* state, with a lower *proposed-TS* and a lower *accepted-TS*, then we simply return an *Ack* to the proposer. This is correct because returning a *Seen-lower-acc* reply and the *Ack* tell the proposer the same exact thing: broadcast accepts using your TS for the RMW you are doing.

## 8.4 How: Retrying

An RMW may attempt to retry. This happens when receiving a *Seen-higher-prop/Seen-higher-acc* to a propose/accept or when receiving a *Log-too-high* reply.

Here are the steps to retry: First we check if our RMW has already been committed (i.e., registered), (because it may have gotten help). In that case, we need to broadcast commits, to ensure that we will not respond to the client, unless we know that the RMW has been committed by a majority.

Beyond that, there are two cases where the RMW can retry: 1) the *KV-pair* is “still-proposed” i.e., it's in *Proposed* state for the RMW that is attempting to retry or 2) the *KV-pair* is “still-accepted” i.e., it's in *Accepted* state for the RMW that is attempting to retry or, 3) the *KV-pair* is in *Invalid* state, but in a greater *log-no*.

If the *KV-pair* is still-proposed, or *Invalid*, then we simply grab the *KV-pair* transitioning to *Proposed* state (if it's *Invalid*), then we transition the *Local-entry* to *Proposed* state and broadcast proposes. The propose uses a “higher TS” if it

has been triggered by a *Seen-higher-prop/Seen-higher-acc* reply.

**Helping myself.** The case where *KV-pair* is still-accepted is more subtle. The *Local-entry* still transitions to *Proposed* state and broadcasts proposes, but the *KV-pair* remains in *Accepted* state. The *Local-entry* treats this as *Seen-lower-acc* reply, but transitions its *helping-flag* to *Propose-locally-accepted* state.

As with any *Seen-lower-acc* reply, the *Local-entry* records the *accepted-TS* of the local *KV-pair* and compares it with any other *accepted-TS* of any incoming *Seen-lower-acc* reply. If no *Seen-lower-acc* reply, with a higher *accepted-TS* is received, and when inspecting the propose replies, the *Seen-lower-acc* reply is triggered (i.e., we have not received a *Rmw-id-committed/Log-too-low* or a majority of *Acks*) then the *Local-entry*, realizes it must “help itself” and therefore it acts as if it had received a majority of *Acks*. It updates the *accepted-TS* in the *KV-pair* and it broadcasts accepts for its RMW.

Notably, if a *Seen-lower-acc* reply with a higher *accepted-TS* is received then the *helping-flag* is transitioned back to default (which is *Not-helping*).

## 8.5 How: Accepting locally

First let's assume the case where the RMW is not helping, i.e., it's trying to get the *KV-pair* to accept itself, not some other RMW.

**Not helping.** We first check the registered RMWs, if the RMW is already committed then the *Local-entry* transitions to *Bcast-commits* state, so that commits will be broadcast using the *accepted-value* and *accepted-log-no* of the *Local-entry*.

The *KV-pair* will accept locally (i.e., transition to *Accepted* state) if it is still in *Accepted* or *Proposed* state for this RMW, i.e., its *rmw-id* matches, its *log-no* has not moved and its *proposed-TS* is the same as that of the *Local-entry*. If that happens, we then decide the new value created by the RMW, and we store it both in the *Local-entry* and in the *KV-pair* (as *accepted-value*).

If the *KV-pair* does not accept locally, it can be because another propose/accept has been received for a different RMW (or even for the same RMW, if someone else is helping us) with a higher TS. Or it can be because, the *log-no* has been committed. In any case, the *Local-entry* will transition to *Needs-KV-pair* state, so that it attempts to grab the *KV-pair* the next time it is inspected.

**Helping.** Assume that l-RMW helps h-RMW. We are trying to locally accept h-RMW, because we have received a *Seen-lower-acc* reply to a propose message for l-RMW. Note that the *Seen-lower-acc* reply may have originated from the local *KV-pair* in the case where h-RMW is already locally accepted, but was stuck and the backoff counter timed out, so we issued proposes for l-RMW.

We don't need to check the registered RMWs here – it's not wrong to do so– but if h-RMW has been committed, it has to be in the present *log-no* and we will not be able to accept locally anyway. We will not need to broadcast commits for h-RMW, if we learn it already has been committed. In all cases where the local accept of h-RMW fails, we will simply stop helping and revert to working on l-RMW, by transitioning the *Local-entry* to *Needs-KV-pair* state and lowering its *helping-flag*.

There are four cases where we should accept h-RMW locally. 1) The *KV-pair* is still in *Proposed* state for our RMW and exactly as we left it, 2) The *KV-pair* is in *Invalid* state but without having advanced its *last-committed-log-no* (which is extremely rare but possible, as explained in § 8.1), 3) The backoff counter has timeout for an accepted RMW and the *KV-pair* is still in *Accepted* state, for that same RMW, 4) Finally, it is possible that l-RMW got locally accepted, then it retried, then it received a *Seen-lower-acc* reply for h-RMW from a remote machine. If h-RMW has a higher *accepted-TS* than the *accepted-TS* that l-RMW was originally locally accepted with, then l-RMW must help h-RMW.

In all four cases, we accept locally: i.e., the *KV-pair* transitions to *Accepted* state, and it updates its *rmw-id*, its *accepted-TS* and its *accepted-value*. The *Local-entry* transitions to *Accepted* state, and broadcasts accept messages that contain the value and *rmw-id* of h-RMW, but the *TS* used by l-RMW (that's just Paxos rules of helping).

If none of the four cases occur, we immediately stop helping and we transition the *Local-entry* to *Needs-KV-pair* state, such that it will start the l-RMW from scratch. The reason is that if we cannot get h-RMW to be accepted locally, it must mean that the *KV-pair* is grabbed by some other RMW, or the *log-no* has been taken, in either case we need not bother with h-RMW anymore.

## 8.6 How: Committing

**Optimization.** If we receive all the acks (one from every machine including the local one) for the accept, then we do not include the value in the commit message. The optimization is triggered always in All-aboard (so almost in all RMWs in that case), but it has no impact on performance, even though it substantially reduces network usage. The reason is that we are bottlenecked by the CPU and not the network.

The reason we only apply the optimization when *all* remote machines have acked, is that otherwise we would have to send different commit messages to different machines (no value to those who have acked the accept!). However, this is not a good idea: the commit gets inserted in a larger coalesced broadcast message with other commits, accepts, writes etc. That message is then sent to allow remote machines. The invariant is that all remote machines receive the same message. If commits would have variable size w.r.t. the receiver, then we would lose the ability to create these large coalesced messages.

We could also avoid sending the *rmw-id*. However, then the *KV-pair* would have to remember the last-accepted-rmw-id, which is not a big issue (we would just add that field to the *KV-pair*). We would then register that last-accepted-rmw-id, which would also be fine: in case the last-accepted-rmw-id has been updated in the interim, that has to mean that someone else committed the RMW. However, since there was no performance benefit to removing the value from the RMW, we chose to keep the *rmw-id* to help with debugging.

## 8.7 Why: Log-too-high

Here we discuss 1) why a commit message is triggered after receiving repeated *Log-too-high* replies and 2) how committing an RMW only after receiving acks for its commit message helps.

It is possible that Machine M1 dies while issuing commits. M2 is the only machine, that gets a commit and thus M2 blissfully issues a propose in the next slot, but there is a chance that the propose will never get acked: if no other machine is trying to do an RMW on it, then the other machines will not learn of the commit. M2 detects this case after receiving a few consecutive *Log-too-high* replies for the same propose and broadcasts a commit for the previously committed RMW, by reverting the local-entry to *Bcast-commits-from-help* and filling the help from *KV-pair* with the last committed RMW (its *log-no*, value and RMW-id are all known as they must be stored on the *KV-pair*).

Note that commits are only triggered by proposes that get *Log-too-high* nacks. This is because an accept that gets a *Log-too-high* will result in sending propose messages again (i.e., the *Local-entry* will go to *Retry-with-higher-TS* state).

**Pathological case.** It can be that we end up in the following pathological case: a propose gets acked but the accepts gets *Log-too-high*, resulting in retrying proposes which get acked and so on for ever. This is very unlikely, because it would mean that repeatedly proposes get acked by a certain majority and accepts by a different one. One could handle this case by triggering commits of the *last-committed-log-no*. We have not.

**Committing after commit acks.** To mitigate the performance impact, we must ensure that machines do not try to propose too fast (we focus on proposes as they come before accepts and are thus more likely to be “early”). If machines propose timely, after commits have been delivered, then these replies will not trigger in most cases and therefore no harm, no foul.

When Machine M1 broadcasts commits for *log-no* = *X*, this allows M2 to receive and apply the commit and then propose on *log-no* = *X* + 1. However, in the common case, there is enough time for the rest of the machines to have received the commit before they receive M2's propose. Indeed, this is the case in our setup. However, other local threads/sessions within M1, can see the commit much earlier and waste resources on issuing proposes that will receive a *Log-too-high* reply. For this reason, M1 applies the commit to its local *KV-pair*, only

after it has received one commit ack (we actually do it after a majority of acks, but 1 ack will not be any different).

As a result, there is no performance penalty from adding the *Log-too-high* replies. Roughly, a single *Log-too-high* reply is sent once for every 3 thousand committed RMWs.

## 9 All-Aboard

Here we will describe an optimization to Classic Paxos, called All-aboard. All-aboard was proposed in Howard's thesis [3]. We will give a brief description of it, focusing on how to implement it over our existing CP specification. In a nutshell the All-aboard optimization does the following: instead of broadcasting proposes, you can propose locally and then accept locally and broadcast accepts. There are two catches: 1) the accept is successful only if acked by all 2) the accept must have a lower TS than any propose that gets broadcast for the same key and *log-no*.

We will attempt this only in the first time an RMW runs; if the RMW is not successful in any way (i.e., it fails to grab the *KV-pair*, or it fails to accept locally, or it fails to gather all the acks), then it simply executes Paxos instead. 99.7% of completed RMWs are All-aboard, so there is really no reason to try adding it in more places.

When All-aboard triggers, commits are also thinner, as they do not include the value. Despite removing propose messages and cutting most of the payload of commits the per-machine throughput only improves roughly from 5.5 million RMWs per second to 7.5. This is because the CPU is the bottleneck.

### 9.1 All-Aboard Theory

The theory of All-Aboard is presented in Howard's thesis. We will summarize here. The theory is underpinned by a rule called Flexible Paxos and its refinement.

**Flexible Paxos.** It is only required for propose quorums to intersect with accept quorums.

**Refinement.** Actually, it suffices for proposes to intersect *only* with accepts that have smaller TSes.

All-Aboard takes advantage. It sets up a threshold  $TS_{\text{version}}$ . Proposes with a lower TS than the threshold are seen only locally, while accepts must gather all acks. Proposes and accepts with a higher TS than the threshold are seen by a majority.

There are two cases when reasoning a propose must always intersect with accepts with a lower TS: 1) a propose with a TS lower than the threshold will always intersect with "lower accepts", because "lower accepts" are guaranteed to receive all acks. 2) a propose with a TS higher than the threshold will always intersect with "lower accepts", because the propose must be acked by a majority and thus overlaps with all accepts.

This allows us to optimistically start the Paxos command using a low TS (lower than the threshold) and avoid broadcasting proposes, going instead directly to accepts. If we are

not able to gather all ack for the accept, we can then broadcast proposes using a higher TS (higher than the threshold).

**An aside: Singleton Paxos.** Howard's thesis refers also to the dual of this: gathering all acks for proposes and only accepting locally. This does not actually work when we are not able to gather all acks for a propose. Here's why. If we gather all acks for proposes for  $TS = X$  (similarly to all-aboard), then any propose with  $TS > X$ , must gather all acks, otherwise it will not intersect with accepts of  $TS = X$ . Therefore, we cannot do the trick where we start from low TSes, and then revert to bigger ones. Can we then do the opposite? Start from high timestamps and then if the propose fails revert to lower ones? No, because once nodes see a propose, they cannot ack any propose/accept with lower TS (this is often referred to as "promise").

Therefore, for the Singleton algorithm we do not have a way to revert to Classic Paxos, and thus it is not really useful.

### 9.2 Specification

Upon grabbing the *KV-pair* for an RMW for the very first time we transition the *KV-pair* and the *Local-entry* to *Accepted* state, and we perform the actions necessary to accept locally: we calculate the result of the RMW and we update the respective fields of the *KV-pair* and the *Local-entry* (such as *accepted-value*, *accepted-TS*, *rmw-id* etc.). We also raise a flag *all-aboard* inside the *Local-entry*.

Since the *Local-entry* is now in *Accepted* state we inspect its replies periodically as before. We know the *Local-entry* can only move forward if it receives *Acks* from *all* machines, because its *all-aboard* flag is raised. We handle accept replies as follows.

If we find any nack, we transition the *Local-entry* according to the guide in Section 4.5. This is because we need *all* remote machines to *Ack* the accept. Therefore any nack must trigger its effect without waiting for more replies, e.g., any of *Seen-higher-prop/Seen-higher-acc/Log-too-high* will transition the *Local-entry* to *Retry-with-higher-TS* state. If we have received only *Acks* but not all of them, then we simply increment an a special counter called *all-aboard-time-out-counter*, which is a field of the *Local-entry*. If *all-aboard-time-out-counter* reaches a predetermined threshold, then we transition to *Retry-with-higher-TS* state. In the case where the local *KV-pair* is still locally accepted for the current RMW, then the "helping myself" case (described in Section 8.4), gets triggered.

The *all-aboard-time-out-counter* ensures that if a machine is slow (or failed) we will not indefinitely wait for it. Notably this timeout, can be arbitrarily small without any potential to violate correctness. Avoiding false positives however is advisable for performance.

**All-aboard TS.** When executing all-aboard, we must guarantee that the used TS is smaller than the TS of every other propose. This is fairly straightforward; we always use a  $TS_{\text{version}}$

= 2, when running all-aboard. Remember that for CP, we always use a  $TS.version = 3$ . If all-aboard is not successful, then it will run CP (broadcasting proposes) and will have to use a  $TS.version \geq 3$ .

**Note.** A final note – obvious in hindsight – but easy to miss, is that if we already suspect a remote machine to be slow/failed, we should avoid triggering the All-aboard optimization all together. If we do not, and then a machine is unresponsive for a few seconds, then during that time we will be constantly having to wait for the *all-aboard-time-out-counter* to expire. This would be terrible for performance. It is very easy to avoid this. If we have not recently heard of every other machine, we simply execute Classic Paxos for every RMW.

## 10 Adding writes

In this section, we will discuss the implementation specification of ABD writes using carstamps.

Writes need not solve consensus; in contrast concurrent writes can execute and be serialized post-hoc deterministically in each node using logical timestmaps (i.e., TSes). Therefore, writes can fundamentally be implemented more efficiently than RMWs. We can do roughly 5.5 million Classic Paxos RMW/s per machine (5 machines). All-aboard does roughly 7.5 million and ABD writes reach 12 million.

Therefore, there is benefit coupling them together. If the client wants 90% of the time to do simple writes and 10% of the time to do RMWs, then we can increase performance, by using ABD 90% of the time.

To do so, we use a technique called *carstamps* to be able to serialize Paxos RMWs and ABD Writes. Below we dive into how carstamps interact with the Paxos protocol and what are the necessary additions to support them.

**Basic idea.** The *KV-pair* will have a *base-TS*. Writes will increase the *base-TS*. RMWs will choose both a *log-no* and a *base-TS*. The *base-TS* (which is Lamport clock same as all TSes we have discussed) along with the *log-no* comprise the carstamp. For example imagine key A and machines M1, M2, M3, M4, M5. M1 writes A increasing its *base-TS* to 1, M1. M2 performs an RMW on A choosing *log-no* = 1 and *base-TS* = 1, M1. M3 writes A before M2 finishes, increasing the *base-TS* to 2, M3. M2's RMW precedes M3's writes and as such will only be applied by machines that have not seen M3's write. Assume that M2's RMW fails because M4's RMW wins out. M4 may choose *base-TS* = 2, M3, and *log-no* = 1.

### 10.1 Invariants

Firstly, we need to make sure that the issuer of an RMW reads the correct value, and the RMW overwrites the same value in all machines, despite getting helped. We will achieve that effect in the same spirit as we have done so far. When the RMW gets accepted locally, then it selects its *base-TS*. If it is to get committed – help or not – it must be that it gets

committed with that same *base-TS*. This guarantees, that an RMW always gets committed with the same *base-TS* in all machines.

Secondly, we need to make sure that the RMW overwrites the most recently committed write. Here is why this is slightly harder.

**Enforcing the Invariants.** When accepting locally, we choose the value that will be read. If the RMW is committed in that same *log-no*, then that value is set in stone, i.e., helpers cannot change it. For example, imagine that Machine-1 accepts locally, then it immediately becomes unresponsive and Machine-2 helps it. When Machine-1 becomes responsive again it realizes its RMW has been helped, but it needs to know what value to read. There are two invariants of Paxos that help us then read the correct value: the RMW can only be helped in a *log-no* where it accepted locally and the value that was helped was the one accepted locally. Therefore, it is safe to read the locally accepted value. This is described in detail in Section 7.

Therefore, to ensure the second invariant i.e., to maintain linearizability with writes and RMWs, it must be that when accepting locally, the RMW uses a *base-TS* that is bigger or equal than that of any write that had completed before the RMW was issued. I.e. it must never be that a write completed before the RMW began executing, but was not seen by the RMW.

To achieve this we include the *base-TS*, in propose messages, asking other machines if they have seen any more writes. When a remote machine intends to ack the propose (i.e., none of the nack conditions are triggered), then it also inspects its *base-TS*; if the propose's *base-TS* is smaller than what locally stored then the propose is still acked, but the *Ack* reply contains the value and its *base-TS*. This allows the proposer to ensure that it will use a *base-TS* that is at least as big as the biggest write that has completed. Sometimes the propose, will find the ts of writes that have not yet completed. This is okay: if the RMW completes it is as if its base write is also completed (similarly to the second round of an ABD write or read).

### 10.2 What about All-aboard

Unfortunately, All-aboard immediately selects the value that it will overwrite, as it gets immediately accepted locally. This however violated the second invariant that denotes that the RMW must overwrite any completed write. The problem is that there may be completed writes that have not been received locally. The fix would be to add a broadcast round that reads timestamps similar to what ABD-writes do. However, that seems awfully close to having proposes and thus beats the purpose of All-aboard in the first place. The inverse optimization, Singleton Paxos, where proposes are acked by all instead of accepts, would naturally work with carstamps. As

explained, in Section 9, Singleton Paxos is not possible to deploy if all acks cannot always be gathered.

We have not added the first round to All-aboard.

### 10.3 Specification Changes

#### Metadata changes.

- The *KV-pair* has a *base-TS* field, to be used by ABD writes for serialization, but also by RMWs, to serialize with writes.
- The *KV-pair* has an *acc-base-TS* field, which is used to return to proposes when sending them *Seen-lower-acc* replies, along with the *accepted-TS* the *accepted-value* and the *rmw-id* (which is the *rmw-id* of the last accepted RMW). This will allow helpers, to commit an RMW with the correct *base-TS*.
- *Log-too-low* replies include the *base-TS* of the last committed RMW. Finally this is also useful when receiving commits without any value (see §8.6), to know which *base-TS* to commit.
- The *Local-entry* needs a *base-TS* field that specifies the chosen *base-TS* at local-accept time.

#### 10.3.1 Protocol changes

**Sending Proposes.** Proposes now include a *base-TS* field.

**Proposes replies.** On sending propose replies we make the following changes

- The *Seen-lower-acc* reply includes the *acc-base-TS* of the *KV-pair*, which is the *base-TS* of the RMW that may be helped.
- The *Log-too-low* replies also include the *base-TS* of the *KV-pair*.
- A new type of ack is introduced: *Ack-base-TS-stale* which is triggered every time a propose is received that can be acked but contains a low *base-TS* compared to the locally stored. The payload of the reply includes a value and a *base-TS*.

On receiving a propose reply of type *Ack-base-TS-stale*, the proposer overwrites its locally stored value and its *base-TS*, if the *base-TS* of the *Ack-base-TS-stale* is bigger than the locally stored *base-TS*.

**Optimization.** As an optimization, we note on the *Local-entry* (by raising a flag) that the RMW has looked for a fresh *base-TS* and therefore subsequent proposes from this propose (because it may get retried), need not inspect the *base-TS* of remote nodes

**Accepts.** Accept messages include the *base-TS*. This gets written in the *KV-pair*'s *acc-base-TS* field. There is no change to accept replies.

**Commits.** Finally, commits include the *base-TS* of the RMW to be committed. Recall on Section 8.6, we discussed an

optimization where commits can be broadcast without their value, if the accept has seen *Acks* from all machines. Those commits need not include *base-TS* either. The receiver side has stored the *base-TS* in its *acc-base-TS* field of the *KV-pair* and will use that. A potential pitfall: if the *KV-pair* is not still in *Accepted* state, but it has progressed, its *acc-base-TS* should not be used.

## 11 Adding Reads

Using Paxos to perform reads is far from ideal. Paxos delivers exclusive access to a key to one session. This is 1) very costly, as it is fundamentally hard to do and 2) it hinders concurrency. Reads can be fully concurrent and need no coordination with each other. To give you a sense, we can do about 140 million ABD reads per second, in contrast to 27 million CP RMWs.

Below, we will explain ABD reads (from this point on simply “reads”) with carstamps.

### 11.1 Protocol

Firstly the reader broadcasts its locally stored carstamp (*log-no* plus a *base-TS*). The receiver inspects the *base-TS* and *last-committed-log-no* of the *KV-pair* and it sends back one of three possible replies:

1. *Carstamp-too-low*: if the incoming carstamp is lower than the locally stored, then the reply includes the locally stored carstamp along with the locally stored value.
2. *Carstamp-equal*: if the incoming carstamp is equal to the locally stored, then the reply has no payload
3. *Carstamp-too-high*: if the incoming carstamp is higher to the locally stored, then the reply has no payload

The reader collects a majority of answers and reads the highest carstamp received. If the reader is not certain that a majority of machines store the value that is about to be read, then before reading, it first broadcasts a commit (as in a Paxos commit) including the carstamp and the value that will be read.

**Commits and reads.** Earlier we said that commits are needed so that reads know what is committed. What would the alternative be? Could reads simply look at the values that are accepted and read them? After all if a value is accepted by a majority then it is in principle also committed. However, the accepting majority and the majority that replies to a read may not be the exact same. In which case reads will be unable to infer whether a value that is accepted in one of the replies, is indeed accepted by a majority. To solve this problem the reads themselves would have to “help” accepted values by broadcasting accepts. That would complicate the read protocol, because as we know the path from sending accepts to getting a value committed can be pretty complex.

We want to avoid having reads running any part of the Paxos protocol and that’s where commit messages are so

useful. Notably, reads may be forced to broadcast commits, but commits can only get acked by remote nodes.

## 12 Conclusion

This document described in detail a high-performance implementation of Classic Paxos. We also saw how it can be combined with ABD writes and reads and how it can be optimized to include All-aboard Paxos.

## References

- [1] M. Burke, A. Cheng, and W. Lloyd, “Gryff: Unifying consensus and shared registers,” in *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*. Santa Clara, CA: USENIX Association, Feb. 2020, pp. 591–617. [Online]. Available: <https://www.usenix.org/conference/nsdi20/presentation/burke>
- [2] V. Gavrielatos, A. Katsarakis, V. Nagarajan, B. Grot, and A. Joshi, “Kite: Efficient and Available Release Consistency for the Datacenter,” in *Proceedings of the 25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, ser. PPOPP ’20. New York, NY, USA: Association for Computing Machinery, 2020, p. 1–16. [Online]. Available: <https://doi.org/10.1145/3332466.3374516>
- [3] H. Howard, “Distributed consensus revised,” Ph.D. dissertation, University of Cambridge, 2019.
- [4] H. Howard, D. Malkhi, and A. Spiegelman, “Flexible Paxos: Quorum intersection revisited,” 2016.
- [5] L. Lamport, “The part-time parliament,” *ACM Transactions on Computer Systems (TOCS)*, vol. 16, no. 2, pp. 133–169, 1998.
- [6] L. Lamport *et al.*, “Paxos made simple,” *ACM Sigact News*, vol. 32, no. 4, pp. 18–25, 2001.
- [7] I. Moraru, D. G. Andersen, and M. Kaminsky, “There is More Consensus in Egalitarian Parliaments,” in *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*, ser. SOSP ’13. New York, NY, USA: ACM, 2013, pp. 358–372. [Online]. Available: <http://doi.acm.org/10.1145/2517349.2517350>
- [8] D. Ongaro and J. Ousterhout, “In Search of an Understandable Consensus Algorithm,” in *Proceedings of the 2014 USENIX Conference on USENIX Annual Technical Conference*, ser. USENIX ATC’14. Berkeley, CA, USA: USENIX Association, 2014, pp. 305–320. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2643634.2643666>
- [9] B. Reed and F. P. Junqueira, “A Simple Totally Ordered Broadcast Protocol,” in *Proceedings of the 2Nd Workshop on Large-Scale Distributed Systems and Middleware*, ser. LADIS ’08. New York, NY, USA: ACM, 2008, pp. 2:1–2:6. [Online]. Available: <http://doi.acm.org/10.1145/1529974.1529978>