

On Barriers and the Gap between Active and Passive Replication (Full Version)^{*}

Flavio P. Junqueira¹ and Marco Serafini²

¹ Microsoft Research, Cambridge, UK
fpj@apache.org

² Yahoo! Research, Barcelona, Spain
serafini@yahoo-inc.com

Abstract. Active replication is commonly built on top of the atomic broadcast primitive. Passive replication, which has been recently used in the popular ZooKeeper coordination system, can be naturally built on top of the primary-order atomic broadcast primitive. Passive replication differs from active replication in that it requires processes to cross a *barrier* before they become primaries and start broadcasting messages. In this paper, we propose a barrier function τ that explains and encapsulates the differences between existing primary-order atomic broadcast algorithms, namely semi-passive replication and Zookeeper atomic broadcast (Zab), as well as the differences between Paxos and Zab. We also show that implementing primary-order atomic broadcast on top of a generic consensus primitive and τ inherently results in higher time complexity than atomic broadcast, as witnessed by existing algorithms. We overcome this problem by presenting an alternative, primary-order atomic broadcast implementation that builds on top of a generic consensus primitive and uses consensus itself to form a barrier. This algorithm is modular and matches the time complexity of existing τ -based algorithms.

1 Introduction

Passive replication is a popular approach to achieve fault tolerance in practical systems [3]. Systems like ZooKeeper [8] or Megastore [1] use primary replicas to produce state updates or state mutations. Passive replication uses two types of replicas: primaries and backups. A primary replica executes client operations, without assuming that the execution is deterministic, and produces state updates. Backups apply state updates in the order generated by the primary. With active replication, by contrast, all replicas execute all client operations, assuming that the execution is deterministic. Replicas execute a sequence of consensus instances on client operations to agree on a single execution sequence using atomic broadcast (*abcast*). Passive replication has a few advantages such as simplifying the design of replicated systems with non-deterministic operations, *e.g.*, those depending on timeouts or interrupts.

It has been observed by Junqueira *et al.* [9] and Birman *et al.* [2] that using atomic broadcast for passive, instead of active, replication requires taking care of specific constraints. State updates must be applied in the exact sequence in which they have been generated: if a primary is in state A and executes an operation making it transition to state update B , the resulting state update δ_{AB} must be applied to state A . Applying it to a different state $C \neq A$ is not safe because it might lead to an incorrect state, which is inconsistent with the history observed by the primary and potentially the clients. Because a state update is the difference between a new state and the previous, there is a causal dependency between state updates. Unfortunately, passive replication algorithms on top of atomic broadcast (*abcast*) do not necessarily preserve this dependency: if multiple primaries are concurrently present in the system, they may generate conflicting state updates that followers end up applying in the wrong order. Primary-order atomic broadcast (*POabcast*) algorithms, like Zab [9], have additional safety properties that solve this problem. In particular, it implements a *barrier*, the *isPrimary* predicate, which must be crossed by processes that want to broadcast messages.

Interestingly, the only existing passive replication algorithm using consensus as a communication primitive, the semi-passive replication algorithm of Defago *et al.* [7], has linear time complexity in the number of concurrently submitted requests. Recent algorithms for passive replication have constant complexity but they directly implement *POabcast* without building on top of consensus [2,9].

^{*} A conference version of this work (without appendices) appears in the proceedings of the 27th International Symposium on Distributed Computing (DISC) 2013 conference.

Table 1. Time complexity of *POabcast* algorithms presented in this paper - see Sect. 6.4 and 7.2 for detail. We consider the use of Paxos as the underlying consensus algorithm since it has optimal latency [11]. However, only the third solution *requires* the use of Paxos; the other algorithms can use any implementation of consensus. For the latency analysis only, we assume that message delays are equal to Δ . The *Stable periods* column reports the time, in a passive replication system, between the receipt of a client request and its delivery by a single broadcasting primary/leader (c is the number of clients). The *Leader change* column reports idle time after a new single leader is elected by Ω and before it can broadcast new messages.

	Stable periods	Leader change
Atomic broadcast [10]	2Δ	2Δ
τ -based <i>POabcast</i> (Sect. 6.1)	$2\Delta \cdot c$	4Δ
τ -based <i>POabcast</i> with white-box Paxos (Sect. 6.3)	2Δ	4Δ
τ -free <i>POabcast</i> (Sect. 7)	2Δ	4Δ

During our work on the ZooKeeper coordination system [8] we have realized that it is still not clear how these algorithms relate, and whether this trade-off between modularity and time complexity is inherent. This paper shows that existing implementations of passive replication can be seen as instances of the same unified consensus-based *POabcast* algorithm, which is basically an atomic broadcast algorithm with a barrier predicate implemented through a *barrier function* τ we define in this work. The τ function outputs the identifier of the consensus instance a leader process must decide on before becoming a primary.

Existing algorithms constitute alternative implementations of τ ; the discriminant is whether they consider the underlying consensus algorithm as a black-box whose internal state cannot be observed. Our τ -based algorithm exposes an inherent trade off. We show that if one implements τ while considering the consensus implementation as a black box, it is *necessary* to execute consensus instances sequentially, resulting in higher time complexity. This algorithm corresponds to semi-passive replication.

If the τ implementation can observe the internal state of the consensus primitive, we can avoid the impossibility and execute parallel instances. For example, Zab is similar to the instance of our unified algorithm that uses Paxos as the underlying consensus algorithm and implements the barrier by reading the internal state of the Paxos protocol. We experimentally evaluate that using parallel instances almost doubles the maximum throughput of passive replication in stable periods, even considering optimizations such as batching. Abstracting away these two alternatives and their inherent limitations regarding time complexity and modularity is one of the main observations of this paper.

Finally, we devise a τ -free *POabcast* algorithm that makes this trade off unnecessary, since it enables running parallel consensus instances using an unmodified consensus primitive as a black box. Unlike barrier-based algorithms, a process becomes a primary by proposing a special value in the next available consensus instances; this value marks the end of the sequence of accepted messages from old primaries. Table 1 compares the different PO abcast algorithms we discuss in our paper.

Our barrier-free algorithm shows that both active and passive replication can be implemented on top of a black-box consensus primitive with small and well understood changes and without compromising performance.

Differences with the DISC’13 version. This version includes gap handling in Algorithm 3. In addition, Section 6.3 gives more details on how Paxos handles gaps.

2 Related Work

Traditional work on passive replication and the primary-backup approach assumes synchronous links [3]. Group communication has been used to support primary-backup systems; it assumes a $\diamond P$ failure detector for liveness [6]. Both atomic broadcast and *POabcast* can be implemented in a weaker system model, *i.e.*, an asynchronous system equipped with an Ω leader oracle [4]. For example, our algorithms do not need to agree on a new view every time a non-primary process crashes.

Some papers have addressed the problem of reconfiguration: dynamically changing the set of processes participating to the state machine replication group. Vertical Paxos supports reconfiguration by using an external master, which can be a replicated state machine [12]. This supports *primary-backup* systems, defined as replicated systems where

write quorums consist of all processes and each single process is a read quorum. Vertical Paxos does not address the issues of passive replication and considers systems where commands, not state updates, are agreed upon by replicas. Virtually Synchronous Paxos (VS Paxos) aims at combining virtual synchrony and Paxos for reconfiguration [2]. Our work assumes a fixed set of processes and does not consider the problem of reconfiguring the set of processes participating to consensus. Shraer *et al.* have recently shown that reconfiguration can be implemented *on top* of a *POabcast* construction as the ones we present in this paper, making it an orthogonal topic [14].

While there has been a large body of work on group communication, only few algorithms implement passive replication in asynchronous systems with Ω failure detectors: semi-passive replication [7], Zab [9] and Virtually synchronous Paxos [2]. We relate these algorithms with our barrier-based algorithms in Sect. 6.5.

Pronto is an algorithm for database replication that shares several design choices with our τ -free algorithm and has the same time complexity in stable periods [13]. Both algorithms elect a primary using an unreliable failure detector and have a similar notion of epochs, which are associated to a single primary. Epoch changes are determined using an agreement protocol, and values from old epochs that are agreed upon after a new epoch has been agreed upon are ignored. Pronto, however, is an active replication protocol: all replicas execute transactions, and non-determinism is handled by agreeing on a per-transaction log of non-deterministic choices that are application specific. Our work focuses on passive replication algorithms, their difference with active replication protocols, and on the notion of barriers in their implementation.

3 System Model and Primitives

Throughout the paper, we consider an asynchronous system composed of a set $\Pi = \{p_1, \dots, p_n\}$ of processes that can fail by crashing. They implement a passive replication algorithm, executing requests obtained by an unbounded number of client processes, which can also fail by crashing. *Correct* processes are those that never crash. Processes are equipped with an Ω failure detector oracle.

Definition 1 (Leader election oracle) *A leader election oracle Ω operating on a set of processes Π outputs the identifier of some process $p \in \Pi$. Instances of the oracle running on different processes can return different outputs. Eventually, all instances of correct processes permanently output the same correct process.*

Our algorithms build on top of (uniform) consensus, which has the following properties.

Definition 2 (Consensus) *A consensus primitive consists of two operations: $propose(v)$ and $decide(v)$ of a value v . It satisfies the following properties:*

Termination. *If some correct process proposes a value, every correct process eventually decides some value.*

Validity. *If a process decides a value, this value was proposed by some process.*

Integrity. *Every correct process decides at most one value.*

Agreement. *No two processes decide differently.*

Since our algorithms use multiple instances of consensus, *propose* and *decide* have an additional parameter denoting the identifier of the consensus instance.

Primary order atomic broadcast (*POabcast*) is an intermediate abstraction used by our unified passive replication algorithm. *POabcast* provides a broadcast primitive *POabcast* and a delivery primitive *POdeliver*. *POabcast* satisfies all safety properties of atomic broadcast.

Definition 3 (Atomic broadcast) *An atomic broadcast primitive consists of two operations: broadcast and deliver of a value. It satisfies the following properties:*

Integrity *If some process delivers v then some process has broadcast v .*

Total order *If some process delivers v before v' then any process that delivers v' must deliver v before v' .*

Agreement *If some process p_i delivers v and some other process p_j delivers v' , then either p_i delivers v' or p_j delivers v .³*

³ We modified the traditional formulation of agreement to state it as a safety property only.

```

initially
  | replied( $c, t$ ) returns false for all inputs;
  |  $initialized \leftarrow false$ ;
upon isPrimary()
  |  $\Theta \leftarrow \Sigma$ ;
  |  $initialized \leftarrow true$ ;
upon  $\neg$ isPrimary()
  |  $initialized \leftarrow false$ ;
upon receive  $\langle c, t, o \rangle \wedge \neg$ replied( $c, t$ )  $\wedge$  isPrimary()  $\wedge$   $initialized$ 
  |  $\Theta \xrightarrow{o} \langle r, \delta \rangle$ ;
  |  $apply(\delta, \Theta)$ ;
  | POabcast( $\langle \delta, r, c, t \rangle$ );
upon receive operation  $\langle c, t, o \rangle \wedge$  replied( $c, t$ )
  | send stored  $\langle c, t, r \rangle$  to  $c$ ;
upon POdeliver( $\langle \delta, r, c, t \rangle$ )
  |  $\Sigma \leftarrow apply(\delta, \Sigma)$ ;
  | set  $\langle c, t \rangle$  as replied and store  $r$  as last reply to  $c$ ;
  | send  $\langle c, t, r \rangle$  to  $c$ ;

```

Algorithm 1: Passive replication based on *POabcast* - replica

```

initially
  |  $t \leftarrow 0$ ;
upon execute operation  $o$ 
  |  $t \leftarrow t + 1$ ;
  | reliably send  $\langle c, t, o \rangle$  to all replicas;
  | wait for  $\langle c, t, r \rangle$  from any replica;

```

Algorithm 2: Passive replication based on *POabcast* - client c

POabcast extends atomic broadcast by introducing the concept of primary and a *barrier*: the additional *isPrimary()* primitive, which *POabcast* uses to signal when a process is ready to broadcast state updates. This predicate resembles *Prmys* in the specification of Budhiraja *et al.* [3]. However, as failure detectors are unreliable in our model, primary election is also unreliable: there might be multiple concurrent primaries at any given time, unlike in [3].

A *primary epoch* for a process p , or simply a *primary*, is a continuous period of time during which *isPrimary()* is true at p . Multiple primaries can be present at any given time: the *isPrimary()* predicate is local to a single process and multiple primary epochs can overlap in time. Let P be the set of primaries such that at least one value they propose is ever delivered by some process in the current run. A *primary mapping* Λ is a function that maps each primary in P to a unique *primary identifier* λ , which we also use to denote the process executing the primary role. We consider primaries as logical processes: saying that event ϵ occurs at primary λ is equivalent to saying that ϵ occurs at some process p during a primary epoch for p having primary identifier λ .

Definition 4 (Primary order atomic broadcast) A *primary order atomic broadcast primitive* consists of two operations *broadcast*(v) and *deliver*(v), and of a binary *isPrimary()* predicate, which indicates whether a process is a primary and is allowed to broadcast a value. Let Λ be a primary mapping and \prec_{Λ} a total order relation among primary identifiers. Primary order broadcast satisfies the Integrity, Total order, and Agreement properties of atomic broadcast; furthermore, it also satisfies the following additional properties:

Local primary order If λ broadcasts v before v' , then a process that delivers v' delivers v before v' .

Global primary order If λ broadcasts v , λ' broadcasts v' , $\lambda \prec_{\Lambda} \lambda'$, and some process p delivers v and v' , then p delivers v before v' .

Primary integrity If λ broadcasts v , λ' broadcasts v' , $\lambda \prec_{\Lambda} \lambda'$, and some process delivers v , then λ' delivers v before it broadcasts v' .

These properties are partially overlapping, as we show in Appendix A. For example, global primary order is very useful in reasoning about the behaviour of *POabcast*, but it can be implied from the other *POabcast* properties. It is also worth noting that local primary order is weaker than the single-sender FIFO property, since it only holds within a single primary epoch.

The above properties focus on safety. For liveness, it is sufficient to require the following:

Definition 5 (Eventual Single Primary) *There exists a correct process such that eventually it is elected primary infinitely often and all messages it broadcasts are delivered by some process.*

Definition 6 (Delivery Liveness) *If a process delivers v then eventually every correct process delivers v .*

4 Passive Replication from *POabcast*

Before describing our unified *POabcast* algorithm, we briefly show how to implement passive replication on top of *POabcast*.

In passive replication systems, all replicas keep a copy of an object and receive operations from clients. Each client waits for a reply to its last operation op before it submits the following one. Each operation has a unique identifier, comprising the client identifier c and a counter t local to the client.

The pseudocode of the passive replication algorithm is illustrated in Algorithms 1 and 2. This algorithm captures key aspects of practical systems, like the ZooKeeper coordination system. In the case of ZooKeeper, for example, the object is the data tree used to store client information.

A primary replica is a process whose *isPrimary* predicate is true. This predicate is determined by the underlying *POabcast* algorithm. We will see later in the paper that becoming primary entails crossing a *barrier*.

Replicas keep two states: a *committed state* Σ and a *tentative state* Θ . The primary (tentatively) executes an operation o on the tentative state Θ ; the primary remembers the identifier of the last executed operation from each client (modeled by the *replied*(c, t) predicate) and the corresponding reply r . The execution of a new operation generates a state update δ , which the primary broadcasts to the other replicas together with the reply and the unique operation identifier $\langle c, t \rangle$. We use $S \xrightarrow{op} \langle r, \delta_{SQ} \rangle$ to denote that executing op on the state S produces a reply r and a state update δ_{SQ} , which determines a transition from S to a new state Q . When a replica delivers state updates, it applies them onto its committed state Σ .

We show in Appendix B that this algorithm implements a linearizable replicated object. We also discuss a counterexample showing why *POabcast* cannot be replaced with atomic broadcast due to its lack of barriers; it is a generalization of the counterexample discussed in [9].

5 Unified *POabcast* Algorithm Using the Barrier Function

After motivating the use of *POabcast* to obtain natural implementations of passive replication and discussing the role of the barrier predicate *isPrimary*, we introduce our unified τ -based *POabcast* algorithm (Algorithm 3). It uses three underlying primitives: consensus, the Ω leader oracle, and a new *barrier function* τ we will define shortly.

Like typical atomic broadcast algorithms, our *POabcast* algorithm runs a sequence of consensus instances, each associated with an instance identifier [5]. Broadcast values are proposed using increasing consensus instance identifiers, tracked using the *prop* counter. Values are decided and delivered following the consensus instance order: if the last decided instance was dec , only the event $decide(v, dec + 1)$ can be activated, resulting in an invocation of *POdeliver*. This abstracts the buffering of out-of-order decisions between the consensus primitive and our algorithm.

The most important difference between our algorithm and an implementation of atomic broadcast is that it imposes an additional barrier condition for broadcasting messages: it must hold *isPrimary*. In particular, it is necessary for safety that $dec \geq \tau$. The condition that a primary must be a leader according to Ω is only required for the liveness property of *POabcast*: it ensures that eventually there is exactly one correct primary in the system. The barrier function τ returns an integer and is defined as follows.

```

initially
|  $dec \leftarrow 0;$ 
|  $prop \leftarrow 0;$ 
upon POabcast( $v$ )  $\wedge$  isPrimary()
|  $prop \leftarrow \max(prop + 1, dec + 1);$ 
| propose( $v, prop$ );
upon decide( $v, dec + 1$ )
|  $dec \leftarrow dec + 1;$ 
| POdeliver( $v$ );
function isPrimary()
| return ( $dec \geq \tau$ )  $\wedge$  ( $\Omega = p$ );

/* Gap handling */
upon  $\Omega$  changes from  $q \neq p$  to  $p$ 
| forall the  $i \in [dec + 1, \tau]$  do
| | propose( $skip(\tau), i$ );
upon decide( $skip(k), dec + 1$ )
|  $dec \leftarrow k;$ 

```

Algorithm 3: POabcast based on the barrier function and consensus - process p

Definition 7 (Barrier function) Let σ be an infinite execution, Λ a primary mapping in σ , \prec_Λ a total order among the primary identifiers, and λ a primary such that at least one value it proposes is delivered in σ . A barrier function τ for λ returns:⁴

$$\tau = \max\{i : \exists v, p, \lambda' \text{ s.t. } decide_p(v, i) \in \sigma \wedge propose_{\lambda'}(v, i) \in \sigma \wedge \lambda' \prec_\Lambda \lambda\}$$

An actual implementation of the τ function can only observe the finite prefix of σ preceding its invocation; however, it must make sure that its outputs are valid in any infinite extension of the current execution. If none of the values proposed by a primary during a primary epoch are ever delivered, τ can return arbitrary values.

We show in Appendix C that this definition of τ is sufficient to guarantee the additional properties of POabcast compared to atomic broadcast. In particular, it is key to guarantee that the primary integrity property is respected. Local primary order is obtained by delivering elements in the order in which they are proposed and decided.

The key to defining a barrier function is identifying a primary mapping Λ and a total order of primary identifiers \prec_Λ that satisfy the barrier property, as we will show in the following section. There are two important observations to do here. First, we use the same primary mapping Λ and total order \prec_Λ for the barrier function and for POabcast. If no value proposed by a primary p is ever decided, then the primary has no identifier in Λ and is not ordered by \prec_Λ , so the output of τ returned to p is not constrained by Definition 7. This is fine because values proposed by p are irrelevant for the correctness of passive replication: they are never delivered in the POabcast Algorithm 3 and therefore the corresponding state updates are never observed by clients in the passive replication Algorithms 1 and 2. Note also that a primary might not know its identifier λ : this is only needed for the correctness argument.

We call τ a “barrier” because of the way it is used in Alg. 3. Intuitively, a new leader waits until it determines either that some value proposed by a previous primary is decided at some point in time or that no such value will ever be delivered. Until the leader makes such a decision, it does not make further progress. Also, although this blocking mechanism is mainly implemented by the leader, the overall goal is to have all processes agreeing on the outcome for a value, so it resembles the classic notion of a barrier.

Consensus does not guarantee the termination of instances in which only a faulty primary has made a proposal. Therefore, a new primary proposes a special $skip(\tau)$ value to ensure progress for all instances in which no decision has been reached yet, that is, those with instance number between $dec + 1$ and τ . If the skip value is decided, all decision events on instances up to τ are ignored, and therefore the values decided in these instances are not POdelivered.

⁴ Subscripts denote the process that executed the *propose* or *decide* steps.

6 Implementations of the Barrier Function τ

6.1 Barriers with Black-Box Consensus

We first show how to implement τ using the consensus primitive as a black box. This solution is modular but imposes the use of sequential consensus instances: a primary is allowed to have at most one outstanding broadcast at a time. This corresponds to the semi-passive replication algorithm [7].

Let $prop$ and dec be the variables used in Algorithm 3, and let τ_{seq} be equal to $\max(prop, dec)$. We have the following result:

Theorem 8 *The function τ_{seq} is a barrier function.*

Proof. We define Λ as follows: if a leader process p proposes a value $v_{i,p}$ for consensus instance i and $v_{i,p}$ is decided, p has primary identifier $\lambda = i$. A primary has only one identifier: after $v_{i,p}$ is broadcast, it holds $prop > dec$ and $dec < \tau_{seq}$, so $isPrimary()$ stops evaluating to true at p . The order \prec_Λ is defined by ordering primary identifiers as regular integers.

If a process p proposes a value v for instance $i = \max(prop + 1, dec + 1)$ in Algorithm 3, it observes $\tau_{seq} = \max(prop, dec) = i - 1$ when it becomes a primary. If v is decided, p has primary identifier $\lambda = i$. All primaries preceding λ in \prec_Λ have proposed values for instances preceding i , so τ_{seq} meets the requirements of barrier functions. \square

6.2 Impossibility

One might wonder if this limitation of sequential instances is inherent. Indeed, this is the case as we now show.

Theorem 9 *Let Π be a set of two or more processes executing the τ -based $POabcast$ algorithm with an underlying consensus implementation C that can only be accessed through its propose and decide calls. There is no local implementation of τ for C allowing a primary p to propose a value for instance i before p reaches a decision for instance $i - 1$.*

Proof. The proof is by contradiction: we assume that a barrier function τ_c allowing primaries to propose values for multiple concurrent consensus instances exists.

Run σ_1 : The oracle Ω outputs some process p as the only leader in the system from the beginning of the run. Assume that p broadcasts two values v_1 and v_2 at the beginning of the run. For liveness of $POabcast$, p must eventually propose values for consensus instances 1 and 2. By assumption, τ_c allows p to start consensus instance 2 before a decision for instance 1 is reached. Therefore p observes $\tau_c = 0$ when it proposes v_1 and v_2 . The output of τ_c must be independent from the internal events of the underlying consensus implementation C , since τ_c cannot observe them. We can therefore assume that no process receives any message before p proposes v_2 .

Run σ'_1 : The prefix of σ_1 that finishes immediately after p proposes v_2 . No process receives any message.

Run σ_2 : Similar to σ_1 , but the only leader is $p' \neq p$ and the proposed values are v'_1 and v'_2 . Process p' observes $\tau_c = 0$ when it proposes v'_1 and v'_2 .

Run σ'_2 : The prefix of σ_2 that finishes immediately after p' proposes v'_2 . No process receives any message.

Run σ_3 : The beginning of this run is the union of all events in the runs σ'_1 and σ'_2 . No process receives any message until the end of the union of σ'_1 and σ'_2 . The Ω oracle is allowed to elect two distinct leaders for a finite time. Process p (resp. p') cannot distinguish between run σ'_1 (resp. σ'_2) and the corresponding local prefix of σ_3 based on the outputs of the consensus primitive and of the leader oracle. After the events of σ'_1 and σ'_2 have occurred, some process decides v'_1 for consensus instance 1 and v_2 for consensus instance 2.

Regardless of the definition of Λ and \prec_Λ , the output of τ_c in σ_3 is incorrect. Let p and p' have primary identifiers λ and λ' when they proposed v_2 and v'_1 , respectively. If $\lambda \prec_\Lambda \lambda'$, τ_c should have returned 2 instead of 0 when p' became primary. If $\lambda' \prec_\Lambda \lambda$, τ_c should have returned 1 instead of 0 when p became primary. \square

6.3 Barriers with White-Box Paxos

An alternative, corresponding to Zab [9], to avoid the aforementioned impossibility is to consider the internal states of the underlying consensus algorithm. We exemplify this approach considering the popular Paxos algorithm [10]. A detailed discussion of Paxos is out of the scope of this work and we only present a summary for completeness.

Overview of Paxos. In Paxos each process keeps, for every consensus instance, an *accepted value*, which is the most current value it is aware of that might have been decided. A process p elected leader must first read, for each instance, the value that may have been decided upon for this instance, if any. To obtain this value, the leader selects a unique *ballot number* b and executes a *read phase* by sending a read message to all other processes. Processes that have not yet received messages from a leader with a higher ballot number b reply by sending their current accepted value for the instance. Each accepted value is sent attached to the ballot number of the previous leader that proposed that value. The other processes also promise not to accept any message from leaders with ballots lower than b . When p receives accepted values from a majority of processes, it *picks* for each instance the accepted value with the highest attached ballot.

After completing the read phase, the new leader proposes the values it picked as well as its own values for the instances for which no value was decided. The leader proposes values in a *write phase*: it sends them to all processes together with the current ballot number b . Processes accept proposed values only if they have not already received messages from a leader with a ballot number $b' > b$. After they accept a proposed value, they send an acknowledgement to the leader proposing it. When a value has been written with the same ballot at a majority of processes, it is decided.

In a nutshell, the correctness argument of Paxos boils down to the following argument. If a value v has been decided, a majority of processes have accepted it with a given ballot number b ; we say that the proposal $\langle v, b \rangle$ is *chosen*. If the proposal is chosen, no process in the majority will accept a value from a leader with a ballot number lower than b . At the same time, every leader with a ballot number higher than b will read the chosen proposal in the read phase, and will also propose the v .

Integrating the Barrier Function. We modify Paxos to incorporate the barrier function. If a process is not a leader, there is no reason for evaluating τ . Whenever a process is elected leader, it executes the read phase. Given a process p such that $\Omega = p$, let $read(p)$ be the maximum consensus instance for which any value has been picked in the last read phase executed by p . The barrier function is implemented as follows:

$$\tau_{\text{Paxos}} = \begin{cases} \top & \text{iff } \Omega \neq p \vee p \text{ is in read phase} \\ read(p) & \text{iff } \Omega = p \wedge p \text{ is in write phase} \end{cases}$$

The output value \top is such that $dec \geq \tau_{\text{Paxos}}$ never holds for any value of dec . This prevents leaders from becoming primaries until a correct output for τ_{Paxos} is determined.

The sequence of values picked by the new leader p can have gaps. This occurs when p can pick some value for an instance i but it cannot pick any value for an instance $j < i$. The new leader fills such gaps with *no-op* values. This is analogous to the gap handling mechanism in Algorithm 3, where the new primary proposes $skip(\tau)$ values to fill gaps. Therefore, whenever Paxos decides a value *no-op* for an instance i , Algorithm 3 treats this decision as if consensus decided $skip(\tau)$ for i .

We now show that this τ implementation is correct. The proof relies on the correctness argument of Paxos.

Theorem 10 *The function τ_{Paxos} is a barrier function.*

Proof. By the definition of τ_{Paxos} , a process becomes a primary if it is a leader and has completed the read phase. Let λ associate a primary with the unique ballot number it uses in the Paxos read phase if some values it proposes is ever decided, and let \prec_λ be the ballot number order. Paxos guarantees that if any process ever decides a value v proposed by a leader with ballot number smaller than the one of λ , then v is picked by λ in the read phase [10]. This is sufficient to meet the requirements of τ . \square

6.4 Time Complexity of τ -Based *POabcast* with Different Barrier Functions

We now explain the second and third row of Table 1. Just for the analysis, we assume that there are c clients in the system, that the communication delay is Δ , and that Paxos is used as underlying consensus protocol in all our implementations since it is optimal [11].

We first consider the barrier function of Sect. 6.1. If a primary receives requests from all clients at the same time, it will broadcast and deliver the corresponding state updates sequentially. Delivering a message requires 2Δ , the latency of the write phase of Paxos. Since each message will take 2Δ time to be delivered, the last message will be delivered in $2\Delta \cdot c$ time. During leader change, Paxos takes 2Δ time to execute the read phase and 2Δ to execute the write phase if a proposal by the old primary has been chosen and potentially decided in the last consensus instance.

With the barrier function of Sect. 6.3, consensus instances are executed in parallel with a latency of 2Δ . The complexity for leader changes is the same, since the write phase is executed in parallel for all instances up to τ .

Note that the longer leader change time of *POabcast* algorithms compared to atomic broadcast (see Table 1) is due to the barrier: before it becomes a primary, a process must *decide* on all values that have been proposed by the previous primaries and potentially decided (chosen). This is equivalent to executing read and write phases that require 4Δ time. In atomic broadcast, it is sufficient that a new leader *proposes* chosen values from previous leaders.

6.5 Relationship between τ Functions and Existing *POabcast* Algorithms

The *POabcast* algorithm with the barrier function of Sect. 6.1 is similar to semi-passive replication [7] since both enforce the same constraint: primaries only keep one outstanding consensus instance at a time. The time complexity of the two protocols using Paxos as the underlying consensus protocol is the same (Table 1, second row).

If the barrier function implementation selects a specific consensus protocol and assumes that it can access its internal state, as discussed in Sect. 6.1, our barrier-based *POabcast* algorithm can broadcast state updates in the presence of multiple outstanding consensus instances. This is the same approach as *Zab*, and indeed there are many parallels with this algorithm. The time complexity in stable periods is the same (see Table 1, third row). A closer look shows that also the leader change complexity is equal, apart from specific optimizations of the *Zab* protocol. In *Zab*, the read phase of Paxos corresponds to the *discovery phase*; the CEPOCH message is used to implement leader election and to speed up the selection of a unique ballot (or epoch, in *Zab* terms) number that is higher than any previous epoch numbers [9]. After the read phase is completed, the leader decides on all consensus instances until the instance identifier returned by τ_{Paxos} - this is the *synchronization phase*, which corresponds to a write phase in Paxos; in our implementation, the barrier function returns and the leader waits until enough consensus instances are decided. At this point, the necessary condition $dec \geq \tau_{Paxos}$ of our generic *POabcast* construction is fulfilled, so the leader crosses the barrier, becomes a primary, and can proceed with proposing values for new instances. In *Zab*, this corresponds to the *broadcast phase*.

Virtually-synchronous Paxos is also a modified version of Paxos that implements *POabcast* and the τ_{Paxos} barrier function, but it has the additional property of making the set of participating processes dynamic [2]. It has the same time complexity during stable periods and leader changes as in Table 1.

7 *POabcast* Using Consensus Instead of τ for the Barrier

The previous section shows an inherent tradeoff in τ implementations between modularity, which can be achieved by using sequential consensus instances and using consensus as a black box, and performance, which can be increased by integrating the implementation of the barrier function in a specific consensus protocol. In this section, we show that this tradeoff can be avoided through the use of an alternative *POabcast* algorithm.

7.1 Algorithm

Our τ -free algorithm (see Algorithm 4) implements *POabcast*, so it is an alternative to Algorithm 3. The algorithm is built upon a leader election oracle Ω and consensus. The main difference with Algorithm 3 is that the barrier predicate *isPrimary* is implemented using consensus instead of τ : consensus instances are used to agree not only on values,

```

1 initially
2    $tent\text{-}epoch, dec, deliv\text{-}seqno, prop, seqno \leftarrow 0;$ 
3    $epoch \leftarrow \perp;$ 
4    $primary \leftarrow \text{false};$ 
5 upon  $\Omega$  changes from  $q \neq p$  to  $p$ 
6   try-primary();
7 procedure try-primary()
8    $tent\text{-}epoch \leftarrow$  new unique epoch number;
9   propose( $\langle$ NEW-EPOCH,  $tent\text{-}epoch$  $\rangle$ ,  $dec$ );
10 upon decide( $\langle$ NEW-EPOCH,  $tent\text{-}epoch_m$  $\rangle$ ,  $dec$ )
11    $dec \leftarrow dec+1;$ 
12    $epoch \leftarrow tent\text{-}epoch_m;$ 
13    $dec\text{-}array \leftarrow$  empty array;
14    $prop\text{-}array \leftarrow$  empty array;
15    $deliv\text{-}seqno \leftarrow dec;$ 
16   if  $\Omega = p$  then
17     if  $tent\text{-}epoch = tent\text{-}epoch_m$  then
18        $prop \leftarrow dec;$ 
19        $seqno \leftarrow dec;$ 
20        $primary \leftarrow \text{true};$ 
21     else
22        $primary \leftarrow \text{false};$ 
23       try-primary();
24 upon POabcast(  $v$  )
25   propose( $\langle$ VAL,  $v, epoch, seqno$  $\rangle$ ,  $prop$ );
26    $prop\text{-}array[prop] \leftarrow \langle v, seqno \rangle;$ 
27    $prop \leftarrow prop+1;$ 
28    $seqno \leftarrow seqno+1;$ 
29 upon decide( $\langle$ VAL,  $v, epoch_m, seqno_m$  $\rangle$ ,  $dec$ )
30   if  $epoch_m = epoch$  then
31      $dec\text{-}array[seqno_m] \leftarrow v;$ 
32     while  $dec\text{-}array[deliv\text{-}seqno] \neq \perp$  do
33       PDeliver( $dec\text{-}array[deliv\text{-}seqno]$ );
34        $deliv\text{-}seqno \leftarrow deliv\text{-}seqno+1;$ 
35   if  $primary \wedge epoch_m \neq epoch$ 
36      $\wedge prop \geq dec$  then
37      $\langle v', seqno' \rangle \leftarrow prop\text{-}array[dec];$ 
38      $prop\text{-}array[prop] \leftarrow prop\text{-}array[dec];$ 
39     propose( $\langle$ VAL,  $v', epoch, seqno'$  $\rangle$ ,  $prop$ );
40      $prop \leftarrow prop+1;$ 
41   if  $\neg primary \wedge \Omega = p$  then
42     try-primary();
43      $dec \leftarrow dec+1;$ 
44 upon  $\Omega$  changes from  $p$  to  $q \neq p$ 
45    $primary \leftarrow \text{false};$ 
46 function isPrimary()
47   return  $primary;$ 

```

Algorithm 4: Barrier-free *POabcast* using black-box consensus - process p

but also on primary election information. Another difference is that some decided value may not be delivered. This requires the use of additional buffering, which slightly increases the complexity of the implementation.

When a process p becomes leader, it picks a unique epoch number $tent\text{-}epoch$ and proposes a \langle NEW-EPOCH, $tent\text{-}epoch$ \rangle value in the smallest consensus instance dec in which p has not yet reached a decision (lines 5-9). Like in Algorithm 3, we use multiple consensus instances. All replicas keep a decision counter dec , which indicates the current instance in which a consensus decision is awaited, and a proposal counter $prop$, which indicates the next available instance for proposing a value. Another similarity with Algorithm 3 is that decision events are processed following the order of consensus instances, tracked using the variable dec (see lines 10 and 29). Out-of-order decision events are buffered, although this is omitted in the pseudocode.

Every time a NEW-EPOCH tuple is decided, the sender of the message is elected primary and its epoch $tent\text{-}epoch$ is *established* (lines 10-23). When a new epoch is established, processes set their current epoch counter $epoch$ to $tent\text{-}epoch$. If the process delivering the NEW-EPOCH tuple is a leader, it checks whether the epoch that has been just established is its own tentative epoch. If this is the case, the process considers itself as a primary and sets $primary$ to true; else, it tries to become a primary again.

When p becomes a primary, it can start to broadcast values by proposing VAL tuples in the next consensus instances, in parallel (lines 24-28). Ensuring that followers are in a state consistent with the new primary does not require using barriers: all processes establishing $tent\text{-}epoch$ in consensus instance i have decided and delivered the same sequence of values in the instances preceding i . This guarantees that the primary integrity property of *POabcast* is respected.

Processes only PDeliver VAL tuples of the last established epoch until a different epoch is established (lines 29-33, see in particular condition $epoch_m = epoch$). The algorithm establishes the following total order \prec_A of primary identifiers: given two different primaries λ and λ' which picked epoch numbers e and e' respectively, we say that $\lambda \prec_A \lambda'$ if and only if a tuple \langle NEW-EPOCH, e \rangle is decided for a consensus instance n , a tuple \langle NEW-EPOCH, e' \rangle is decided for a consensus instance m , and $n < m$. Suppose that p is the primary λ with epoch number e_λ elected in consensus instance dec_λ . All processes set their current epoch variable e to e_λ after deciding in instance dec_λ .

From consensus instance number $dec_\lambda + 1$ to the next consensus instance in which a NEW-EPOCH tuple is decided, processes decide and deliver only values that are sent from λ and included in VAL tuples with $epoch_m = e_\lambda$. Replicas thus deliver messages following the order \prec_A of the primaries that sent them, fulfilling the global primary order property of *POabcast*.

The additional complexity in handling VAL tuples is necessary to guarantee the local primary order property of *POabcast*. VAL tuples of an epoch are not necessarily decided in the same order as they are proposed. This is why primaries include a sequence number *seqno* in VAL tuples. In some consensus instance, the tuples proposed by the current primary might not be the ones decided. This can happen in the presence of concurrent primaries, since primaries send proposals for multiple overlapping consensus instances without waiting for decisions. If a primary is demoted, values from old and new primaries could be interleaved in the sequence of decided values for a finite number of instances. All processes agree on the current epoch of every instance, so they do not deliver messages from other primaries with different epoch numbers. However, it is necessary to buffer out-of-order values from the current primary to deliver them later. That is why processes store decided values from the current primary in the *dec-array* array (line 31), and deliver them only if a continuous sequence of sequence numbers, tracked by *deliv-seqno*, can be delivered (lines 32-34).

Primaries also need to resend VAL tuples that could not be decided in the correct order. When values are proposed, they are stored in the *prop-array* following the sequence number order; this buffer is reset to the next ongoing consensus instance every time a new primary is elected. Primaries resend VAL tuples in lines 35-40. Primaries keep a proposal instance counter *prop*, indicating the next consensus instance in which values can be proposed. If an established primary has outstanding proposals for the currently decided instance *dec*, it holds $prop \geq dec$. In this case, if the decided VAL tuple is not one such outstanding proposal but has instead been sent by a previous primary, it holds that $epoch_m \neq epoch$. If all the previous conditions hold, the established primary must resend the value that has been skipped, $prop\text{-array}[dec].v'$, using the same original sequence number $prop\text{-array}[dec].seqno'$ in the next available consensus instance, which is *prop*.

The arrays *dec-array* and *prop-array* do not need to grow indefinitely. Elements of *dec-array* (resp. *prop-array*) with position smaller than *deliv-seqno* (resp. *dec*) can be garbage-collected.

For liveness, a leader which is not a primary keeps trying to become a primary by sending a NEW-EPOCH tuple for every consensus instance (lines 22-23). The *primary* variable is true if a leader is an established primary. It stops being true if the primary is not a leader any longer (lines 44-45).

Algorithm 4 correctly implements *POabcast*, as shown in the Appendix D.

7.2 Time Complexity

As before, we use Paxos for the consensus algorithm and assume a communication delay of Δ . During stable periods, the time to deliver a value is 2Δ , which is the time needed to execute a Paxos write phase. When a new leader is elected, it first executes the read phase, which takes 2Δ . Next, it executes the write phase for all instances in which values have been read but not yet decided, and for one additional instance for its NEW-EPOCH tuple. All these instances are executed in parallel, so they finish within 2Δ time. After this time, the new leader crosses the barrier, becomes a primary, and starts broadcasting new values.

8 Experimental Evaluation

Our τ -free algorithm combines modularity with constant time complexity. Since our work was motivated by our work on systems like ZooKeeper, one might wonder whether this improvement has a practical impact. Current implementations of replicated systems can reach some degree of parallelism even if they execute consensus instances sequentially. This is achieved through an optimization called *batching*: multiple clients requests are aggregated in a batch and agreed upon together using a single instance. Even in presence of batching, we found that there is a substantial advantage of running multiple consensus instances in parallel.

We implemented two variants of the Paxos algorithm, one with sequential consensus instances and one with parallel ones, and measured the performance of running our *POabcast* algorithms on top of it. We consider fault-free runs in which the leader election oracle outputs the same leader to all processes from the beginning. We used three replicas




Fig. 1. Latency and throughput with micro benchmarks. Request and state update sizes were set to 1kb, which is the typical size observed in ZooKeeper. Both protocols use batching

and additional dedicated machines for the clients; all servers are quad-core 2.5 GHz CPU servers with 16 GB of RAM connected through a Gigabit network.

The experiments consist of micro-benchmarks, in which the replicated object does nothing and has an empty state. These benchmarks are commonly used in the evaluation of replication algorithms because they reproduce a scenario in which the replication protocol, rather than execution, is the bottleneck of the system. In systems where execution is the bottleneck, using a more efficient replication algorithm has no impact on the performance of the system.

We used batching in all our experiments. With sequential consensus instances, we batch all requests received while a previous instance is ongoing. In the pipelined version, we start a new consensus instance when either the previous instance is completed or b requests have been batched. We found $b = 50$ to be optimal. Every measurement was repeated five times at steady state, and variances were negligible.

Figure 1 reports the performance of the two variants with messages (requests and state updates) of size 1 kB, which is a common state update size for ZooKeeper and Zab [9]. Each point in the plot corresponds of a different number of clients concurrently issuing requests.

The peak throughput with the parallel consensus instances is almost two times the one with sequential instances. The same holds with messages of size 4 kB. The difference decreases with smaller updates than the ones we observe in practical systems like ZooKeeper. In the extreme case of empty requests and state updates, the two approaches have virtually the same request latency and throughput: they both achieve a maximum throughput of more than 110 kops/sec and a minimum latency of less than 0.5 ms.

These results show that low time complexity (see Table 1) is very important for high-performance passive replication. When there is little load in the system, the difference in latency between the two variants is negligible. In fact, due to the use of batching, running parallel consensus instances is not needed. Since clients can only submit one request at a time, we increase load by increasing the number of clients concurrently submitting requests; this is equivalent to increasing the parameter c of Table 1. As the number of clients increase, latency grows faster in the sequential case, as predicted by our analysis. With sequential consensus instances, a larger latency also results in significantly worse throughput compared to the parallel variant due to lower network and CPU utilization.

9 Conclusions

Some popular systems such as ZooKeeper have used passive replication to mask crash faults. We have shown that the barrier predicate *isPrimary* is a key element differentiating active and passive replication, and its implementation significantly impacts the complexity and modularity of the algorithm. We have shown how to implement passive replication on top of *POabcast*. By making leader processes cross the barrier before becoming primaries, we prevent state updates from being decided and applied out of order.

We then extracted a unified algorithm for implementing *POabcast* using the barrier function that abstract existing approaches. The barrier function is a simple way to understand the difference between passive and active replication, as well as the characteristics of existing *POabcast* algorithms, but it imposes a tradeoff between parallelism and modularity. We have proposed an algorithm that does present such a limitation by not relying upon a barrier function and yet it guarantees the order of state updates according to the succession of primaries over time. This algorithm is different from existing ones in its use of consensus, instead of barrier functions, for primary election.

Acknowledgement We would like to express our gratitude to Alex Shraer and Benjamin Reed for the insightful feedback on previous versions of the paper, and to Daniel Gómez Ferro for helping out with the experiments.

References

1. Jason Baker, Chris Bond, James Corbett, JJ Furman, Andrey Khorlin, James Larson, Jean-Michel Léon, Yawei Li, Alexander Lloyd, and Vadim Yushprakh. Megastore: Providing scalable, highly available storage for interactive services. In *CIDR*, volume 11, pages 223–234, 2011.
2. Ken Birman, Dahlia Malkhi, and Robbert Van Renesse. Virtually synchronous methodology for dynamic service replication. Technical Report MSR-TR-2010-151, Microsoft Research, 2010.
3. Navin Budhiraja, Keith Marzullo, Fred B. Schneider, and Sam Toueg. *The primary-backup approach*, pages 199–216. ACM Press/Addison-Wesley, 1993.
4. Tushar Deepak Chandra, Vassos Hadzilacos, and Sam Toueg. The weakest failure detector for solving consensus. *Journal of the ACM*, 43(4):685–722, 1996.
5. Tushar Deepak Chandra and Sam Toueg. Unreliable failure detectors for reliable distributed systems. *Journal of the ACM*, 43(2):225–267, March 1996.
6. Gregory V. Chockler, Idit Keidar, and Roman Vitenberg. Group communication specifications: a comprehensive study. *ACM Computing Surveys*, 33(4):427–469, 2001.
7. Xavier Défago and André Schiper. Semi-passive replication and lazy consensus. *Journal of Parallel and Distributed Computing*, 64(12):1380–1398, 2004.
8. Patrick Hunt, Mahadev Konar, Flavio P. Junqueira, and Ben Reed. Zookeeper: Wait-free coordination for Internet-scale systems. In *USENIX Annual Technical Conference*, pages 145–158, 2010.
9. Flavio P. Junqueira, Benjamin Reed, and Marco Serafini. Zab: High-performance broadcast for primary-backup systems. In *IEEE Conference on Dependable Systems and Networks*, pages 245–256, 2011.
10. Leslie Lamport. The part-time parliament. *ACM Transactions on Computing Systems (TOCS)*, 16(2):133–169, 1998.
11. Leslie Lamport. Lower bounds for asynchronous consensus. In *Future Directions in Distributed Computing Workshop*, pages 22–23, October 2003.
12. Leslie Lamport, Dahlia Malkhi, and Lidong Zhou. Vertical paxos and primary-backup replication. In *ACM Symposium on Principles Of Distributed Computing*, pages 312–313, 2009.
13. Fernando Pedone and Svend Frolund. Pronto: A fast failover protocol for off-the-shelf commercial databases. In *IEEE Symposium on Reliable Distributed Systems*, pages 176–185, 2000.
14. Alexander Shraer, Benjamin Reed, Dahlia Malkhi, and Flavio Junqueira. Dynamic reconfiguration of primary/backup clusters. In *USENIX Annual Technical Conference*, pages 425–438, 2012.

A Alternative Definitions of *POabcast*

There are alternative formulation of *POabcast*. For example, while global primary order is important for understanding the properties of *POabcast*, it can be implied by the other properties of *POabcast*.

Theorem 11 Global primary order can be implied from the other properties of *POabcast*.

Proof. Assume that a process p_i delivers two values v and v' , which were broadcast by the primaries λ and λ' respectively. Also, assume that $\lambda \prec_A \lambda'$. Global primary order requires that p_i delivers first v and then v' . Assume by contradiction that p_i delivers the two values in the opposite order. By primary integrity, λ' delivers v before broadcasting v' . By total order, λ' delivers v' before v . However, by integrity, λ' cannot deliver v' since it did not broadcast it yet, a contradiction. \square

B Correctness of Passive Replication on Top of *POabcast*

This section shows the correctness of Algorithm 1.

Formalism and rationale. Replicas keep a copy of a shared object. The shared object can be nondeterministic: there may be multiple state updates $\delta_{SQ_1}, \delta_{SQ_2}, \dots$ and replies r_1, r_2, \dots such that $S \xrightarrow{OR} \langle \delta_{SQ_i}, r_i \rangle$. We omit subscripts from state updates when these are not needed.

Applying state updates responds to the constraints we discussed in the introduction: state updates can only be applied on the state from which they were generated. Formally, given a state update δ_{SQ} such that $S \xrightarrow{OR} \langle r, \delta_{SQ} \rangle$ for any value of r , we have the following.

$$apply(\delta_{SQ}, R) = \begin{cases} Q & \text{iff } S = R \\ \perp & \text{otherwise} \end{cases}$$

We also use the notation $S \xrightarrow{OP} Q$ to denote the tentative execution of the operation op on a tentative state S resulting in a transition to a tentative state Q .

The notation \perp denotes an undefined state. A correct passive replication system shall never reach such undefined state. This formalizes the difference between agreeing on state updates and operations. The execution of operations never leads the system to an undefined state. The order of their execution may be constrained by consistency requirements, for example linearizability, but not by the semantics of the replicated object. State updates are different: they need to be applied exactly on the state in which they were generated or an undefined state \perp can be reached.

The reason for adopting *POabcast* instead of regular atomic broadcast in Algorithm 1 can be illustrated with an example (see Fig. 2). Consider an algorithm similar to Algorithm 1 such that atomic broadcast is used instead of *POabcast*. Assume also that *isPrimary* is implemented using reliable leader election, that is, only one primary process exists from the beginning, and it is a correct process. Consider a run in which the initial state of all processes is A . A process p_i is elected leader; it tentatively executes two operations op_1 and op_2 such that $A \xrightarrow{op_1} B \xrightarrow{op_2} C$, and it atomically broadcasts state updates δ_{AB} and δ_{BC} . Assume now that p_i crashes and p_j becomes leader. p_j did not deliver any message so far, so it is still in state A . Process p_j executes operation op_3 such that $A \xrightarrow{op_3} D$, and broadcasts δ_{AD} .

In this run there are no concurrent primaries. Yet, atomic broadcast leads to an incorrect behavior since it could deliver the sequence of state updates $\langle \delta_{AD}, \delta_{BC} \rangle$. Since all delivered state updates are executed, a backup would execute $apply(\delta_{AD}, A)$, transitioning to state D , and then execute $apply(\delta_{BC}, D)$, transitioning to an undefined state \perp that may not be reachable in correct runs. Note that even using FIFO atomic broadcast would not solve the problem. It guarantees that if a message from a process is delivered then all messages previously broadcast from that process are previously delivered. In our example, FIFO atomic broadcast could still deliver the sequence of state updates $\langle \delta_{AD}, \delta_{AB}, \delta_{BC} \rangle$, which also leads to \perp .

POabcast prevents this incorrect run. First, it determines a total order between primary identifiers: for example, this could order the primary identifier of p_i before the one of p_j . Global primary order ensures that followers apply state updates following this order: the updates sent by p_j are delivered after the ones sent by p_i . Primary integrity guarantees

that, before broadcasting a state update, a primary delivers all state updates from previous primaries, thus making sure that all followers can safely apply the state updates it broadcasts. For example, p_j would apply the sequence of state updates $\langle \delta_{AB}, \delta_{BC} \rangle$, reaching state C before it executes op_3 and starts sending state updates: therefore, it will not broadcast the state update δ_{AC} .

Correctness. The following is a more general correctness result.

Theorem 12 *Passive replication based on primary order atomic broadcast satisfies linearizability.*

Proof. Since *POabcast* ensures the agreement and total order properties of atomic broadcast, it delivers state updates in a consistent total order \prec_{Δ} . We only need to consider delivered state updates since they are the only ones whose effects are visible to clients.

Consider a state update $\delta = \delta_{SQ}$ that is delivered by a process p and has position i in the sequence of state updates delivered by p . Let δ be broadcast by some primary λ that has generated δ_{SQ} by tentatively executing some operation op on its local tentative state $\Theta = S$. We show that, for any value of $i \geq 0$ that:

- (i) No state update preceding δ in \prec_{Δ} is generated by tentatively executing op ;
- (ii) S is the state p reaches after delivering the state updates preceding δ in \prec_{Δ} .

Predicate (i) establishes a mapping of delivered state updates to operations; therefore, \prec_{Δ} can be used to determine an execution order \prec_o of operations. Predicate (ii) shows that replies to clients are legal and consistent with \prec_o . If both predicates hold, then replicas transition from the state S to the state Q upon delivering the state update δ_{SQ} . Since $S \xrightarrow{op} \langle Q, r \rangle$, Q is a legal state and r a legal reply. Therefore, linearizability follows by showing these two properties for any value of $i \geq 0$.

Both properties (i) and (ii) trivially hold for $i = 0$ since all processes start from the same initial state. We now consider the case $i > 0$.

Let δ' be the last state update delivered by λ before POabcasting any value. It follows from the integrity property of atomic broadcast that $\delta' \prec_{\Delta} \delta$. Let $k < i$ be the position of δ' in the sequence of operations delivered by λ . We use the notation $(\delta', \delta)_{\prec_{\Delta}}$ to indicate all operations delivered by λ between δ' and δ , which have positions in (k, i) .

We now consider two cases.

Case 1: All state updates in $(\delta', \delta)_{\prec_{\Delta}}$ are broadcast by λ . In our algorithm, when *POabcast* delivers a state update, it also delivers the identifier of the request that was executed to obtain the update. Therefore, λ knows the identifier of all requests related to all state updates preceding and including δ' in the delivery order \prec_{Δ} , and consider them as replied. Primaries generate state updates only for operations that have not been replied, so op is not executed to generate any state updates up to and including δ' in the delivery order \prec_{Δ} . The primary λ also knows the identifiers of all the operations it executed to generate the state updates in $(\delta', \delta]_{\prec_{\Delta}}$. It executes op and broadcasts δ only if it knows that op was not already executed, ensuring property (i).

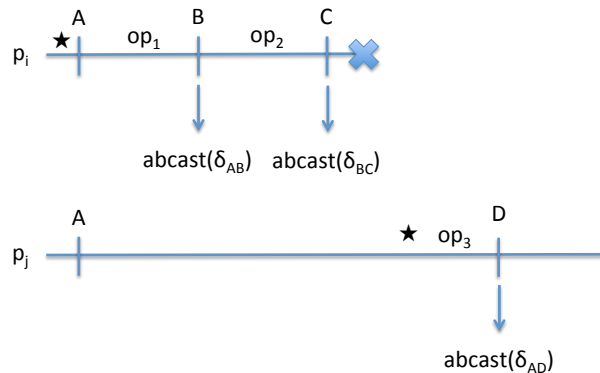


Fig. 2. In this run, atomic broadcast allows delivering the incorrect sequence of state updates $\langle \delta_{AD}, \delta_{BC} \rangle$. The star denotes becoming a primary. A, B, C, D are tentative states.

For property (ii), both p and λ deliver and apply the same sequence of state updates preceding and including δ' in \prec_{Δ} . It follows from Local Primary Order that p commits all state updates in $(\delta', \delta]_{\prec_{\Delta}}$ in the same order as they are generated by λ .

Case 2: There exists a state update in $(\delta', \delta)_{\prec_{\Delta}}$ that has been broadcast by another primary $\lambda' \neq \lambda$. This case leads to a contradiction. If $\lambda \prec_{\Delta} \lambda'$, a contradiction follows from global primary order, which requires that p deliver δ before δ' . If $\lambda' \prec_{\Delta} \lambda$, a contradiction follows from primary integrity: since p delivers δ'' before δ , λ must deliver δ'' before δ' , so $\delta'' \prec_{\Delta} \delta'$. \square

Liveness directly follows from the liveness property of *POabcast*. All requests from correct clients eventually complete, since correct clients resend their requests to the current primary.

C Correctness of the Unified *POabcast* Algorithm

We now show the correctness of Algorithm 3.

Theorem 13 *Primary order atomic broadcast based on the barrier function is correct.*

Proof. Processes exchange values using consensus instances; this ensures that all atomic broadcast properties are respected. Local primary order follows by three observations. First, if a process proposes a value for an instance i , it will never propose a value for an instance $j < i$ since *prop* only increases. Second, values are decided and delivered according to the order of consensus instance identifiers. Three, skip values do not determine gaps in the sequence of values proposed by a primary.

For primary integrity, consider that in order to broadcast a value, a process needs to become primary first. Let λ be the primary identifier of a process. By definition of barrier function, the last invocation of τ returns to λ a value i greater or equal than the maximum identifier of a consensus instance in which a value proposed by any previous primary $\lambda' \prec_{\Delta} \lambda$ is decided. Before broadcasting any message, λ decides on all instances up to i and delivers all values proposed by previous primaries that are delivered. If λ proposes a skip value and this is decided for an instance number $j < i$, no value from previous primaries decided between j and i is delivered.

Global primary order follows from the other properties of *POabcast*.

The liveness properties of the algorithm follow from the Termination property of the underlying consensus primitive and from the Leader election oracle property of Ω . If some previous primary is faulty, the consensus instances it initiated may not terminate. In this case, a new correct primary proposes skip values to guarantee progress. Since the new primary proposes values starting from instance τ , skipping to τ is sufficient to guarantee liveness when eventually a correct primary is elected. \square

D Correctness of the Barrier-Free Algorithm

This section shows that Algorithm 4 correctly implements *POabcast*.

Definition 14 (Total order of primaries) *Given two primaries λ and λ' with epoch numbers e and e' respectively, we say that $\lambda \prec_{\Delta} \lambda'$ if and only if a tuple $\langle \text{NEW-EPOCH}, \text{epoch} \rangle$ is decided for consensus instance number n , a tuple $\langle \text{NEW-EPOCH}, \text{epoch}' \rangle$ is decided for consensus instance number m , and $n < m$.*

The fact that this is a total order derives directly from the following two observations. First, all primaries are ordered: a leader becomes a primary with epoch number *epoch* only after its *primary* variable is set to true, and this occurs only after the $\langle \text{NEW-EPOCH}, \text{epoch} \rangle$ is decided for some consensus instance. Second, a primary λ proposes a $\langle \text{NEW-EPOCH}, \text{epoch} \rangle$ tuple only once and in a single consensus instance.

In the following, we use the notation $\lambda(\text{epoch})$ to denote the primary with epoch number *epoch*.

Lemma 1. *If a process p enters a value v sent by a primary $\lambda(\text{tent-epoch}_m)$ in $\text{da}[\text{seqno}_m]$ upon deciding on instance number n , then $\langle \text{NEW-EPOCH}, \text{tent-epoch}_m \rangle$ is the latest NEW-EPOCH tuple decided before instance n and the tuple $\langle \text{VAL}, v, \text{tent-epoch}_m, \text{seqno}_m \rangle$ is decided for instance n .*

Proof. Process p adds v to da upon deciding a $\langle \text{VAL}, v, \text{epoch}_m, \text{seqno}_m \rangle$ tuple such that $\text{epoch}_m = \text{epoch}$. The variable epoch is set to tent-epoch_m upon deciding on the last $\langle \text{NEW-EPOCH}, \text{tent-epoch}_m \rangle$ tuple, so $\text{epoch}_m = \text{epoch} = \text{tent-epoch}_m$. \square

Lemma 2. *If a process delivers a value v upon deciding on consensus instance n , every process that decides on n delivers v .*

Proof. All processes decide the same values in the same consensus instance order. A value v is delivered only upon deciding on a VAL tuple containing v . The decision of whether to deliver a value or not only depends on variables whose value is consistent across processes, since they are modified consistently by all processes every time a decision for the next instance is reached. Therefore, every process deciding on instance n delivers the same VAL tuple and takes a consistent decision on delivering v . \square

Lemma 3. *A process delivers a value v in at most one consensus instance.*

Proof. Consider all values that are broadcast to be distinct. A value v can only be delivered upon deciding a VAL tuple that contains v . The primary λ that has broadcast v can include v in more than one VAL tuple. In fact, λ stores v in $ta[l_i]$ every time it proposes v for consensus instance l_i . The value $ta[l_i]$ is proposed, however, only if v is not delivered upon deciding on instance l_i . \square

Lemma 4. *Algorithm 4 satisfies Integrity: If some process delivers v then some process has broadcast v .*

Proof. A value v is delivered only upon deciding on a VAL tuple containing v . These tuples are proposed either when v is broadcast, or when another VAL tuple is decided. In the latter case, v is taken from the array ta . However, only values that are broadcast are entered in ta . \square

Lemma 5. *Algorithm 4 satisfies Total Order: If some process delivers v before v' then any process that delivers v' must deliver v before v' .*

Proof. Let p_i deliver v before v' and p_j deliver v' . We now show that p_j delivers v before v' .

Processes deliver values only upon deciding on a consensus instance. From Lemma 3, this happens only once per value. Since p_i delivers v and v' , this occurs upon deciding on instances n and m respectively, with $n < m$. Since p_j delivers v' , it decides on instance m ; this happens only after deciding on instance n . From Lemma 2, p_j delivers v upon deciding on instance n . \square

Lemma 6. *Algorithm 4 satisfies Agreement: If some process p_i delivers v and some other process p_j delivers v' , then either p_i delivers v' or p_j delivers v .*

Proof. Processes deliver values only upon deciding on a consensus instance. From Lemma 3, this happens only once per value. Let n (resp. m) be the consensus instance in which p_i (resp. p_j) decides upon and deliver v (resp. v'). If $n < m$, then p_j decides on n before deciding on m , so it delivers v from Lemma 2. Otherwise, p_i decides on m before n and delivers v' . \square

Lemma 7. *Algorithm 4 satisfies Local Primary Order: If a primary λ broadcasts v before v' , then a process that delivers v' delivers v before v' .*

Proof. Assume by contradiction that a primary λ broadcasts a value v followed by a value v' , and a process p delivers v' but not v . Processes deliver values only upon deciding on a consensus instance. From Lemma 3, this happens only once per value; from Lemma 2, this happens, given a value, when deciding on the same consensus instance for all processes. Let n be the consensus instance whose decision triggers the delivery of v' , and let $\langle \text{VAL}, v', \text{epoch}_m, \text{seqno}_m \rangle$ be the tuple decided for n . The value v' is stored in the array $da[l]$ for some index l ; the array da was reset the last time a NEW-EPOCH tuple was decided, say at consensus instance $m < n$.

Upon deciding on m , the value of s is set to $m+1$. After deciding on m and until the decision on n , s is incremented only when the value $da[s]$ is delivered. Therefore, before delivering v' , a process has delivered all elements of da with indexes in $[m+1, l]$. From Lemma 1, all values v'' added in $da[\text{seqno}_m]$ are sent by λ , which includes seqno_m in

the VAL tuple based on its counter ls - this holds even if the VAL tuple is resent more than once by λ , since the value of the $seqno_m$ field is determined by reading the original index stored in the ta array.

λ sets ls to $m + 1$ upon election and increments every time a new tuple is broadcast. If λ broadcasts v before v' , it has sent a VAL tuple with $seqno_m \in [m + 1, l - 1]$. The value of $da[seqno_m]$ held by a process after deciding on m and before deciding on l can only be v . Therefore, if a process delivers v' , it also delivers v . \square

Lemma 8. *Algorithm 4 satisfies Global Primary Order: If a primary λ broadcasts v , λ' broadcasts v' , and $\lambda \prec_A \lambda'$, then a process that delivers both v and v' delivers v before v' .*

Proof. By definition, if $\lambda \prec_A \lambda'$ then the NEW-EPOCH tuple of λ is decided for a consensus instance n preceding the one, m , in which the NEW-EPOCH tuple of λ' is decided. Each primary sends a NEW-EPOCH tuple only once. When a process delivers v' , it has it stored in the da array, which is reset every time a NEW-EPOCH tuple is delivered. From Lemma 1, v' is delivered upon deciding on a consensus instance following m . Similarly, a process delivers v upon deciding on a consensus instance $l \in [n + 1, m - 1]$. Since decisions on consensus instances are taken sequentially, v is delivered before v' . \square

Lemma 9. *Algorithm 4 satisfies Primary Integrity: If a primary λ broadcasts v , λ' broadcasts v' , $\lambda \prec_A \lambda'$, and some process delivers v , then λ' delivers v before it broadcasts v' .*

Proof. By definition, if $\lambda \prec_A \lambda'$ then the NEW-EPOCH tuple of λ is decided for a consensus instance n preceding the one, m , in which the NEW-EPOCH tuple of λ' is decided. Each primary sends a NEW-EPOCH tuple only once. When a process delivers v , it has it stored in the da array, which is reset every time a NEW-EPOCH tuple is delivered. From Lemma 1, v' is delivered upon deciding on a consensus instance $l \in [n + 1, m - 1]$. The primary λ' proposes its NEW-EPOCH tuple for instance m after deciding on instance l . By Lemma 2, λ' delivers v upon deciding on instance l , before sending its NEW-EPOCH tuple, and thus before becoming a primary and broadcasting values. \square

Lemma 10. *Eventually, there exists a correct primary λ such that for any other primary λ' ever existing, $\lambda' \prec_A \lambda$*

Proof. Let p be a correct process that is eventually indicated as a leader by Ω . Upon becoming a leader, p chooses a new epoch number $tent\text{-}epoch$ and sends a NEW-EPOCH tuple for its current instance d .

Process p continues sending NEW-EPOCH tuples for all consensus instances from d on until its tuple is accepted and it becomes a primary. We show this by induction. For each instance $n \geq d$ in which its NEW-EPOCH tuple is not decided, the following can happen. If another NEW-EPOCH tuple is decided with an epoch number $tent\text{-}epoch_m \neq tent\text{-}epoch$, then p tries to become a primary in the next instance $n + 1$. If a VAL tuple is decided, p is not yet a primary so $primary$ is false. Also in this case, p tries to become a primary in the next instance $n + 1$. Eventually, p remains the only leader in the system. All other processes either crash or have $primary$ set to false because they are not currently leader. Therefore, eventually a NEW-EPOCH tuple from p is decided, p becomes the primary λ , and it is the last NEW-EPOCH tuple decided in any consensus instance. According to the definition of the total order of primaries, λ is the maximum. \square

Lemma 11. *Algorithm 4 satisfies Eventual Single Primary: There eventually exists a correct primary λ such that there is no subsequent primary λ' with $\lambda \prec_A \lambda'$ and all messages λ broadcasts are eventually delivered by all correct processes.*

Proof. It follows from Lemma 10 that there exists a $\lambda = \lambda(epoch)$ that is maximum in the total order of primaries. By definition, its $\langle \text{NEW-EPOCH}, epoch \rangle$ tuple is the last NEW-EPOCH tuple decided for an instance n . After deciding on n , λ proposes all tuples that are broadcast using a sequence number ls that starts from $n + 1$ and is increased by one every time a new value is broadcast.

We now show that all tuples proposed by λ are eventually delivered by induction on ls , starting from $n + 1$. Assume by induction that all values proposed by λ with sequence number up to $ls - 1$ are eventually delivered (trivially true if $ls = n + 1$ since no value is proposed by λ with a lower sequence number). The primary λ proposes v for ls for the first time in consensus instance li , and sets $ta[m] = \langle v, ls \rangle$. Note that $li \geq ls$ since li is incremented every time ls is.

Let d be the value of li at λ when v was proposed. If $\langle \text{VAL}, v, epoch, ls \rangle$ is decided in instance d , then $da[ls]$ is set to v . By induction, all values from n to $ls - 1$ are eventually included delivered, so they are added to da . When all elements of da with indexes in $[n + 1, ls - 1]$ are filled, the element $da[ls]$ is delivered too.

If a different VAL tuple is decided in instance d , then it is not from the λ , since λ proposes only one value per instance li . Therefore, $epoch_m \neq epoch, primary$ is true for λ , and $li \geq d$, so the leader proposes a $\langle VAL, v, epoch, ls \rangle$ tuple again, this time for a new consensus instance li . This resending can happen several time. Eventually, however, only λ proposes VAL tuples, so only its tuples are decided. \square

Lemma 12. *Algorithm 4 satisfies Delivery Liveness.*

Proof. Assume that a process delivers v after deciding it in consensus instance i . By the liveness property of consensus, every correct process eventually decides v in instance i . Since the decision on delivering v deterministically depends on the consistent order of decided values, all processes deliver v . \square

The following Theorem 15 derives directly from these lemmas and from Lemma 12.

Theorem 15 *Barrier-free primary order atomic broadcast is correct.*