

Replicated Data Types: Specification, Verification, Optimality

Sebastian Burckhardt
Microsoft Research

Alexey Gotsman
IMDEA Software Institute

Hongseok Yang
University of Oxford

Marek Zawirski
INRIA & UPMC-LIP6

Abstract

Geographically distributed systems often rely on replicated eventually consistent data stores to achieve availability and performance. To resolve conflicting updates at different replicas, researchers and practitioners have proposed specialized consistency protocols, called replicated data types, that implement objects such as registers, counters, sets or lists. Reasoning about replicated data types has however not been on par with comparable work on abstract data types and concurrent data types, lacking specifications, correctness proofs, and optimality results.

To fill in this gap, we propose a framework for specifying replicated data types using relations over events and verifying their implementations using replication-aware simulations. We apply it to 7 existing implementations of 4 data types with nontrivial conflict-resolution strategies and optimizations (last-writer-wins register, counter, multi-value register and observed-remove set). We also present a novel technique for obtaining lower bounds on the worst-case space overhead of data type implementations and use it to prove optimality of 4 implementations. Finally, we show how to specify consistency of replicated stores with multiple objects axiomatically, in analogy to prior work on weak memory models. Overall, our work provides foundational reasoning tools to support research on replicated eventually consistent stores.

Categories and Subject Descriptors D.2.4 [Software Engineering]: Software/Program Verification; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs

Keywords Replication; eventual consistency; weak memory

1. Introduction

To achieve availability and scalability, many networked computing systems rely on *replicated stores*, allowing multiple clients to issue operations on shared data on a number of *replicas*, which communicate changes to each other using message passing. For example, large-scale Internet services rely on *geo-replication*, which places data replicas in geographically distinct locations, and applications for mobile devices store replicas locally to support offline use. One benefit of such architectures is that the replicas remain locally available to clients even when network connections fail. Unfortunately, the famous CAP theorem [18] shows that such high Availability and tolerance to network Partitions are incompatible with *strong Consistency*, i.e., the illusion of a single centralized replica handling all operations. For this reason, modern replicated stores often

provide weaker forms of consistency, commonly dubbed *eventual consistency* [35]. ‘Eventual’ usually refers to the guarantee that

if clients stop issuing update requests, then the replicas (1) will eventually reach a consistent state.

Eventual consistency is a hot research area, and new replicated stores implementing it appear every year [1, 12, 15, 17, 22, 26, 32, 33, 36]. Unfortunately, their semantics is poorly understood: the very term eventual consistency is a catch-all buzzword, and different stores claiming to be eventually consistent actually provide subtly different guarantees. The property (1), which is a form of *quiescent consistency*, is too weak to capture these. Although it requires the replicas to converge to the same state eventually, it doesn’t say which one it will be. Furthermore, (1) does not provide any guarantees in realistic scenarios when updates never stop arriving. The difficulty of reasoning about the behavior of eventually consistent stores comes from a multitude of choices to be made in their design, some of which we now explain.

Allowing the replicas to be temporarily inconsistent enables eventually consistent stores to satisfy clients’ requests from the local replica immediately, and broadcast the changes to the other replicas only after the fact, when the network connection permits this. However, this means that clients can concurrently issue conflicting operations on the same data item at different replicas; furthermore, if the replicas are out-of-sync, these operations will be applied to its copies in different states. For example, two users sharing an online store account can write two different zip codes into the delivery address; the same users connected to replicas with different views of the shopping cart can also add and concurrently remove the same product. In such situations the store needs to ensure that, after the replicas exchange updates, the changes by different clients will be merged and all conflicts will be resolved in a meaningful way. Furthermore, to ensure eventual consistency (1), the conflict resolution has to be uniform across replicas, so that, in the end, they converge to the same state.

The protocols achieving this are commonly encapsulated within *replicated data types* [1, 10, 15, 17, 30, 32, 33] that implement *objects*, such as registers, counters, sets or lists, with various conflict-resolution strategies. The strategies can be as simple as establishing a total order on all operations using timestamps and letting the last writer win, but can also be much more subtle. Thus, a data type can detect the presence of a conflict and let the client deal with it: e.g., the *multi-value register* used in Amazon’s Dynamo key-value store [17] would return both conflicting zip codes in the above example. A data type can also resolve the conflict in an application-specific way. For example, the *observed-remove set* [7, 31] processes concurrent operations trying to add and remove the same element so that an add always wins, an outcome that may be appropriate for a shopping cart.

Replicated data type implementations are often nontrivial, since they have to maintain not only client-observable object state, but also *metadata* needed to detect and resolve conflicts and to handle network failures. This makes reasoning about their behavior challenging. The situation gets only worse if we consider multi-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

POPL ’14, January 22–24, 2014, San Diego, CA, USA.
Copyright © 2014 ACM 978-1-4503-2544-8/14/01...\$15.00.
<http://dx.doi.org/10.1145/2535838.2535848>

ple replicated objects: in this case, asynchronous propagation of updates between replicas may lead to counterintuitive behaviors—*anomalies*, in database terminology. The following code illustrates an anomaly happening in real replicated stores [1, 17]:

$$\text{Replica } r_1 \rightarrow x.\text{wr}(\text{post}) \quad \left\| \begin{array}{l} i = y.\text{rd} \text{ // comment} \leftarrow \text{Replica } r_2 \\ y.\text{wr}(\text{comment}) \quad \left\| \begin{array}{l} j = x.\text{rd} \text{ // empty} \end{array} \right. \end{array} \right. \quad (2)$$

We have two clients reading from and writing to register objects x and y at two different replicas; i and j are client-local variables. The first client makes a post by writing to x at replica r_1 and then comments on the post by writing to y . After every write, replica r_1 might send a message with the update to replica r_2 . If the messages carrying the writes of post to x and comment to y arrive to replica r_2 out of the order they were issued in, the second client can see the comment, but not the post. Different replicated stores may allow such an anomaly or not, and this has to be taken into account when reasoning about them.

In this paper, we propose techniques for reasoning about eventually consistent replicated stores in the following three areas.

1. Specification. We propose a comprehensive framework for specifying the semantics of replicated stores. Its key novel component is *replicated data type specifications* (§3), which provide the first way of specifying the semantics of replicated objects with advanced conflict resolution declaratively, like abstract data types [24]. We achieve this by defining the result of a data type operation not by a function of states, but of *operation contexts*—sets of events affecting the result of the operation, together with some relationships between them. We show that our specifications are sufficiently flexible to handle data types representing a variety of conflict-resolution strategies: last-write-wins register, counter, multi-value register and observed-remove set.

We then specify the semantics of a whole store with multiple objects, possibly of different types, by *consistency axioms* (§7), which constrain the way the store processes incoming requests in the style of weak shared-memory models [2] and thus define the anomalies allowed. As an illustration, we define consistency models used in existing replicated stores, including a weak form of eventual consistency [1, 17] and different kinds of causal consistency [22, 26, 32, 33]. We find that, when specialized to last-writer-wins registers, these specifications are very close to fragments of the C/C++ memory model [5]. Thus, our specification framework generalizes axiomatic shared-memory models to replicated stores with nontrivial conflict resolution.

2. Verification. We propose a method for proving the correctness of replicated data type implementations with respect to our specifications and apply it to seven existing implementations of the four data types mentioned above, including those with nontrivial optimizations. Reasoning about the implementations is difficult due to the highly concurrent nature of a replicated store, with multiple replicas simultaneously updating their object copies and exchanging messages. We address this challenge by proposing *replication-aware simulations* (§5). Like classical simulations from data refinement [20], these associate a concrete state of an implementation with its abstract description—structures on events, in our case. To combat the complexity of replication, they consider the state of an object at a single replica or a message in transit separately and associate it with abstract descriptions of only those events that led to it. Verifying an implementation then requires only reasoning about an instance of its code running at a single replica.

Here, however, we have to deal with another challenge: code at a single replica can access both the state of an object and a message at the same time, e.g., when updating the former upon receiving the latter. To reason about such code, we often need to rely on certain *agreement properties* correlating the abstract descriptions of the message and the object state. Establishing these properties re-

quires global reasoning. Fortunately, we find that agreement properties needed to prove realistic implementations depend only on basic facts about their messaging behavior and can thus be established once for broad classes of data types. Then a particular implementation within such a class can be verified by reasoning purely locally.

By carefully structuring reasoning in this way, we achieve easy and intuitive proofs of single data type implementations. We then lift these results to stores with multiple objects of different types by showing how consistency axioms can be proved given properties of the transport layer and data type implementations (§7).

3. Optimality. Replicated data type designers strive to optimize their implementations; knowing that one is optimal can help guide such efforts in the most promising direction. However, proving optimality is challengingly broad as it requires quantifying over all possible implementations satisfying the same specification.

For most data types we studied, the primary optimization target is the size of the metadata needed to resolve conflicts or handle network failures. To establish optimality of metadata size, we present a novel method for proving lower bounds on the *worst-case metadata overhead* of replicated data types—the proportion of metadata relative to the client-observable content. The main idea is to find a large family of executions of an arbitrary correct implementation such that, given the results of data type operations from a certain fixed point in any of the executions, we can recover the previous execution history. This implies that, across executions, the states at this point are distinct and thus must have some minimal size.

Using our method, we prove that four of the implementations we verified have an optimal worst-case metadata overhead among all implementations satisfying the same specification. Two of these (counter, last-writer-wins register) are well-known; one (optimized observed-remove set [6]) is a recently proposed nontrivial optimization; and one (optimized multi-value register) is a small improvement of a known implementation [32] that we discovered during a failed attempt to prove optimality of the latter. We summarize all the bounds we proved in Fig. 10.

We hope that the theoretical foundations we develop will help in exploring the design space of replicated data types and replicated eventually consistent stores in a systematic way.

2. Replicated Data Types


We now describe our formal model for replicated stores and introduce *replicated data type implementations*, which implement operations on a single object at a replica and the protocol used by replicas to exchange updates to this object. Our formalism follows closely the models used by replicated data type designers [32].

A replicated store is organized as a collection of named *objects* $\text{Obj} = \{x, y, z, \dots\}$. Each object is hosted at all *replicas* $r, s \in \text{ReplicaID}$. The sets of objects and replicas may be infinite, to model their dynamic creation. Clients interact with the store by performing *operations* on objects at a specified replica. Each object $x \in \text{Obj}$ has a *type* $\tau = \text{type}(x) \in \text{Type}$, whose *type signature* $(\text{Op}_\tau, \text{Val}_\tau)$ determines the set of supported operations Op_τ (ranged over by o) and the set of their return values Val_τ (ranged over by a, b, c, d). We assume that a special value $\perp \in \text{Val}_\tau$ belongs to all sets Val_τ and is used for operations that return no value. For example, we can define a counter data type ctr and an integer register type intreg with operations for reading, incrementing or writing an integer a : $\text{Val}_{\text{ctr}} = \text{Val}_{\text{intreg}} = \mathbb{Z} \cup \{\perp\}$, $\text{Op}_{\text{ctr}} = \{\text{rd}, \text{inc}\}$ and $\text{Op}_{\text{intreg}} = \{\text{rd}\} \cup \{\text{wr}(a) \mid a \in \mathbb{Z}\}$.

We also assume sets Message of messages (ranged over by m) and timestamps Timestamp (ranged over by t). For simplicity, we let timestamps be positive integers: $\text{Timestamp} = \mathbb{N}_1$.

DEFINITION 1. A *replicated data type implementation* for a data type τ is a tuple $\mathcal{D}_\tau = (\Sigma, \vec{\sigma}_0, M, \text{do}, \text{send}, \text{receive})$, where $\vec{\sigma}_0$:

Figure 1. Illustrations of a concrete (a) and two abstract executions (b, c)



$\text{ReplicaID} \rightarrow \Sigma$, $M \subseteq \text{Message}$ and

$$\begin{aligned} \text{do} &: \text{Op}_\tau \times \Sigma \times \text{Timestamp} \rightarrow \Sigma \times \text{Val}_\tau; \\ \text{send} &: \Sigma \rightarrow \Sigma \times M; \quad \text{receive} : \Sigma \times M \rightarrow \Sigma. \end{aligned}$$

We denote a component of \mathcal{D}_τ , such as do , by $\mathcal{D}_\tau.\text{do}$. A tuple \mathcal{D}_τ defines the class of implementations of objects with type τ , meant to be instantiated for every such object in the store. Σ is the set of states (ranged over by σ) used to represent the current state of the object, including metadata, at a single replica. The initial state at every replica is given by $\bar{\sigma}_0$.

\mathcal{D}_τ provides three methods that the rest of the store implementation can call at a given replica; we assume that these methods execute atomically. We visualize store executions resulting from repeated calls to the methods as in Fig. 1(a), by arranging the calls on several vertical timelines corresponding to replicas at which they occur and denoting the delivery of messages by diagonal arrows. In §4, we formalize them as sequences of transitions called *concrete executions* and define the store semantics by their sets; the intuition given by Fig. 1(a) should suffice for the following discussion.

A client request to perform an operation $o \in \text{Op}_\tau$ triggers the call $\text{do}(o, \sigma, t)$ (e.g., event 1 in Fig. 1(a)). This takes the current state $\sigma \in \Sigma$ of the object at the replica where the request is issued and a timestamp $t \in \text{Timestamp}$ provided by the rest of the store implementation and produces the updated object state and the return value of the operation. The data type implementation can use the timestamp provided, e.g., to implement the last-writer-wins conflict-resolution strategy mentioned in §1, but is free to ignore it.

Nondeterministically, in moments when the network is able to accept messages, a replica calls send . Given the current state of the object at the replica, send produces a message in M to broadcast to all other replicas (event 2 in Fig. 1(a)); sometimes send also alters the state of the object. Using broadcast rather than point-to-point communication does not limit generality, since we can always tag messages with the intended receiver. Another replica that receives the message generated by send calls receive to merge the enclosed update into its copy of the object state (event 3 in Fig. 1(a)).

We now reproduce three replicated data type implementations due to Shapiro et al. [32]. They fall into two categories: in *op-based* implementations, each message carries a description of the latest operations that the sender has performed, and in *state-based* implementations, a description of *all* operations it knows about.

Op-based counter (ctr). Fig. 2(a) shows an implementation of the ctr data type. A replica stores a pair $\langle a, d \rangle$, where a is the current value of the counter, and d is the number of increments performed since the last broadcast (we use angle brackets for tuples representing states and messages). The send method returns d and resets it; the receive method adds the content of the message to a . This implementation is correct, as long as each message is delivered exactly once (we show how to prove this in §5). Since inc operations commute, they never conflict: applying them in different orders at different replicas yields the same final state.

State-based counter (ctr). The implementation in Fig. 2(b) summarizes the currently known history by recording the contri-

Figure 2. Three replicated data type implementations

(a) Op-based counter (ctr)	
$\Sigma = \mathbb{N}_0 \times \mathbb{N}_0$	$\text{do}(\text{rd}, \langle a, d \rangle, t) = (\langle a, d \rangle, a)$
$M = \mathbb{N}_0$	$\text{do}(\text{inc}, \langle a, d \rangle, t) = (\langle a + 1, d + 1 \rangle, \perp)$
$\bar{\sigma}_0 = \lambda r. \langle 0, 0 \rangle$	$\text{send}(\langle a, d \rangle) = (\langle a, 0 \rangle, d)$
	$\text{receive}(\langle a, d \rangle, d') = \langle a + d', d \rangle$
(b) State-based counter (ctr)	
Σ	$= \text{ReplicaID} \times (\text{ReplicaID} \rightarrow \mathbb{N}_0)$
σ_0	$= \lambda r. \langle r, \lambda s. 0 \rangle$
M	$= \text{ReplicaID} \rightarrow \mathbb{N}_0$
$\text{do}(\text{rd}, \langle r, v \rangle, t)$	$= (\langle r, v \rangle, \sum \{v(s) \mid s \in \text{ReplicaID}\})$
$\text{do}(\text{inc}, \langle r, v \rangle, t)$	$= (\langle r, v[r \mapsto v(r) + 1] \rangle, \perp)$
$\text{send}(\langle r, v \rangle)$	$= (\langle r, v \rangle, v)$
$\text{receive}(\langle r, v \rangle, v')$	$= \langle r, (\lambda s. \max\{v(s), v'(s)\}) \rangle$
(c) State-based last-writer-wins register (intreg)	
Σ	$= \mathbb{Z} \times (\text{Timestamp} \uplus \{0\})$
$\bar{\sigma}_0$	$= \lambda r. \langle 0, 0 \rangle$
M	$= \Sigma$
$\text{do}(\text{rd}, \langle a, t \rangle, t')$	$= (\langle a, t \rangle, a)$
$\text{do}(\text{wr}(a'), \langle a, t \rangle, t')$	$= \text{if } t < t' \text{ then } (\langle a', t' \rangle, \perp) \text{ else } (\langle a, t \rangle, \perp)$
$\text{send}(\langle a, t \rangle)$	$= (\langle a, t \rangle, \langle a, t \rangle)$
$\text{receive}(\langle a, t \rangle, \langle a', t' \rangle)$	$= \text{if } t < t' \text{ then } \langle a', t' \rangle \text{ else } \langle a, t \rangle$

bution of every replica to the counter value separately (reminiscent of vector clocks [28]). A replica stores its identifier r and a vector v such that for each replica s the entry $v(s)$ gives the number of increments made by clients at s that have been received by r . A rd operation returns the sum of all entries in the vector. An inc operation increments the entry for the current replica. We denote by $v[i \mapsto j]$ the function that has the same value as v everywhere, except for i , where it has the value j . The send method returns the vector, and the receive method takes the maximum of each entry in the vectors v and v' given to it. This is correct because an entry for s in either vector reflects a prefix of the sequence of increments done at replica s . Hence, we know that $\min\{v(s), v'(s)\}$ increments by s are taken into account both in $v(s)$ and in $v'(s)$.

State-based last-writer-wins (LWW) register (intreg). Unlike counters, registers have update operations that are not commutative. To resolve conflicts, the implementation in Fig. 2 uses the last-writer-wins strategy, creating a total order on writes by associating a unique timestamp with each of them. A state contains the current value, returned by rd , and the timestamp at which it was written (initially, we have 0 instead of a timestamp). A $\text{wr}(a')$ compares its timestamp t' with the timestamp t of the current value a and sets the value to the one with the highest timestamp. Note that here we have to allow for $t' < t$, since we do not make any assumptions about timestamps apart from uniqueness: e.g., the rest of the store implementation can compute them using physical or Lamport clocks [21]. We show how to state assumptions about timestamps in §4. The send method just returns the state, and the receive method chooses the winning value by comparing the timestamps in the current state and the message, like wr .

State-based vs. op-based. State-based implementations converge to a consistent state faster than op-based implementations because they are *transitively delivering*, meaning that they can propagate updates indirectly. For example, when using the counter in Fig. 2(b), in the execution in Fig. 1(a) the read at r_3 (event 7) returns 2, even though the message from r_1 has not arrived yet, because r_3 learns about r_1 's update via r_2 . State-based implementations are also resilient against transport failures like message *loss*, *reordering*, or *duplication*. Op-based implementations require the replicated store using them to mask such failures (e.g., using message sequence numbers, retransmission buffers, or reorder buffers).

The potential weakness of state-based implementations is the size of states and messages, which motivates our examination of space optimality in §6. For example, we show that the counter in Fig. 2(b) is optimal, meaning that no counter implementation satisfying the same requirements (transitive delivery and resilience against message loss, reordering, and duplication) can do better.

3. Specifying Replicated Data Types and Stores

Consider the concrete execution in Fig. 1(a). What are valid return values for the read in event 7? Intuitively, 1 or 2 can be justifiable, but not 100. We now present a framework for specifying the expected outcome declaratively, without referring to implementation details. For example, we give a specification of a replicated counter that is satisfied by both implementations in Fig. 2(a, b).

In presenting the framework, we rely on the intuitive understanding of the way a replicated store executes given in §2. Later we define the store semantics formally (§4), which lets us state what it means for a store to satisfy our specifications (§4 and §7).

3.1 Abstract Executions and Specification Structure

We define our specifications on *abstract executions*, which include only user-visible events (corresponding to do calls) and describe the other information about the store processing in an implementation-independent form. Informally, we consider a concrete execution correct if it can be justified by an abstract execution satisfying the specifications that is “similar” to it and, in particular, has the same operations and return values.

Abstract executions are inspired by axiomatic definitions of weak shared-memory models [2]. In particular, we use their previously proposed reformulation with visibility and arbitration relations [12], which are similar to the reads-from and coherence relations from weak shared-memory models. We provide a comparison with shared-memory models in §7 and with [12] in §8.

DEFINITION 2. An *abstract execution* is a tuple $A = (E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar})$, where

- $E \subseteq \text{Event}$ is a set of events from a countable universe Event ;
- each event $e \in E$ describes a replica $\text{repl}(e) \in \text{ReplicaID}$ performing an operation $\text{oper}(e) \in \text{Op}_{\text{type}(\text{obj}(e))}$ on an object $\text{obj}(e) \in \text{Obj}$, which returns the value $\text{rval}(e) \in \text{Val}_{\text{type}(\text{obj}(e))}$;
- $\text{ro} \subseteq E \times E$ is a **replica order**, which is a union of transitive, irreflexive and total orders on events at each replica;
- $\text{vis} \subseteq E \times E$ is an **acyclic visibility relation** such that $\forall e, f \in E. e \xrightarrow{\text{vis}} f \implies \text{obj}(e) = \text{obj}(f)$;
- $\text{ar} \subseteq E \times E$ is an **arbitration relation**, which is a union of transitive, irreflexive and total orders on events on each object.

We also require that ro , vis and ar be well-founded.

In the following, we denote components of A and similar structures as in $A.\text{repl}$. We also use $(e, f) \in r$ and $e \xrightarrow{r} f$ interchangeably.

Informally, $e \xrightarrow{\text{vis}} f$ means that f is aware of e and thus e 's effect can influence f 's return value. In implementation terms, this may be the case if the update performed by e has been delivered to the replica performing f before f is issued. The exact meaning of “delivered”, however, depends on how much information messages carry in the implementation. For example, as we explain in §3.2, the return value of a read from a counter is equal to the number of inc operations visible to it. Then, as we formalize in §4, the abstract execution illustrated in Fig. 1(b) justifies the op-based implementation in Fig. 2(a) reading 1 in the concrete execution in Fig. 1(a). The abstract execution in Fig. 1(c) justifies the state-based implementation in Fig. 2(b) reading 2 due to transitive delivery (§2). There is no abstract execution that would justify reading 100.

The ar relation represents the ordering information provided by the store, e.g., via timestamps. On the right we show an abstract execution corresponding to a variant of the anomaly (2). The ar edge means that any replica that sees both writes to x should assume that post overwrites empty .

We give a store specification by two components, constraining abstract executions:

1. **Replicated data type specifications** determine return values of operations in an abstract execution in terms of its vis and ar relations, and thus define conflict-resolution policies for individual objects in the store. The specifications are the key novel component of our framework, and we discuss them next.
2. **Consistency axioms** constrain vis and ar and thereby disallow anomalies and extend the semantics of individual objects to that of the entire store. We defer their discussion to §7. See Fig. 13 for their flavor; in particular, COCV prohibits the anomaly above.

Each of these components can be varied separately, and our specifications will define the semantics of any possible combination. Given a specification of a store, we can determine whether a set of events can be observed by its users by checking if there is an abstract execution with this set of events satisfying the data type specifications and consistency axioms.

3.2 Replicated Data Type Specifications

In a sequential setting, the semantics of a data type τ can be specified by a function $\mathcal{S}_\tau : \text{Op}_\tau^+ \rightarrow \text{Val}_\tau$, which, given a non-empty sequence of operations performed on an object, specifies the return value of the last operation. For a register, read operations return the value of the last preceding write, or zero if there is no prior write. For a counter, read operations return the number of preceding increments. Thus, for any sequence of operations ξ :

$$\begin{aligned} \mathcal{S}_{\text{intreg}}(\xi \text{ rd}) &= a, \text{ if } \text{wr}(0) \xi = \xi_1 \text{ wr}(a) \xi_2 \text{ and} \\ &\quad \xi_2 \text{ does not contain wr operations;} \\ \mathcal{S}_{\text{ctr}}(\xi \text{ rd}) &= (\text{the number of inc operations in } \xi); \\ \mathcal{S}_{\text{intreg}}(\xi \text{ inc}) &= \mathcal{S}_{\text{ctr}}(\xi \text{ wr}(a)) = \perp. \end{aligned}$$

In a replicated store, the story is more interesting. We specify a data type τ by a function \mathcal{F}_τ , generalizing \mathcal{S}_τ . Just like \mathcal{S}_τ , this determines the return value of an operation based on prior operations performed on the object. However, \mathcal{F}_τ takes as a parameter not a sequence, but an *operation context*, which includes all we need to know about a store execution to determine the return value of a given operation o —the set E of all events that are visible to o , together with the operations performed by the events and visibility and arbitration relations on them.

DEFINITION 3. An *operation context* for a data type τ is a tuple $L = (o, E, \text{oper}, \text{vis}, \text{ar})$, where $o \in \text{Op}_\tau$, E is a finite subset of Event , $\text{oper} : E \rightarrow \text{Op}_\tau$, $\text{vis} \subseteq E \times E$ is acyclic and $\text{ar} \subseteq E \times E$ is transitive, irreflexive and total.

We can extract the context of an event $e \in A.E$ in an abstract execution A by selecting all events visible to it according to $A.\text{vis}$:

$$\text{ctx}(A, e) = (A.\text{oper}(e), G, (A.\text{oper})|_G, (A.\text{vis})|_G, (A.\text{ar})|_G),$$

where $G = (A.\text{vis})^{-1}(e)$ and $\cdot|_G$ is the restriction to events in G . Thus, in the abstract execution in Fig. 1(b), the operation context of the read from x includes only one increment event; in the execution in Fig. 1(c) it includes two.

DEFINITION 4. A **replicated data type specification** for a type τ is a function \mathcal{F}_τ that, given an operation context L for τ , specifies a return value $\mathcal{F}_\tau(L) \in \text{Val}_\tau$.

Note that $\mathcal{F}_\tau(o, \emptyset, \dots)$ returns the value resulting from performing o on the initial state for the data type (e.g., 0 for the LWW-register).

We specify multiple data types used in a replicated store by a partial function \mathbb{F} mapping them to data type specifications.

DEFINITION 5. An abstract execution A *satisfies* \mathbb{F} , written $A \models \mathbb{F}$, if the return value of every event in A is computed on its context by the specification for the type of the object the event accesses:

$$\forall e \in A.E. (A.\text{rval}(e) = \mathbb{F}(\text{type}(A.\text{obj}(e)))(\text{ctxt}(A, e))).$$

We specify a whole store by \mathbb{F} and a set of consistency axioms (§7). This lets us determine if its users can observe a given set of events by checking if there is an abstract execution with these events that satisfies \mathbb{F} according to the above definition, as well as the axioms.

Note that \mathcal{F}_τ is deterministic. This does not mean that so is an outcome of an operation on a store; rather, that all the non-determinism arising due to its distributed nature is resolved by vis and ar in the context passed to \mathcal{F}_τ . These relations are chosen arbitrarily subject to consistency axioms. Due to the determinacy property, two events that perform the same operation and see the same set of events produce the same return values. As we show in §7, this property ensures that our specifications can formalize eventual consistency in the sense of (1).

We now give four examples of data type specifications, corresponding to the four conflict-resolution strategies mentioned in §1 and §2: (1) operations commute, so no conflicts arise; (2) last writer wins; (3) all conflicting values are returned; and (4) conflicts are resolved in an application-specific way. We start by specifying the data types whose implementations we presented in §2.

1. Counter (ctr) is defined by

$$\begin{aligned} \mathcal{F}_{\text{ctr}}(\text{inc}, E, \text{oper}, \text{vis}, \text{ar}) &= \perp; \\ \mathcal{F}_{\text{ctr}}(\text{rd}, E, \text{oper}, \text{vis}, \text{ar}) &= \{e \in E \mid \text{oper}(e) = \text{inc}\}. \end{aligned} \quad (3)$$

Thus, according to Def. 5 the executions in Fig. 1(b) and 1(c) satisfy the counter specification: both 1 and 2 are valid return values for the read from x when there are two concurrent increments.

2. LWW-register (lwwreg) is defined by

$$\mathcal{F}_{\text{lwwreg}}(o, E, \text{oper}, \text{vis}, \text{ar}) = \mathcal{S}_{\text{lwwreg}}(E^{\text{ar}} o), \quad (4)$$

where E^{ar} denotes the sequence obtained by ordering the operations performed by the events in E according to ar . Thus, the return value is determined by establishing a total order of the visible operations and applying the regular sequential semantics. For example, by Def. 5 in the example execution from §3.1 the read from x has to return *empty*; if we had a vis edge from the write of *post* to the read from x , then the read would have to return *post*. As we show in §7, weak shared-memory models are obtained by specializing our framework to stores with only LWW-registers.

We can obtain a concurrent semantics \mathcal{F}_τ of any data type τ based on its sequential semantics \mathcal{S}_τ similarly to (4). For example, \mathcal{F}_{ctr} defined above is equivalent to what we obtain using this generic construction. The next two examples go beyond this.

3. Multi-value register (mvr). This register [1, 17] has the same operations as the LWW-register, but its reads return a set of values:

$$\begin{aligned} \mathcal{F}_{\text{mvr}}(\text{rd}, E, \text{oper}, \text{vis}, \text{ar}) &= \{a \mid \exists e \in E. \text{oper}(e) = \text{wr}(a) \\ &\wedge \neg \exists f \in E. \text{oper}(f) = \text{wr}(-) \wedge e \xrightarrow{\text{vis}} f\}. \end{aligned}$$

(We write $_$ for an expression whose value is irrelevant.) A read returns the values written by currently conflicting writes, defined as those that are not superseded in vis by later writes; ar is not used. For example, a rd would return $\{2, 3\}$ in the context on the right.




Figure 3. The set of configurations Config and the transition relation $\longrightarrow_{\mathbb{D}}$: $\text{Config} \times \text{Event} \times \text{Config}$ for a data type library \mathbb{D} . We use $e : \{h_1 = u_1, h_2 = u_2\}$ to abbreviate $h_1(e) = u_1$ and $h_2(e) = u_2$. We uncurry $R \in \text{RState}$ where convenient.

$$\begin{aligned} \text{Obj}_\tau &= \{x \in \text{Obj} \mid \text{type}(x) = \tau\} \\ \text{RState} &= \bigcup_{X \subseteq \text{Obj}} \prod_{x \in X} (\text{ReplicaID} \rightarrow \mathbb{D}(\text{type}(x)).\Sigma) \\ \text{TState} &= \text{MessageID} \rightarrow \bigcup_{\tau \in \text{Type}} (\text{ReplicaID} \times \text{Obj}_\tau \times \mathbb{D}(\tau).M) \\ \text{Config} &= \text{RState} \times \text{TState} \\ \\ \mathbb{D}(\text{type}(x)).\text{do}(o, \sigma, t) &= (\sigma', a) \\ e : \{\text{act} = \text{do}, \text{obj} = x, \text{repl} = r, \text{oper} = o, \text{time} = t, \text{rval} = a\} \\ \hline (R[(x, r) \mapsto \sigma], T) &\xrightarrow{e}_{\mathbb{D}} (R[(x, r) \mapsto \sigma'], T) \\ \\ \mathbb{D}(\text{type}(x)).\text{send}(\sigma) &= (\sigma', m) \quad \text{mid} \notin \text{dom}(T) \\ e : \{\text{act} = \text{send}, \text{obj} = x, \text{repl} = r, \text{msg} = \text{mid}\} \\ \hline (R[(x, r) \mapsto \sigma], T) &\xrightarrow{e}_{\mathbb{D}} (R[(x, r) \mapsto \sigma'], T[\text{mid} \mapsto (r, x, m)]) \\ \\ \mathbb{D}(\text{type}(x)).\text{receive}(\sigma, m) &= \sigma' \quad r \neq r' \\ e : \{\text{act} = \text{receive}, \text{obj} = x, \text{repl} = r, \text{srepl} = r', \text{msg} = \text{mid}\} \\ \hline (R[(x, r) \mapsto \sigma], T[\text{mid} \mapsto (r', x, m)]) &\xrightarrow{e}_{\mathbb{D}} (R[(x, r) \mapsto \sigma'], T[\text{mid} \mapsto (r', x, m)]) \end{aligned}$$

4. Observed-remove set (orset). How do we specify a replicated set of integers? The operations of adding and removing different elements commute and thus do not conflict. Conflicts arise from concurrently adding and removing the same element. For example, we need to decide what rd will return as the contents of the set in the context $(\text{rd}, \{e, f\}, \text{oper}, \text{vis}, \text{ar})$, where $\text{oper}(e) = \text{add}(42)$ and $\text{oper}(f) = \text{remove}(42)$. If we use the generic construction from the LWW-register, the result will depend on the arbitration relation: \emptyset if $e \xrightarrow{\text{ar}} f$, and $\{42\}$ otherwise. An application may require a more consistent behavior, e.g., that an add operation always win against concurrent remove operations. Observed-remove (OR) set [7, 31] achieves this by mandating that remove operations cancel only the add operations that are visible to them:

$$\begin{aligned} \mathcal{F}_{\text{orset}}(\text{rd}, E, \text{oper}, \text{vis}, \text{ar}) &= \{a \mid \exists e \in E. \text{oper}(e) = \text{add}(a) \\ &\wedge \neg \exists f \in E. \text{oper}(f) = \text{remove}(a) \wedge e \xrightarrow{\text{vis}} f\}, \end{aligned} \quad (5)$$

In the above operation context rd will return \emptyset if $e \xrightarrow{\text{vis}} f$, and $\{42\}$ otherwise. The rationale is that, in the former case, $\text{add}(42)$ and $\text{remove}(42)$ are not concurrent: the user who issued the $\text{remove}(42)$ knew that 42 was in the set and thus meant to remove it. In the latter case, the two operations are concurrent and thus add wins.

As the above examples illustrate, our specifications can describe the semantics of data types and their conflict-resolution policies declaratively, without referring to the internals of their implementations. In this sense the specifications generalize the concept of an *abstract data type* [24] to the replicated setting.

4. Store Semantics and Data Type Correctness

A *data type library* \mathbb{D} is a partial mapping from types τ to data type implementations $\mathbb{D}(\tau)$ from Def. 1. We now define the semantics of a replicated store with a data type library \mathbb{D} as a set of its *concrete executions*, previously introduced informally by Fig. 1(a). We then state what it means for data type implementations of §2 to satisfy their specifications of §3.2 by requiring their concrete executions to be justified by abstract ones. In §7 we generalize this to the correctness of the whole store with multiple object with respect to both data type specifications and consistency axioms.

Semantics. We define the semantics using the relation $\longrightarrow_{\mathbb{D}}$: $\text{Config} \times \text{Event} \times \text{Config}$ in Fig. 3, which describes a single

step of the store execution. The relation transforms **configurations** $(R, T) \in \text{Config}$ describing the store state: R gives the object state at each replica, and T the set of messages in transit between them, each identified by a message identifier $mid \in \text{MessageID}$. A message is annotated by the origin replica and the object to which it pertains. We allow the store to contain only some objects from Obj and thus allow R to be partial on them. We use a number of functions on events, such as act , obj , etc., to record the information about the corresponding transitions, so that $\rightarrow_{\mathbb{D}}$ is implicitly parameterized by them; we give their full list in Def. 6 below.

The first rule in Fig. 3 describes a replica r performing an operation o on an object x using the do method of the corresponding data type implementation. We record the return value using the function rval . To communicate the change to other replicas, we can at any time perform a transition defined by the second rule, which puts a new message m created by a call to send into the set of messages in transit. The third rule describes the delivery of such a message to a replica r other than the origin replica r' , which triggers a call to receive . Note that the relation $\rightarrow_{\mathbb{D}}$ does not make any assumptions about message delivery: messages can be delivered in any order, multiple times, or not at all. These assumptions can be introduced separately, as we show later in this section. A concrete execution can be thought of as a finite or infinite sequence of transitions:

$$(R_0, T_0) \xrightarrow{e_1} (R_1, T_1) \xrightarrow{e_2} \dots \xrightarrow{e_n} (R_n, T_n) \dots,$$

where all events e_i are distinct. To ease mapping between concrete and abstract executions in the future, we formalize it as a structure on events, similarly to Def. 2.

DEFINITION 6. A **concrete execution** of a store with a data type library \mathbb{D} is a tuple

$$C = (E, \text{eo}, \text{pre}, \text{post}, \text{act}, \text{obj}, \text{repl}, \text{oper}, \text{time}, \text{rval}, \text{msg}, \text{srepl}).$$

Here $E \subseteq \text{Event}$, the **execution order** eo is a well-founded, transitive, irreflexive and total order on E , relating the events according to the order of the transitions they describe, time is injective and $\text{pre}, \text{post} : E \rightarrow \text{Config}$ form a valid sequence of transitions:

$$\begin{aligned} & (\forall e \in E. \text{pre}(e) \xrightarrow{e}_{\mathbb{D}} \text{post}(e)) \wedge \\ & (\forall e, f \in E. e \xrightarrow{\text{eo}} f \wedge \neg \exists g. e \xrightarrow{\text{eo}} g \xrightarrow{\text{eo}} f \implies \text{post}(e) = \text{pre}(f)). \end{aligned}$$

We have omitted the types of functions on events, which are easily inferred from Fig. 3: e.g., $\text{act} : E \rightarrow \{\text{do}, \text{send}, \text{receive}\}$ and $\text{time} : E \rightarrow \text{Timestamp}$, defined only on e with $\text{act}(e) = \text{do}$.

We denote the initial configuration of C by $\text{init}(C) = C.\text{pre}(e_0)$, where e_0 is the minimal event in $C.\text{eo}$. If $C.E$ is finite, we denote the final configuration of C by $\text{final}(C) = C.\text{post}(e_f)$, where e_f is the maximal event in $C.\text{eo}$. The **semantics** $\llbracket \mathbb{D} \rrbracket$ of \mathbb{D} is the set of all its concrete executions C that start in a configuration with an empty set of messages and all objects in initial states, i.e.,

$$\exists X \subseteq \text{Obj}. \text{init}(C) = ((\lambda x \in X. \mathbb{D}(\text{type}(x)).\vec{\sigma}_0), []),$$

where $[]$ is the everywhere-undefined function.

Transport layer specifications. Data type implementations such as the op-based counter in Fig. 2(a) can rely on some guarantees concerning the delivery of messages ensured by the rest of the store implementation. They may similarly assume certain properties of timestamps other than uniqueness (guaranteed by the injectivity of time in Def. 6). We take such assumptions into account by admitting only a subset of executions from $\llbracket \mathbb{D} \rrbracket$ that satisfy a **transport layer specification** \mathcal{T} , which is a predicate on concrete executions. Thus, we consider a **replicated store** to be defined by a pair $(\mathbb{D}, \mathcal{T})$ and the set of its executions be $\llbracket \mathbb{D} \rrbracket \cap \mathcal{T}$.

Even though our definition of \mathcal{T} lets it potentially restrict data type implementation internals, the particular instantiations we use

only restrict message delivery and timestamps. For technical reasons, we assume that \mathcal{T} always satisfies certain closure properties: for every $C \in \mathcal{T}$, the projection of C onto events on a given object or a subset of events forming a prefix in the eo order is also in \mathcal{T} .

As an example, we define a transport layer specification ensuring that a message is delivered to any single replica at most once, as required by the implementation in Fig. 2(a). Let the **delivery relation** $\text{del}(C) \in C.E \times C.E$ pair events sending and receiving the same message:

$$\begin{aligned} e \xrightarrow{\text{del}(C)} f & \iff e \xrightarrow{C.\text{eo}} f \wedge C.\text{act}(e) = \text{send} \wedge \\ & C.\text{act}(f) = \text{receive} \wedge C.\text{msg}(e) = C.\text{msg}(f). \end{aligned}$$

Then the desired condition on concrete executions C is

$$\begin{aligned} \forall e, f, g \in C.E. e \xrightarrow{\text{del}(C)} f \wedge e \xrightarrow{\text{del}(C)} g \wedge \\ C.\text{repl}(f) = C.\text{repl}(g) \implies f = g. \quad (\text{T-Unique}) \end{aligned}$$

Data type implementation correctness. We now state what it means for an implementation \mathcal{D}_τ of type τ from Def. 1 to satisfy a specification \mathcal{F}_τ from Def. 4. To this end, we consider the behavior of \mathcal{D}_τ under the “most general” client and transport layer, performing all possible operations and message deliveries. Formally, let $\llbracket \mathcal{D}_\tau \rrbracket$ be the set of executions $C \in \llbracket [\tau \mapsto \mathcal{D}_\tau] \rrbracket$ of a store containing a single object x of a type τ with the implementation \mathcal{D}_τ , i.e., $\text{init}(C) = (R, [])$ for some R such that $\text{dom}(R) = \{x\}$.

Then \mathcal{D}_τ should satisfy \mathcal{F}_τ under a transport specification \mathcal{T} if for every concrete execution $C \in \llbracket \mathcal{D}_\tau \rrbracket \cap \mathcal{T}$ we can find a “similar” abstract execution satisfying \mathcal{F}_τ and, in particular, having the same operations and return values. As it happens, all components of the abstract execution except visibility are straightforwardly determined by C ; as explained in §3.1, we have some freedom in choosing visibility. We define the choice using a **visibility witness** \mathcal{V} , which maps a concrete execution $C \in \llbracket \mathcal{D}_\tau \rrbracket$ to an acyclic relation on $(C.E)_{|\text{do}}$ defining visibility (here $\cdot_{|\text{do}}$ is the restriction to events e with $C.\text{act}(e) = \text{do}$). Let

$$\begin{aligned} e \xrightarrow{\text{ro}(C)} f & \iff e \xrightarrow{C.\text{eo}} f \wedge C.\text{repl}(e) = C.\text{repl}(f); \\ e \xrightarrow{\text{ar}(C)} f & \iff e, f \in (C.E)_{|\text{do}} \wedge \\ & C.\text{obj}(e) = C.\text{obj}(f) \wedge C.\text{time}(e) < C.\text{time}(f). \end{aligned}$$

Then the abstract execution justifying $C \in \llbracket \mathcal{D}_\tau \rrbracket$ is defined by

$$\text{abs}(C, \mathcal{V}) = (C.E_{|\text{do}}, E.\text{repl}_{|\text{do}}, E.\text{obj}_{|\text{do}}, E.\text{oper}_{|\text{do}}, E.\text{rval}_{|\text{do}}, \text{ro}(C)_{|\text{do}}, \mathcal{V}(C), \text{ar}(C)).$$

DEFINITION 7. A data type implementation \mathcal{D}_τ **satisfies** a specification \mathcal{F}_τ with respect to \mathcal{V} and \mathcal{T} , written $\mathcal{D}_\tau \text{ sat}[\mathcal{V}, \mathcal{T}] \mathcal{F}_\tau$, if $\forall C \in \llbracket \mathcal{D}_\tau \rrbracket \cap \mathcal{T}. (\text{abs}(C, \mathcal{V}) \models [\tau \mapsto \mathcal{F}_\tau])$, where \models is defined in Def. 5.

As we explained informally in §3.1, the visibility witness depends on how much information the implementation puts into messages. Since state-based implementations, such as the ones in Fig. 2(b, c), are transitively delivering (§2), for them we use the witness $\mathcal{V}^{\text{state}}(C) = (\text{ro}(C) \cup \text{del}(C))^+_{|\text{do}}$. By the definition of $\text{ro}(C)$ and $\text{del}(C)$, $(\text{ro}(C) \cup \text{del}(C))$ is acyclic, so $\mathcal{V}^{\text{state}}$ is well-defined. State-based implementations do not make any assumptions about the transport layer: in this case we write $\mathcal{T} = \text{T-Any}$. In contrast, op-based implementations, such as the one in Fig. 2(a), require $\mathcal{T} = \text{T-Unique}$. Since such implementations are not transitively delivering, the witness $\mathcal{V}^{\text{state}}$ is not appropriate for them. We could attempt to define a witness for them by straightforwardly lifting the delivery relation:

$$\begin{aligned} \mathcal{V}^{\text{op}}(C) & = \text{ro}(C)_{|\text{do}} \cup \{(e, f) \mid e, f \in (C.E)_{|\text{do}} \wedge \\ & \exists e', f'. e \xrightarrow{\text{ro}(C)} e' \xrightarrow{\text{del}(C)} f' \xrightarrow{\text{ro}(C)} f\}. \end{aligned}$$

However, we need to be more careful, since for op-based implementations $e \xrightarrow{\text{ro}(C)} e' \xrightarrow{\text{del}(C)} f' \xrightarrow{\text{ro}(C)} f$ does not ensure that the update of e is taken into account by f : if there is another send event e'' in between e and e' , then e'' will capture the update of e and e' will not. Hence, we define the witness as:

$$\begin{aligned} \mathcal{V}^{\text{op}}(C) &= \text{ro}(C)|_{\text{do}} \cup \{(e, f) \mid e, f \in (C.E)|_{\text{do}}\} \\ &\wedge \exists e', f'. e \xrightarrow{\text{ro}(C)} e' \xrightarrow{\text{del}(C)} f' \xrightarrow{\text{ro}(C)} f \\ &\wedge \neg \exists e''. e \xrightarrow{\text{ro}(C)} e'' \xrightarrow{\text{ro}(C)} e' \wedge C.\text{act}(e'') = \text{send}. \end{aligned}$$

We next present a method for proving data type implementation correctness in the sense of Def. 7. In §7 we lift this to stores with multiple objects and take into account consistency axioms.


5. Proving Data Type Implementations Correct

The straightforward approach to proving correctness in the sense of Def. 7 would require us to consider global store configurations in executions C , including object states at all replicas and all messages in transit, making the reasoning non-modular and unintuitive. To deal with this challenge, we focus on a single component of a store configuration using *replication-aware simulation relations* \mathcal{R}_r and \mathcal{M} , analogous to simulation (aka coupling) relations used in data refinement [20]. The \mathcal{R}_r relation associates the object state at a replica r with an abstract execution that describes *only those events that led to this state*; \mathcal{M} does the same for a message. For example, when proving \mathcal{D}_{ctr} in Fig. 2(b) with respect to \mathcal{F}_{ctr} in (3), \mathcal{M} associates a message carrying a vector v with executions in which each replica s makes $v(s)$ increments. As part of a proof of \mathcal{D}_τ , we require checking that the effect of its methods, such as $\mathcal{D}_\tau.\text{do}$, can be simulated by appropriately transforming related abstract executions while preserving the relations. We define these transformations using *abstract methods* do^\sharp , send^\sharp and receive^\sharp as illustrated in Fig. 4(a, b). For example, if a replica r executes $\mathcal{D}_\tau.\text{do}$ from a state σ related by \mathcal{R}_r to an abstract execution I (we explain the use of I instead of A later), we need to find an I' related by \mathcal{R}_r to the resulting state σ' . We also need to check that the value returned by $\mathcal{D}_\tau.\text{do}$ on σ is equal to that returned by \mathcal{F}_τ on I .

These conditions consider the behavior of an implementation method on a single state and/or message and its effect on only the relevant part of the abstract execution. However, by localizing the reasoning in this way, we lose some global information that is actually required to verify realistic implementations. In particular, this occurs when discharging the obligation for receive in Fig. 4(b). Taking a global view, σ and m there are meant to come from the same configuration in a concrete execution C ; correspondingly, I and J are meant to be fragments of the same abstract execution $\text{abs}(C, \mathcal{V})$. In this context we may know certain *agreement properties* correlating I and J , e.g., that the union of their visibility relations is itself a well-formed visibility relation and is thus acyclic. Establishing them requires global reasoning about whole executions C and $\text{abs}(C, \mathcal{V})$. Fortunately, we find that this can be done knowing only the abstract methods, not the implementation \mathcal{D}_τ . Furthermore, these methods state basic facts about the messaging behavior of implementations and are thus common to broad classes of them, such as state-based or op-based. This allows us to establish agreement properties using global reasoning once for a given class of implementations; at this stage we can also benefit from the transport layer specification \mathcal{T} and check that the abstract methods construct visibility according to the given witness \mathcal{V} . Then a particular implementation within the class can be verified by discharging local obligations, such as those in Fig. 4(a, b), while assuming agreement properties. This yields easy and intuitive proofs.

To summarize, we deal with the challenge posed by a distributed data type implementation by decomposing reasoning about it into

Figure 4. Diagrams illustrating replication-aware simulations



global reasoning done once for a broad class of implementations and local implementation-specific reasoning. We start by presenting the general form of obligations to be discharged for a single implementation within a certain class (§5.1) and the particular form they take for the class of state-based implementations (§5.2), together with some examples (§5.3). We then formulate the obligations to be discharged for a class of implementations (§5.4), which in particular, establish the agreement properties assumed in the per-implementation obligations. In §B, we give the obligations for op-based implementations, together with a proof of the counter in Fig. 2(a). An impatient reader can move on to §6 after finishing §5.3, and come back to §5.4 later.

Since Def. 7 considers only single-object executions, we fix an object x of type τ and consider only concrete and abstract executions over x , whose sets we denote by $\text{CEx}[x]$ and $\text{AEx}[x]$.

5.1 Replication-Aware Simulations

As is typical for simulation-based proofs, we need to use auxiliary state to record information about the computation history. For this reason, actually our simulation relations associate a state or a message with an *instrumented execution*—a pair $(A, \text{info}) \in \text{IEx}$ of an abstract execution $A \in \text{AEx}[x]$ and a function $\text{info} : A.E \rightarrow \text{Alnfo}$, tagging events with auxiliary information from a set Alnfo . As we show below, Alnfo can be chosen once for a class of data type implementations: e.g., $\text{Alnfo} = \text{Timestamp}$ for state-based ones (§5.2). We use I and J to range over instrumented executions and shorten, e.g., $I.A.E$ to $I.E$. For a partial function h we write $h(x)\downarrow$ for $x \in \text{dom}(h)$, and adopt the convention that $h(x) = y$ implies $h(x)\downarrow$.

DEFINITION 8. A *replication-aware simulation* between \mathcal{D}_τ and \mathcal{F}_τ with respect to info and *abstract methods* do^\sharp , send^\sharp and receive^\sharp is a collection of relations $\{\mathcal{R}_r, \mathcal{M} \mid r \in \text{ReplicaID}\}$ satisfying the conditions in Fig. 5.

Here info and abstract methods are meant to be fixed for a given class of implementations, such as state or op-based. To prove a particular implementation within this class, one needs to find simulation relations satisfying the conditions in Fig. 5. For example, as we show in §5.3, the following relation lets us prove the correctness of the counter in Fig. 2(b) with respect to info and abstract methods appropriate for state-based implementations:

$$\begin{aligned} \langle s, v \rangle [\mathcal{R}_r] I &\iff (r = s) \wedge (v [\mathcal{M}] I); \\ v [\mathcal{M}] ((E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar}), \text{info}) &\iff \\ \forall s. v(s) &= |\{e \in E \mid \text{oper}(e) = \text{inc} \wedge \text{repl}(e) = s\}|. \end{aligned} \quad (6)$$

INIT in Fig. 5 associates the initial state at a replica r with the execution having an empty set of events. DO , SEND and RECEIVE formalize the obligations illustrated in Fig. 4(a, b). Note that do^\sharp is parameterized by an event e (required to be fresh in instantiations) and the information about the operation performed.

The abstract methods are partial and the obligations in Fig. 5 assume that their applications are defined. When instantiating receive^\sharp for a given class of implementations, we let it be defined only when its arguments satisfy the agreement property for this class, which we establish separately (§5.4). While doing this, we can also establish some *execution invariants*, holding of single ex-

Figure 5. Definition of a replication-aware simulation $\{\mathcal{R}_r, \mathcal{M}\}$ between \mathcal{D}_r and \mathcal{F}_r . All free variables in each condition are implicitly universally quantified and have the following types: $\sigma, \sigma' \in \Sigma, m \in M, I, I', J \in \text{IEx}, e \in \text{Event}, r \in \text{ReplicaID}, o \in \text{Op}_r, a \in \text{Val}_r, t \in \text{Timestamp}$.

$$\begin{aligned} \mathcal{R}_r &\subseteq \Sigma \times \text{IEx}, \quad r \in \text{ReplicaID}; & \mathcal{M} &\subseteq M \times \text{IEx} \\ \text{do}^\sharp &: \text{IEx} \times \text{Event} \times \text{ReplicaID} \times \text{Op}_r \times \text{Val}_r \times \text{Timestamp} \rightarrow \text{IEx} \\ \text{send}^\sharp &: \text{IEx} \rightarrow \text{IEx} \times \text{IEx}; & \text{receive}^\sharp &: \text{IEx} \times \text{IEx} \rightarrow \text{IEx} \\ \text{INIT: } &\mathcal{D}_r.\bar{\sigma}_0(r) [\mathcal{R}_r] I_\emptyset, \text{ where } I_\emptyset.E = \emptyset \\ \text{DO: } &(\text{do}^\sharp(I, e, r, o, a, t) = I' \wedge \mathcal{D}_r.\text{do}(o, \sigma, t) = (\sigma', a) \wedge \sigma [\mathcal{R}_r] I) \\ &\implies (\sigma' [\mathcal{R}_r] I' \wedge a = \mathcal{F}_r(\text{ctxt}(I'.A, e))) \\ \text{SEND: } &(\text{send}^\sharp(I) = (I', J) \wedge \mathcal{D}_r.\text{send}(\sigma) = (\sigma', m) \wedge \sigma [\mathcal{R}_r] I) \\ &\implies (\sigma' [\mathcal{R}_r] I' \wedge m [\mathcal{M}] J) \\ \text{RECEIVE: } &(\text{receive}^\sharp(I, J) \downarrow \wedge \sigma [\mathcal{R}_r] I \wedge m [\mathcal{M}] J) \\ &\implies (\mathcal{D}_r.\text{receive}(\sigma, m) [\mathcal{R}_r] \text{receive}^\sharp(I, J)) \end{aligned}$$

ecutions supplied as parameters to do^\sharp and send^\sharp . We similarly assume them in Fig. 5 via the definedness of these abstract methods.

5.2 Instantiation for State-Based Implementations

Fig. 6 defines the domain AInfo and abstract methods appropriate for state-based implementations. In §5.4 we show that the existence of a simulation of Def. 8 with respect to these parameters implies $\mathcal{D}_r \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_r$ (Theorems 9 and 10). The do^\sharp method adds a fresh event e with the given attributes to I ; its timestamp t is recorded in info . In the resulting execution I' , the event e is the last one by its replica, observes all events in I and occupies the place in arbitration consistent with t . The send^\sharp method just returns I , which formalizes the intuition that, in state-based implementations, send returns a message capturing all the information about the object available at the replica. The receive^\sharp method takes the component-wise union $I \sqcup J$ of executions I related to the current state and J related to the message, applied recursively to the components of $I.A$ and $J.A$. We also assume that $I \sqcup J$ recomputes the arbitration relation in the resulting execution from the timestamps. This is the reason for recording them in info : we would not be able to construct $\text{receive}^\sharp(I, J).\text{ar}$ solely from $I.\text{ar}$ and $J.\text{ar}$.

The agreement property $\text{agree}(I, J)$ guarantees that $I \sqcup J$ is well-formed (e.g., its visibility relation is acyclic) and that, for each replica, I describes a computation extending J or vice versa. The latter follows from the observation we made when explaining the state-based counter in §2: a message or a state in a state-based implementation reflects a prefix of the sequence of events performed at a given replica. The first conjunct of the execution invariant inv requires arbitration to be consistent with event timestamps; the second conjunct follows from the definition of $\mathcal{V}^{\text{state}}$ (§4). When discharging the obligations in Fig. 5 with respect to the parameters in Fig. 6 for a particular implementation, we can rely on the agreement property and the execution invariant.

5.3 Examples

We illustrate the use of the instantiation from §5.2 on the state-based counter, LWW-register and OR-set. In §A we also give proofs of two multi-value register implementations.

Counter: \mathcal{D}_{ctr} in Fig. 2(b) and \mathcal{F}_{ctr} in (3). Discharging the obligations in Fig. 5 for the simulation (6) is easy. The key case is **RECEIVE**, where the first conjunct of agree in Fig. 6 ensures that $\min\{v(s), v'(s)\}$ increments by a replica s are taken into account both in $v(s)$ and in $v'(s)$:

$$\begin{aligned} (\langle r, v \rangle [\mathcal{R}_r] I) \wedge (v' [\mathcal{M}] J) &\implies (\forall s. \min\{v(s), v'(s)\} = \\ & \{|e \in I.E \cap J.E \mid I.\text{oper}(e) = \text{inc} \wedge I.\text{repl}(e) = s\}|). \end{aligned}$$

Figure 6. Instantiation for state-based data type implementations. In do^\sharp we omit the straightforward definition of ar' in terms of $I.\text{info}$ and t .

$$\begin{aligned} \text{AInfo} &= \text{Timestamp} \\ \text{agree}(I, J) &\iff \forall r. ((\{e \in I.E \mid I.\text{repl}(e) = r\}, I.\text{ro}) \text{ is a prefix of} \\ & (\{e \in J.E \mid J.\text{repl}(e) = r\}, J.\text{ro}) \text{ or vice versa}) \wedge (I \sqcup J \in \text{IEx}) \\ \text{inv}(I) &\iff (\forall e, f \in I.E. (e, f) \in I.\text{ar} \iff I.\text{info}(e) < I.\text{info}(f)) \\ & \wedge ((I.\text{vis} \cup I.\text{ro})^+ \subseteq I.\text{vis}) \\ \text{do}^\sharp(I, e, r, o, a, t) &= I', \quad \text{if } \text{inv}(I) \wedge e \notin I.E \wedge I' \in \text{IEx} \\ \text{where } I' &= ((I.E \cup \{e\}, I.\text{repl}[e \mapsto r], I.\text{obj}[e \mapsto x], I.\text{oper}[e \mapsto o], \\ & I.\text{rval}[e \mapsto a], I.\text{ro} \cup \{(f, e) \mid f \in I.E \wedge I.\text{repl}(f) = r\}, \\ & I.\text{vis} \cup \{(f, e) \mid f \in I.E\}, \text{ar}'), I.\text{info}[e \mapsto t]) \\ \text{send}^\sharp(I) &= (I, I), \quad \text{if } \text{inv}(I) \\ \text{receive}^\sharp(I, J) &= I \sqcup J, \quad \text{if } \text{inv}(I) \wedge \text{inv}(J) \wedge \text{agree}(I, J) \end{aligned}$$

This allows establishing $\text{receive}(\langle r, v \rangle, v') [\mathcal{R}_r] (I \sqcup J)$, thus formalizing the informal justification of correctness we gave in §2.

LWW-register: $\mathcal{D}_{\text{intreg}}$ in Fig. 2(c) and $\mathcal{F}_{\text{intreg}}$ in (4). We associate a state or a message $\langle a, t \rangle$ with any execution that contains a $\text{wr}(a)$ event with the timestamp t maximal among all other wr events (as per info). By inv in Fig. 6, this event is maximal in arbitration, which implies that rd returns the correct value; the other obligations are also discharged easily. Formally, $\forall r. \mathcal{R}_r = \mathcal{M}$ and

$$\begin{aligned} \langle a, t \rangle [\mathcal{M}] ((E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar}), \text{info}) &\iff \\ (t = 0 \wedge a = 0 \wedge (\neg \exists e \in E. \text{oper}(e) = \text{wr}(-))) \vee \\ (t > 0 \wedge (\exists e \in E. \text{oper}(e) = \text{wr}(a) \wedge \text{info}(e) = t) \\ \wedge (\forall f \in E. \text{oper}(f) = \text{wr}(-) \implies \text{info}(f) \leq t)). \end{aligned}$$

Optimized OR-set: $\mathcal{D}_{\text{orset}}$ in Fig. 7 and $\mathcal{F}_{\text{orset}}$ in (5). A problem with implementing a replicated set is that we often cannot discard the information about an element from a replica state after it has been removed: if another replica unaware of the removal sends us a snapshot of its state containing this element, the semantics of the set may require our receive to keep the element out of the set. As we prove in §6, for the OR-set keeping track of information about removed elements cannot be fully avoided, which makes its space-efficient implementation very challenging. Here we consider a recently-proposed OR-set implementation [6] that, as we show in §6, has an optimal space complexity. It improves on the original implementation [31], whose complexity was suboptimal (we have proved the correctness of the latter as well; see §A).

An additional challenge posed by the OR-set is that, according to $\mathcal{F}_{\text{orset}}$, a remove operation may behave differently with respect to different events adding the same element to the set, depending on whether it sees them or not. This causes the implementation to treat internally each add operation as generating a unique *instance* of the *element* being added, further increasing the space required. To combat this, the implementation concisely summarizes information about instances. An instance is represented by a unique *instance identifier* that is generated when a replica performs an add and consists of the replica identifier and the number of adds (of any elements) performed at the replica until then. In a state $\langle r, V, w \rangle$, the vector w determines the identifiers of all instances that the current replica r has ever observed: for any replica s , the replica r has seen $w(s)$ successive identifiers $(s, 1), (s, 2), \dots, (s, w(s))$ generated at s . To generate a new identifier in $\text{do}(\text{add}(a'))$, the replica r increments $w(r)$. The connection between the vector w in a state or a message and add events $e_{s,k}$ in corresponding executions is formalized in lines 1-3 of the simulation relation, also shown in Fig. 7. In receive we take the pointwise maximum of the two vectors w and w' . Like for the counter, the first conjunct of agree implies that this preserves the clauses in lines 1-3.

Figure 7. Optimized OR-set implementation [6] and its simulation

$$\Sigma = \text{ReplicaID} \times ((\mathbb{Z} \times \text{ReplicaID}) \rightarrow \mathbb{N}_0) \times (\text{ReplicaID} \rightarrow \mathbb{N}_0)$$

$$\vec{\sigma}_0 = \lambda r. \langle r, (\lambda a, s. 0), (\lambda s. 0) \rangle$$

$$M = ((\mathbb{Z} \times \text{ReplicaID}) \rightarrow \mathbb{N}_0) \times (\text{ReplicaID} \rightarrow \mathbb{N}_0)$$

$$\text{do}(\text{add}(a'), \langle r, V, w \rangle, t) = (\langle r, (\lambda a, s. \text{if } a = a' \wedge s = r \\ \text{then } w(r) + 1 \text{ else } V(a, s)), w[r \mapsto w(r) + 1] \rangle, \perp)$$

$$\text{do}(\text{remove}(a'), \langle r, V, w \rangle, t) = (\langle r, (\lambda a, s. \text{if } a = a' \text{ then } 0 \text{ else } V(a, s)), w \rangle, \perp)$$

$$\text{do}(\text{rd}, \langle r, V, w \rangle, t) = (\langle r, V, w \rangle, \{a \mid \exists s. V(a, s) > 0\})$$

$$\text{send}(\langle r, V, w \rangle) = (\langle r, V, w \rangle, \langle V, w \rangle)$$

$$\text{receive}(\langle r, V, w \rangle, \langle V', w' \rangle) = (\langle r, (\lambda a, s. \text{if } (V(a, s) = 0 \wedge w(s) \geq V'(a, s)) \vee \\ (V'(a, s) = 0 \wedge w'(s) \geq V(a, s)) \\ \text{then } 0 \text{ else } \max\{V(a, s), V'(a, s)\}), \\ (\lambda s. \max\{w(s), w'(s)\}) \rangle)$$

$$\langle s, V, w \rangle [\mathcal{R}_r] I \iff (r = s) \wedge (\langle V, w \rangle [\mathcal{M}] I)$$

$$\langle V, w \rangle [\mathcal{M}] ((E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar}), \text{info}) \iff$$

- 1: $\exists \text{ distinct } e_{s,k}. \{\{e_{s,k} \mid s \in \text{ReplicaID} \wedge 1 \leq k \leq w(s)\} \wedge$
- 2: $\{e \in E \mid \text{oper}(e) = \text{add}(\cdot)\} \wedge$
- 3: $(\forall s, k, j. (\text{repl}(e_{s,k}) = s) \wedge (e_{s,j} \xrightarrow{\text{ro}} e_{s,k} \iff j < k)) \wedge$
- 4: $(\forall a, s. (V(a, s) \leq w(s)) \wedge (V(a, s) \neq 0 \implies$
- 5: $(\text{oper}(e_{s,V(a,s)}) = \text{add}(a)) \wedge$
- 6: $(\neg \exists k. V(a, s) < k \leq w(s) \wedge \text{oper}(e_{s,k}) = \text{add}(a)) \wedge$
- 7: $(\neg \exists f \in E. \text{oper}(f) = \text{remove}(a) \wedge e_{s,V(a,s)} \xrightarrow{\text{vis}} f)) \wedge$
- 8: $(\forall a, s, k. e_{s,k} \in E \wedge \text{oper}(e_{s,k}) = \text{add}(a) \implies$
- 9: $(k \leq V(a, s) \vee \exists f \in E. \text{oper}(f) = \text{remove}(a) \wedge e_{s,k} \xrightarrow{\text{vis}} f))$

The component w in $\langle r, V, w \rangle$ records identifiers of both of those instances that have been removed and those that are still in the set (are *active*). The component V serves to distinguish the latter. As it happens, we do not need to store all active instances of an element a : for every replica s , it is enough to keep the last active instance identifier generated by an $\text{add}(a)$ at this replica. If $V(a, s) \neq 0$, this identifier is $(s, V(a, s))$; if $V(a, s) = 0$, all instances of a generated at s that the current replica knows about are inactive. The meaning of V is formalized in the simulation: each instance identifier given by V is covered by w (line 4) and, if $V(a, s) \neq 0$, then the event $e_{s,V(a,s)}$ performs $\text{add}(a)$ (line 5), is the last $\text{add}(a)$ by replica s (line 6) and has not been observed by a $\text{remove}(a)$ (line 7). Finally, the $\text{add}(a)$ events that are not seen by a $\text{remove}(a)$ in the execution are either the events $e_{s,V(a,s)}$ or those superseded by them (lines 8-9). This ensures that returning all elements with an active instance in rd matches $\mathcal{F}_{\text{orset}}$.

When a replica r performs $\text{do}(\text{add}(a'))$, we update $V(a', r)$ to correspond to the new instance identifier. Conversely, in $\text{do}(\text{remove}(a'))$, we clear all entries in $V(a')$, thereby deactivating all instances of a' . However, after this their identifiers are still recorded in w , and so we know that they have been previously removed. This allows us to address the problem with implementing receive we mentioned above: if we receive a message with an active instance $(s, V'(a, s))$ of an element a that is not in the set at our replica ($V(a, s) = 0$), but previously existed ($w(s) \geq V'(a, s)$), this means that the instance has been removed and should not be active in the resulting state (the entry for (a, s) should be 0). We also do the same check with the state and the message swapped.

As the above explanation shows, our simulation relations are useful not only for proving correctness of data type implementations, but also for explaining their designs. Discharging obligations in Fig. 5 requires some work for the OR-set; due to space constraints, we defer this to §A.

Figure 8. Function step that mirrors the effect of an event $e \in C'.E$ from $C' \in \text{CEX}[x]$ in $D \in \text{DEX}$, defined when so is the abstract method used

$$\text{step}(C', e, D) =$$

$$D[r \mapsto \text{do}^\sharp(D(r), e, r, C'.\text{oper}(e), C'.\text{rval}(e), C'.\text{time}(e))],$$

$$\text{if } C'.\text{act}(e) = \text{do} \wedge C'.\text{repl}(e) = r$$

$$\text{step}(C', e, D) = D[r \mapsto I, C'.\text{msg}(e) \mapsto J],$$

$$\text{if } C'.\text{act}(e) = \text{send} \wedge C'.\text{repl}(e) = r \wedge C'.\text{msg}(e) \notin \text{dom}(D) \wedge$$

$$\text{send}^\sharp(D(r)) = (I, J)$$

$$\text{step}(C', e, D) = D[r \mapsto \text{receive}^\sharp(D(r), D(C'.\text{msg}(e)))],$$

$$\text{if } C'.\text{act}(e) = \text{receive} \wedge C'.\text{repl}(e) = r$$

5.4 Soundness and Establishing Agreement Properties

We present conditions on Alfo and abstract methods ensuring the soundness of replication-aware simulations over them and, in particular, establishing the agreement property and execution invariants assumed via the definedness of abstract operations in Fig. 5.

THEOREM 9 (Soundness). *Assume Alfo , do^\sharp , send^\sharp , receive^\sharp , \mathcal{V} and \mathcal{T} that satisfy the conditions in Fig. 9 for some \mathcal{G} . If there exists a replication-aware simulation between \mathcal{D}_τ and \mathcal{F}_τ with respect to these parameters, then $\mathcal{D}_\tau \text{ sat}[\mathcal{V}, \mathcal{T}] \mathcal{F}_\tau$.*

Conditions in Fig. 9 require global reasoning, but can be discharged once for a class of data types. For example, they hold of the instantiation for state-based implementations from §5.2, as well as one for op-based implementations presented in §B.

THEOREM 10. *There exists \mathcal{G} such that, for all \mathcal{D}_τ , the parameters in Fig. 6 satisfy the conditions in Fig. 9 with respect to this \mathcal{G} , $\mathcal{V} = \mathcal{V}^{\text{state}}$, $\mathcal{T} = \text{T-Any}$.*

The proofs of Theorems 9 and 10 are given in §B. To explain the conditions in Fig. 9, here we consider the proof strategy for Theorem 9. To establish $\mathcal{D}_\tau \text{ sat}[\mathcal{V}, \mathcal{T}] \mathcal{F}_\tau$, for any $C \in \llbracket \mathcal{D}_\tau \rrbracket \cap \mathcal{T}$ we need to show $\text{abs}(C, \mathcal{V}) \models [\tau \mapsto \mathcal{F}_\tau]$. We prove this by induction on the length of C . To use the localized conditions in Fig. 5, we require a relation \mathcal{G} associating C with a **decomposed execution**—a partial function $D : (\text{ReplicaID} \cup \text{MessageID}) \rightarrow \text{IEx}$ that gives fragments of $\text{abs}(C, \mathcal{V})$ corresponding to replica states and messages in the final configuration of C . We write DEX for the set of all decomposed executions, so that $\mathcal{G} \subseteq \text{CEX}[x] \times \text{DEX}$. The existence of a decomposed execution D such that $C \llbracket \mathcal{G} \rrbracket D$ forms the core of our induction hypothesis. G-CTXT in Fig. 9 checks that the abstract methods construct visibility according to \mathcal{V} : it requires the context of any event e by a replica r to be the same in $D(r)$ and $\text{abs}(C, \mathcal{V})$. Together with DO in Fig. 5, this ensures $\text{abs}(C, \mathcal{V}) \models [\tau \mapsto \mathcal{F}_\tau]$.

We write $C' \sim (C \xrightarrow{e} (R, T))$ when C' is an extension of C in the following sense: $C'.E = C.E \uplus \{e\}$, the other components of C' are those of C restricted to $C.E$, e is last in $C'.\text{eo}$ and $C'.\text{post}(e) = (R, T)$. For the induction step, assume $C \llbracket \mathcal{G} \rrbracket D$ and $C' \sim (C \xrightarrow{e} (R, T))$; see Fig. 4(c). Then the decomposed execution D' corresponding to C' is given by $\text{step}(C', e, D)$, where the function step in Fig. 8 mirrors the effect of the event e from C' in D using the abstract methods. G-STEP ensures that it preserves the relation \mathcal{G} . Crucially, G-STEP also requires us to establish the definedness of step and thus the corresponding abstract method. This justifies the agreement property and execution invariants encoded by the definedness and allows us to use the conditions in Fig. 5 to complete the induction. We also require G-INIT, which establishes the base case, and G-VIS, which formulates a technical restriction on \mathcal{V} . Finally, the conditions in Fig. 9 allow us to use the transport specification \mathcal{T} by considering only executions C satisfying it.

6. Space Bounds and Implementation Optimality

Object states in replicated data type implementations include not only the current client-observable content, but also metadata

Figure 9. Proof obligations for abstract methods. Free variables are implicitly universally quantified and have the following types: $C, C' \in \text{CEx}[x] \cap \mathcal{T}$, $D \in \text{DEX}$, $r \in \text{ReplicaID}$, $e \in \text{Event}$, $(R, T) \in \text{Config}$.

G-CTXT:	$(C \llbracket \mathcal{G} \rrbracket D \wedge e \in \text{abs}(C, \mathcal{V}).E \wedge \text{abs}(C, \mathcal{V}).\text{repl}(e) = r)$ $\implies \text{ctxt}(D(r).A, e) = \text{ctxt}(\text{abs}(C, \mathcal{V}), e)$
G-STEP:	$(C' \sim (C \xrightarrow{e} (R, T)) \wedge (C \llbracket \mathcal{G} \rrbracket D))$ $\implies (\text{step}(C', e, D) \downarrow \wedge C' \llbracket \mathcal{G} \rrbracket \text{step}(C', e, D))$
G-INIT:	$(C.E = \{e\} \wedge C.\text{pre}(e) = (-, []))$ $\implies (\text{step}(C, e, D_\emptyset) \downarrow \wedge C \llbracket \mathcal{G} \rrbracket \text{step}(C, e, D_\emptyset))$, where D_\emptyset is such that $\text{dom}(D_\emptyset) = \text{ReplicaID} \wedge$ $\forall r \in \text{ReplicaID}. D_\emptyset(r).E = \emptyset$
G-VIS:	$(e \in \text{abs}(C, \mathcal{V}).E \wedge (C \text{ is a prefix of } C' \text{ under } C'.\text{eo}))$ $\implies \text{ctxt}(\text{abs}(C, \mathcal{V}), e) = \text{ctxt}(\text{abs}(C', \mathcal{V}), e)$

needed for conflict resolution or masking network failures. Space taken by this metadata is a major factor determining their efficiency and feasibility. As illustrated by the OR-set in §5.3, this is especially so for state-based implementations, i.e., those that satisfy their data type specifications with respect to the visibility witness $\mathcal{V}^{\text{state}}$ and the transport layer specification T-Any. We now present a general technique for proving lower bounds on this space overhead, which we use to prove optimality of four state-based implementations (we leave other implementation classes for future work; see §9). As in §5, we only consider executions over a fixed object x of type τ .

6.1 Metadata Overhead

To measure space, we need to consider how data are represented. An **encoding** of a set S is an injective function $\text{enc} : S \rightarrow \Lambda^+$, where Λ is some suitably chosen fixed finite set of characters (left unspecified). Sometimes, we clarify the domain being encoded using a subscript: e.g., $\text{enc}_{\mathbb{N}_0}(1)$. For $s \in S$, we let $\text{len}_S(s)$ be the length of $\text{enc}_S(s)$. The length can vary: e.g., for an integer k , $\text{len}_{\mathbb{N}_0}(k) \in \Theta(\lg k)$. We use standard encodings (listed in §C) for return values $\text{enc}_{\text{val}_\tau}$ of the data types τ we consider and assume an arbitrary but fixed encoding of object states $\text{enc}_{\mathcal{D}_\tau, \Sigma}$.

To distinguish metadata from the client-observable content of the object, we assume that each data type has a special `rd` operation that returns the latter, as is the case in the examples considered so far. For a concrete execution $C \in \llbracket \mathcal{D}_\tau \rrbracket$ over the object x and a read event $e \in (C.E)|_{\text{rd}}$, we define $\text{state}(e)$ to be the state of the object accessed at e : $\text{state}(e) = R(x, C.\text{repl}(e))$ for $(R, -) = C.\text{pre}(e)$.

We now define the metadata overhead as a ratio, by dividing the size of the object state by the size of the observable state. We then quantify the worst-case overhead by taking the maximum of this ratio over all read operations in all executions with given numbers of replicas n and update operations m . To define the latter, we assume that each data type τ specifies a set $\text{Upd}_\tau \subset \text{Op}_\tau$ of update operations; for all examples in this paper $\text{Upd}_\tau = \text{Op}_\tau \setminus \{\text{rd}\}$.

DEFINITION 11. *The **maximum metadata overhead** of an execution $C \in \llbracket \mathcal{D}_\tau \rrbracket$ of an implementation \mathcal{D}_τ is*

$$\text{mmo}(\mathcal{D}_\tau, C) = \max \left\{ \frac{\text{len}_{\mathcal{D}_\tau, \Sigma}(\text{state}(e))}{\text{len}_{\text{val}_\tau}(C.\text{rval}(e))} \mid e \in (C.E)|_{\text{rd}} \right\}.$$

*The **worst-case metadata overhead** of an implementation \mathcal{D}_τ over all executions with n replicas and m updates ($2 \leq n \leq m$) is*

$$\begin{aligned} \text{wcmo}(\mathcal{D}_\tau, n, m) &= \max \{ \text{mmo}(\mathcal{D}_\tau, C) \mid C \in \llbracket \mathcal{D}_\tau \rrbracket \wedge \\ &\quad n = |\{C.\text{repl}(e) \mid e \in C.E\}| \wedge \\ &\quad m = |\{e \in C.E \mid C.\text{oper}(e) \in \text{Upd}_\tau\}| \}. \end{aligned}$$

We consider only executions with $m \geq n$, since we are interested in the asymptotic overhead of executions where all replicas are mutated (i.e., perform at least one update operation).

Figure 10. Summary of bounds on metadata overhead for stated-based implementations, as functions of the number of replicas n and updates m

Type	Existing implementation			Any implementation overhead
	algorithm	ref.	overhead	
ctr	Fig. 2(b)	[31]	$\widehat{\Theta}(n)$	$\widehat{\Omega}(n)$
orset	Fig. 7	[6]	$\widehat{\Theta}(n \lg m)$	$\widehat{\Omega}(n \lg m)$
	Fig. 15, §A	[31]	$\widehat{\Theta}(m \lg m)$	
intreg	Fig. 2(c)	[31]	$\widehat{\Theta}(\lg m)^\dagger$	$\widehat{\Omega}(\lg m)$
mvr	Fig. 17, §A	new [‡]	$\widehat{\Theta}(n \lg m)$	$\widehat{\Omega}(n \lg m)$
	Fig. 16, §A	[31]	$\widehat{\Theta}(n^2 \lg m)$	

[†] Assuming timestamp encoding is $O(\lg m)$, satisfied by Lamport clocks.

[‡] An optimization of [31] discovered during the optimality proof.

DEFINITION 12. *Assume \mathcal{D}_τ and a positive function $f(n, m)$.*

- f is an **asymptotic upper bound** ($\mathcal{D}_\tau \in \widehat{O}(f(n, m))$) if*
 $\sup_{n, m \rightarrow \infty} (\text{wcmo}(\mathcal{D}_\tau, n, m) / f(n, m)) < \infty$, i.e.,
 $\exists K > 0. \forall m \geq n \geq 2. \text{wcmo}(\mathcal{D}_\tau, n, m) < K f(n, m)$;
- f is an **asymptotic lower bound** ($\mathcal{D}_\tau \in \widehat{\Omega}(f(n, m))$) if*
 $\lim_{n, m \rightarrow \infty} (\text{wcmo}(\mathcal{D}_\tau, n, m) / f(n, m)) \neq 0$, i.e.,
 $\exists K > 0. \forall m_0 \geq n_0 \geq 2. \exists n \geq n_0, m \geq n_0.$
 $\text{wcmo}(\mathcal{D}_\tau, n, m) > K f(n, m)$;
- f is an **asymptotically tight bound** ($\mathcal{D}_\tau \in \widehat{\Theta}(f(n, m))$) if it is both an upper and a lower asymptotic bound.*

Fig. 10 summarizes our results; as described in §5, we have proved all the implementations correct. Matching lower and upper bounds indicate worst-case optimality of an implementation (note that this is different from optimality in all cases). The derivation of upper bounds relies on standard techniques and is deferred to §C. We now proceed to the main challenge: how to derive lower bounds that apply to *any* implementation of τ . We present proofs for `ctr` and `orset`; `intreg` and `mvr` are covered in §C.

6.2 Experiment Families

The goal is to show that for any correct implementation \mathcal{D}_τ (i.e., such that $\mathcal{D}_\tau \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_\tau$), the object state must store some minimum amount of information. We achieve this by constructing an *experiment family*, which is a collection of executions \mathbb{C}_α , where $\alpha \in Q$ for some index set Q . Each experiment contains a distinguished read event e_α . The experiments are designed in such a way that the object states $\text{state}(e_\alpha)$ must be distinct, which then implies a lower bound $\lfloor \lg_{|\Lambda|} |Q| \rfloor$ on the size of their encoding. To prove that they are distinct, we construct black-box tests that execute the methods of \mathcal{D}_τ on the states and show that the tests must produce different results for each state (e_α) provided \mathcal{D}_τ is correct. Formally, the tests induce a read-back function rb that satisfies $\text{rb}(\text{state}(e_\alpha)) = \alpha$. We encapsulate the core argument in the following lemma.

DEFINITION 13. *An **experiment family** for an implementation \mathcal{D}_τ is a tuple $(Q, n, m, \mathbb{C}, \mathbf{e}, \text{rb})$ where Q is a finite set, $2 \leq n \leq m$, and for each $\alpha \in Q$, $\mathbb{C}_\alpha \in \llbracket \mathcal{D}_\tau \rrbracket$ is an execution with n replicas and m updates, $\mathbf{e}_\alpha \in (\mathbb{C}_\alpha.E)|_{\text{rd}}$ and $\text{rb} : \mathcal{D}_\tau.\Sigma \rightarrow Q$ is a function satisfying $\text{rb}(\text{state}(e_\alpha)) = \alpha$.*

LEMMA 14. *If $(Q, n, m, \mathbb{C}, \mathbf{e}, \text{rb})$ is an experiment family, then*

$$\text{wcmo}(\mathcal{D}_\tau, n, m) \geq \lfloor \lg_{|\Lambda|} |Q| \rfloor / (\max_{\alpha \in Q} \text{len}(\mathbb{C}_\alpha.\text{rval}(\mathbf{e}_\alpha))).$$

PROOF. Since $\text{rb}(\text{state}(e_\alpha)) = \alpha$, the states $\text{state}(e_\alpha)$ are pairwise distinct and so are their encodings $\text{enc}(\text{state}(e_\alpha))$. Since there are fewer than $|Q|$ strings of length strictly less than $\lfloor \lg_{|\Lambda|} |Q| \rfloor$, for

some $\alpha \in Q$ we have $\text{len}(\text{enc}(\text{state}(\mathbf{e}_\alpha))) \geq \lfloor \lg_{|\Lambda|} |Q| \rfloor$. Then

$$\text{wcmo}(\mathcal{D}_\tau, n, m) \geq \text{mmo}(\mathcal{D}_\tau, \mathbb{C}_\alpha) \geq \frac{\text{len}(\text{state}(\mathbf{e}_\alpha))}{\text{len}(\mathbb{C}_\alpha.\text{rval}(\mathbf{e}_\alpha))} \geq \frac{\lfloor \lg_{|\Lambda|} |Q| \rfloor}{\max_{\alpha' \in Q} \text{len}(\mathbb{C}_{\alpha'}.\text{rval}(\mathbf{e}_{\alpha'}))}. \quad \square$$

To apply this lemma to the best effect, we need to find experiment families with $|Q|$ as large as possible and $\text{len}(\mathbb{C}_{\alpha'}.\text{rval}(\mathbf{e}_{\alpha'}))$ as small as possible. Finding such families is challenging, as there is no systematic way to derive them. We relied on intuitions about “which situations force replicas to store a lot of information” when searching for experiment families.

Driver programs. We define experiment families using *driver programs* (e.g., see Fig. 11). These are written in imperative pseudocode and use traditional constructs like loops and conditionals. As they execute, they construct concrete executions of the data type library $[\tau \mapsto \mathcal{D}_\tau]$ by means of the following instructions, each of which triggers a uniquely-determined transition from Fig. 3:

$\text{do}_r o^t$ do operation o on x at replica r with timestamp t
 $u \leftarrow \text{do}_r o^t$ same, but assign the return value to u
 $\text{send}_r(\text{mid})$ send a message for x with identifier mid at r
 $\text{receive}_r(\text{mid})$ receive the message mid at replica r

Programs explicitly supply timestamps for do and message identifiers for send and receive. We require that they do this correctly, e.g., respect uniqueness of timestamps. When a driver program terminates, it may produce a return value. For a program P , an implementation \mathcal{D}_τ , and a configuration (R, T) , we let $\text{exec}(\mathcal{D}_\tau, (R, T), P)$ be the concrete execution of the data type library $[\tau \mapsto \mathcal{D}_\tau]$ starting in (R, T) that results from running P ; we define $\text{result}(\mathcal{D}_\tau, (R, T), P)$ as the return value of P in this run.

6.3 Lower Bound for State-Based Counter (ctr)

THEOREM 15. *If $\mathcal{D}_{\text{ctr}} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{ctr}}$, then \mathcal{D}_{ctr} is $\widehat{\Omega}(n)$.*

We start by formulating a suitable experiment family.

LEMMA 16. *If $\mathcal{D}_{\text{ctr}} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{ctr}}$, $n \geq 2$ and $m \geq n$ is a multiple of $(n-1)$, then tuple $(Q, n, m, \mathbb{C}, \mathbf{e}, \text{rb})$ as defined in the left column of Fig. 11 is an experiment family.*

The idea of the experiments is to force replica 1 to remember one number for each of the other replicas in the system, which then introduces an overhead proportional to n ; cf. the implementation in Fig. 2(b). We show one experiment in Fig. 12. All experiments start with a common initialization phase, defined by *init*, where each of the replicas $2..n$ performs $m/(n-1)$ increments and sends a message after each increment. All messages remain undelivered until the second phase, defined by *exp*(α). There replica 1 receives exactly one message from each replica $r = 2..n$, selected using $\alpha(r)$. An experiment concludes with the read \mathbf{e}_α on the first replica.

The read-back works by performing separate tests for each of the replicas $r = 2..n$, defined by *test*(r). For example, to determine which message was sent by replica 2 during the experiment in Fig. 12, the program *test*(2): reads the counter value at replica 1, getting 12; delivers the final message by replica 2 to it; and reads the counter value at replica 1 again, getting 14. By observing the difference, the program can determine the message sent during the experiment: $\alpha(2) = 5 - (14 - 12) = 3$.


PROOF OF LEMMA 16. The only nontrivial obligation is to prove $\text{rb}(\text{state}(\mathbf{e}_\alpha)) = \alpha$. Let $(R_\alpha, T_\alpha) = \text{final}(\mathbb{C}_\alpha)$. Then

$$\begin{aligned} \alpha(r) &\stackrel{(i)}{=} \text{result}(\mathcal{D}_{\text{ctr}}, (R_0, T_0), (\text{init}; \text{exp}(\alpha); \text{test}(r))) \\ &= \text{result}(\mathcal{D}_{\text{ctr}}, (R_\alpha, T_\alpha), \text{test}(r)) \\ &\stackrel{(ii)}{=} \text{result}(\mathcal{D}_{\text{ctr}}, (R_{\text{init}}[(x, 1) \mapsto R_\alpha(x, 1)], T_{\text{init}}), \text{test}(r)) \\ &= \text{rb}(R_\alpha(x, 1))(r) = \text{rb}(\text{state}(\mathbf{e}_\alpha))(r), \end{aligned}$$

Figure 11. Experiment families $(Q, n, m, \mathbb{C}, \mathbf{e}, \text{rb})$ used in the proofs of Theorem 15 (ctr) and Theorem 17 (orset)

ctr	orset
Conditions on n, m (number of replicas/updates)	
$m \geq n \geq 2$	$m \geq n \geq 2$
$m \bmod (n-1) = 0$	$(m-1) \bmod (n-1) = 0$
Index set Q	
$Q = ([2..n] \rightarrow [1.. \frac{m}{n-1}])$	$Q = ([2..n] \rightarrow [1.. \frac{m-1}{n-1}])$
Family size $ Q $	
$ Q = (\frac{m}{n-1})^{n-1}$	$ Q = (\frac{m-1}{n-1})^{n-1}$
Driver programs	
<pre> procedure <i>init</i> for all $r \in [2..n]$ for all $i \in [1.. \frac{m}{n-1}]$ $\text{do}_r \text{inc}^{r+m+i}$ $\text{send}_r(\text{mid}_{r,i})$ procedure <i>exp</i>(α) for all $r \in [2..n]$ $\text{receive}_1(\text{mid}_{r, \alpha(r)})$ $\text{do}_1 \text{rd}^{(n+2)m}$ // read \mathbf{e}_α procedure <i>test</i>(r) $u \leftarrow \text{do}_1 \text{rd}^{(n+3)m}$ $\text{receive}_1(\text{mid}_{r, \frac{m}{n-1}})$ $u' \leftarrow \text{do}_1 \text{rd}^{(n+4)m}$ return $\frac{m}{n-1} - (u' - u)$ </pre>	<pre> procedure <i>init</i> for all $r \in [2..n]$ for all $i \in [1.. \frac{m-1}{n-1}]$ $\text{do}_r \text{add}(0)^{r+m+i}$ $\text{send}_r(\text{mid}_{r,i})$ procedure <i>exp</i>(α) for all $r \in [2..n]$ $\text{receive}_1(\text{mid}_{r, \alpha(r)})$ $\text{do}_1 \text{remove}(0)^{(n+2)m}$ $\text{do}_1 \text{rd}^{(n+3)m}$ // read \mathbf{e}_α procedure <i>test</i>(r) for all $i \in [1.. \frac{m-1}{n-1}]$ $\text{receive}_1(\text{mid}_{r,i})$ $u \leftarrow \text{do}_1 \text{rd}^{(n+4)m+i}$ if $0 \in u$ return $i - 1$ return $\frac{m-1}{n-1}$ </pre>
Definition of executions \mathbb{C}_α	
$\mathbb{C}_\alpha = \text{exec}(\mathcal{D}_\tau, (R_0, T_0), \text{init}; \text{exp}(\alpha))$	
where $(R_0, T_0) = ([x \mapsto \mathcal{D}_\tau.\sigma_0], \emptyset)$	
Definition of read-back function $\text{rb} : \mathcal{D}_\tau.\Sigma \rightarrow Q$	
$\text{rb}(\sigma) = \lambda r : [2..n]. \text{result}(\mathcal{D}_\tau, (R_{\text{init}}[(x, 1) \mapsto \sigma], T_{\text{init}}), \text{test}(r))$	
where $(R_{\text{init}}, T_{\text{init}}) = \text{post}(\text{exec}(\mathcal{D}_\tau, (R_0, T_0), \text{init}))$	

Figure 12. Example experiment ($n = 4$ and $m = 15$) and test for ctr. Gray dashed lines represent the configuration $(R_{\text{init}}[(x, 1) \mapsto R_\alpha(x, 1)], T_{\text{init}})$ where the *test* driver program is applied.



where:

- (i) This is due to $\mathcal{D}_{\text{ctr}} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{ctr}}$, as we explained informally above. Let

$$C'_\alpha = \text{exec}(\mathcal{D}_{\text{ctr}}, (R_0, T_0), (\text{init}; \text{exp}(\alpha); \text{test}(r))).$$

Then the operation context in $\text{abs}(C'_\alpha, \mathcal{V}^{\text{state}})$ of the first read in *test*(r) contains $\sum_{r=2}^n \alpha(r)$ increments, while that of the second read contains $(m/(n-1)) - \alpha(r)$ more increments.

- (ii) We have $T_\alpha = T_{\text{init}}$ because *exp*(α) does not send any messages. Also, R_α and $R_{\text{init}}[(x, 1) \mapsto R_\alpha(x, 1)]$ can differ only

in the states of the replicas $2..n$. These cannot influence the run of $test(r)$, since it performs events on replica 1 only. \square

PROOF OF THEOREM 15. Given n_0, m_0 , we pick $n = n_0$ and some $m \geq n_0$ such that m is a multiple of $(n - 1)$ and $m \geq n^2$. Take the experiment family $(Q, n, m, \mathbb{C}, e, rb)$ given by Lemma 16. Then for any α , $\mathbb{C}_\alpha.rval(e_\alpha)$ is at most the total number of increments m in \mathbb{C}_α . Using Lemma 14 and $m \geq n^2$, for some constants K_1, K_2, K_3, K independent from n_0, m_0 we get:

$$\begin{aligned} wcmo(\mathcal{D}_{ctr}, n, m) &\geq \lfloor \lg_{|\Lambda|} |Q| \rfloor / (\max_{\alpha \in Q} \text{len}(\mathbb{C}_\alpha.rval(e_\alpha))) \geq \\ K_1 \frac{\lg_{|\Lambda|} \left(\frac{m}{n-1}\right)^{n-1}}{\text{len}_{N_0}(m)} &\geq K_2 \frac{n \lg(m/n)}{\lg m} \geq K_3 \frac{n \lg \sqrt{m}}{\lg m} \geq Kn. \quad \square \end{aligned}$$

6.4 Lower Bound for State-Based OR-Set (orset)

THEOREM 17. *If $\mathcal{D}_{orset} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{orset}$, then \mathcal{D}_{orset} is $\Omega(n \lg m)$.*

LEMMA 18. *If $\mathcal{D}_{orset} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{orset}$, $n \geq 2$ and $m \geq n$ is such that $(m - 1)$ is a multiple of $(n - 1)$, then the tuple $(Q, n, m, \mathbb{C}, e, rb)$ on the right in Fig. 11 is an experiment family.*

The proof is the same as that of Lemma 16, except for obligation (i). We therefore give only informal explanations.

The main idea of the experiments defined in the lemma is to force replica 1 to remember element instances even after they have been removed at that replica; cf. our explanation of the challenges of implementing the OR-set from §5.3. The experiments follow a similar pattern to those for `ctr`, but use different operations. In the common *init* phase, each replica $2..n$ performs $\frac{m-1}{n-1}$ operations adding a designated element 0, which are interleaved with sending messages. In the experiment phase $exp(\alpha)$, one message from each replica $r = 2..n$, selected by $\alpha(r)$, is delivered to replica 1. At the end of execution, replica 1 removes 0 from the set and performs the read e_α . The return value of this read is always the empty set.

To perform the read-back of $\alpha(r)$ for $r = 2..n$, $test(r)$ delivers all messages by replica r to replica 1 in the order they were sent and, after each such delivery, checks if replica 1 now reports the element 0 as part of the set. From $\mathcal{D}_{orset} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{orset}$ and the definition (5) of \mathcal{F}_{orset} , we get that exactly the first $\alpha(r)$ such deliveries will have no effect on the contents of the set: the respective add operations have already been observed by the remove operation that replica 1 performed in the experiment phase. Thus, if 0 appears in the set right after delivering the i -th message of replica r , then $\alpha(r) = i - 1$, and if 0 does not appear by the time the loop is finished, then $\alpha(r) = (m - 1)/(n - 1)$.

PROOF OF THEOREM 17. Given n_0, m_0 , we pick $n = n_0$ and some $m \geq n_0$ such that $(m - 1)$ is a multiple of $(n - 1)$ and $m \geq n^2$. Take the experiment family $(Q, n, m, \mathbb{C}, e, rb)$ given by Lemma 18. For any $\alpha \in Q$, $\mathbb{C}_\alpha.rval(e_\alpha) = \emptyset$, which can be encoded with a constant length. Using Lemma 14 and $m \geq n^2$, for some constants K_1, K_2, K we get:

$$\begin{aligned} wcmo(\mathcal{D}_{orset}, n, m) &\geq \lfloor \lg_{|\Lambda|} |Q| \rfloor / (\max_{\alpha \in Q} \text{len}(\mathbb{C}_\alpha.rval(e_\alpha))) \\ &\geq K_1 n \lg(m/n) \geq K_2 n \lg \sqrt{m} \geq Kn \lg m. \quad \square \end{aligned}$$

7. Store Correctness and Consistency Axioms

Recall that we define a replicated store by a data type library \mathbb{D} and a transport layer specification \mathcal{T} (§4), and we specify its behavior by a function \mathbb{F} from types $\tau \in \text{dom}(\mathbb{D})$ to data type specifications and a set of consistency axioms (§3). The axioms are just constraints over abstract executions, such as those shown in Fig. 13; from now on we denote their sets by X . So far we have concentrated on single data type specifications $\mathbb{F}(\tau)$ and their correspondence to implementations $\mathbb{D}(\tau)$, as stated by Def. 7. In this section we consider consistency axioms and formulate the

notion of correctness of the whole store $(\mathbb{D}, \mathcal{T})$ with respect to its specification (\mathbb{F}, X) .

Our first goal is to lift the statement of correctness given by Def. 7 to a store $(\mathbb{D}, \mathcal{T})$ with multiple objects of different data types. To this end, we assume a function \mathbb{V} mapping each type $\tau \in \text{dom}(\mathbb{D})$ to its visibility witness \mathbb{V} . This allows us to construct the visibility relation for a concrete execution $C \in \llbracket \mathbb{D} \rrbracket \cap \mathcal{T}$ by applying $\mathbb{V}(\tau)$ to its projection onto the events on every object of type τ :

$$\text{witness}(\mathbb{V}) = \lambda C. \bigcup \{ \mathbb{V}(\text{type}(x))(C|_x) \mid x \in \text{Obj} \},$$

where $\cdot|_x$ projects to events over x . Then the correctness of every separate data type τ in the store with respect to $\mathbb{F}(\tau)$ according to Def. 7 automatically ensures that the behavior of the whole store is consistent with \mathbb{F} in the sense of Def. 5.

PROPOSITION 19. $(\forall \tau \in \text{dom}(\mathbb{D}). \mathbb{D}(\tau) \text{ sat}[\mathbb{V}(\tau), \mathcal{T}] \mathbb{F}(\tau)) \implies (\forall C \in \llbracket \mathbb{D} \rrbracket \cap \mathcal{T}. \text{abs}(C, \text{witness}(\mathbb{V})) \models \mathbb{F})$.

This motivates the following definition of store correctness. Let us write $A \models X$ when the abstract execution A satisfies the axioms X .

DEFINITION 20. A store $(\mathbb{D}, \mathcal{T})$ is **correct** with respect to a specification (\mathbb{F}, X) , if for some \mathbb{V} :

- (i) $\forall \tau \in \text{dom}(\mathbb{D}). (\mathbb{D}(\tau) \text{ sat}[\mathbb{V}(\tau), \mathcal{T}] \mathbb{F}(\tau))$; and
- (ii) $\forall C \in \llbracket \mathbb{D} \rrbracket \cap \mathcal{T}. (\text{abs}(C, \text{witness}(\mathbb{V})) \models X)$.

We showed how to discharge (i) in §5. The validity of axioms X required by (ii) most often depends on the transport layer specification \mathcal{T} : e.g., to disallow the anomaly (2) from §1, \mathcal{T} needs to provide guarantees on how messages pertaining to different objects are delivered. However, data type implementations can also enforce axioms by putting enough information into messages: e.g., implementations correct with respect to $\mathcal{V}^{\text{state}}$ from §4 ensure that *vis* is transitive regardless of the behavior of the transport layer. Fortunately, to establish (ii) in practice, we do not need to consider the internals of data type implementations in \mathbb{D} —just knowing the visibility witnesses used in the statements of their correctness is enough, as formulated in the following definition.

DEFINITION 21. A set W of visibility witnesses and a transport layer specification \mathcal{T} **validate** axioms X , if

$$\forall C, \mathbb{V}. (C \in \mathcal{T}) \wedge (\{ \mathbb{V}(\tau) \mid \tau \in \text{dom}(\mathbb{V}) \} \subseteq W) \implies (\text{abs}(C, \text{witness}(\mathbb{V})) \models X).$$

Since visibility witnesses are common to wide classes of data types (e.g., state- or op-based), our proofs of the validity of axioms will not have to be redone if we add new data type implementations to the store from a class already considered.

We next present axioms formalizing several variants of eventual consistency used in replicated stores (Fig. 13 and 14) and W and \mathcal{T} that validate them. We then use this as a basis for discussing connections with weak shared-memory models. Due to space constraints, we defer technical details and proofs to §D.

Basic eventual consistency. **EVENTUAL** and **THINAIR** define a weak form of eventual consistency. **EVENTUAL** ensures that an event cannot be invisible to infinitely many other events on the same object and thus implies (1) from §1: informally, if updates stop, then reads at all replicas will eventually see all updates and will return the same values (§3.2). However, **EVENTUAL** is stronger than quiescent consistency: the latter does not provide any guarantees at all for executions with infinitely many updates to the store, whereas our specification implies that the return values are computed according to $\mathbb{F}(\tau)$ using increasingly up-to-date view of the store state. We formalize these relationships in §D.

THINAIR prohibits values from appearing “out-of-thin-air” [27], like 42 in Fig. 14(a) (recall that registers are initialized to 0). Cycles in $\text{ro} \cup \text{vis}$ that lead to out-of-thin-air usually arise

Figure 13. A selection of consistency axioms over an execution $(E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar})$

Auxiliary relations	
$\text{sameobj}(e, f) \iff \text{obj}(e) = \text{obj}(f)$	
Per-object causality (aka happens-before) order: $\text{hbo} = ((\text{ro} \cap \text{sameobj}) \cup \text{vis})^+$	
Causality (aka happens-before) order: $\text{hb} = (\text{ro} \cup \text{vis})^+$	
Axioms	
EVENTUAL: $\forall e \in E. \neg(\exists \text{ infinitely many } f \in E. \text{sameobj}(e, f) \wedge \neg(e \xrightarrow{\text{vis}} f))$	
THINAIR: $\text{ro} \cup \text{vis}$ is acyclic	
POCV (Per-Object Causal Visibility): $\text{hbo} \subseteq \text{vis}$	
POCA (Per-Object Causal Arbitration): $\text{hbo} \subseteq \text{ar}$	
COCV (Cross-Object Causal Visibility): $(\text{hb} \cap \text{sameobj}) \subseteq \text{vis}$	
COCA (Cross-Object Causal Arbitration): $\text{hb} \cup \text{ar}$ is acyclic	

Figure 14. Anomalies allowed or disallowed by different axioms

(a) Disallowed by THINAIR: $x, y : \text{intreg}$ $i = x.\text{rd} \parallel j = y.\text{rd}$ $y.\text{wr}(i) \parallel x.\text{wr}(j)$	
(b) Disallowed by POCV: $x : \text{orset}$ $x.\text{add}(1) \parallel i = x.\text{rd} \parallel j = x.\text{rd}$ $x.\text{add}(2) \parallel x.\text{add}(3)$	
(c) Allowed by COCV and COCA: $x, y : \text{intreg}$ $x.\text{wr}(1) \parallel y.\text{wr}(1)$ $i = y.\text{rd} \parallel j = x.\text{rd}$	

from effects of speculative computations, which are done by some older replicated stores [35].

THINAIR is validated by $\{\mathcal{V}^{\text{state}}, \mathcal{V}^{\text{op}}\}$ and T-Any, and EVENTUAL by $\{\mathcal{V}^{\text{state}}, \mathcal{V}^{\text{op}}\}$ and the following condition on C ensuring that every message is eventually delivered to all other replicas and every operation is followed by a message generation:

$$\begin{aligned}
 (\forall e \in C.E. \forall r, r'. C.\text{act}(e) = \text{send} \wedge C.\text{repl}(e) = r \wedge r \neq r' \\
 \implies \exists f. C.\text{repl}(f) = r' \wedge e \xrightarrow{\text{del}(C)} f) \wedge \\
 (\forall e \in C.E. C.\text{act}(e) = \text{do} \implies \exists f. \text{act}(f) = \text{send} \wedge e \xrightarrow{\text{roo}(C)} f),
 \end{aligned}$$

where $\text{roo}(C)$ is $\text{ro}(C)$ projected to events on the same object.

Causality guarantees. Many replicated stores achieve availability and partition tolerance while providing stronger guarantees, which we formalize by the other axioms in Fig. 13. We call an execution *per-object*, respectively, *cross-object causally consistent*, if it is eventually consistent (as per above) and satisfies the axioms POCV and POCA, respectively, COCV and COCA. POCV guarantees that an operation sees all operations connected to it by a causal chain of events on the same object; COCV also considers causal chains via different objects. Thus, POCV disallows the execution in Fig. 14(b), and COCV the one in §3.1, corresponding to (2) from §1. POCA and COCA similarly require arbitration to be consistent with causality. The axioms highlight the principle of formalizing stronger consistency models: including more edges into vis and ar , so that clients have more up-to-date information.

Cross-object causal consistency is implemented by, e.g., COPS [26] and Gemini [22]. It is weaker than strong consistency, as it allows reading stale data. For example, it allows the execution in Fig. 14(c), where both reads fetch the initial value of the register, despite writes to it by the other replica. It is easy to check that this

outcome cannot be produced by any interleaving of the events at the two replicas, and is thus not strongly consistent.

An interesting feature of per-object causal consistency is that state-based data types ensure most of it just by the definition of $\mathcal{V}^{\text{state}}$: POCV is validated by $\{\mathcal{V}^{\text{state}}\}$ and T-Any. If the witness set is $\{\mathcal{V}^{\text{state}}, \mathcal{V}^{\text{op}}\}$, then we need \mathcal{T} to guarantee the following: informally, if a send event e and another event f are connected by a causal chain of events on the same object, then the message created by e is delivered to $C.\text{repl}(f)$ by the time f is done. POCA is validated by $\{\mathcal{V}^{\text{state}}, \mathcal{V}^{\text{op}}\}$ and the transport layer specification $(\text{roo}(C) \cup \text{del}(C))^+_{|\text{do}} \subseteq \text{ar}(C)$. This states that timestamps of events on every object behave like a Lamport clock [21]. Conditions for COCV and COCA are similar.

There also exist consistency levels in between basic eventual consistency and per-object causal consistency, defined using so-called *session guarantees* [34]. We cover them in §D.

Comparison with shared-memory consistency models. Interestingly, the specializations of the consistency levels defined by the axioms in Fig. 13 to the type `intreg` of LWW-registers are very close to those adopted by the memory model in the 2011 C and C++ standards [5]. Thus, POCA and POCV define the semantics of the fragment of C/C++ restricted to so-called *relaxed* operations; there this semantics is defined using *coherence* axioms, which are analogous to session guarantees [34]. COCV and COCA are close to the semantics of C/C++ restricted to *release-acquire* operations. However, C/C++ does not have an analog of EVENTUAL and does not validate THINAIR, since it makes the effects of speculations visible to the programmer [4]. We formalize the correspondence to C/C++ in §D. In the future, this correspondence may open the door to applying technology developed for shared-memory models to eventually consistent systems; promising directions include model checking [3, 9], automatic inference of required consistency levels [25] and compositional reasoning [4].

8. Related Work

For a comprehensive overview of replicated data type research we refer the reader to Shapiro et al. [31]. Most papers proposing new data type implementations [6, 30–32] do not provide their formal declarative specifications, save for the expected property (1) of quiescent consistency or first specification attempts for sets [6, 7]. Formalizations of eventual consistency have either expressed quiescent consistency [8] or gave low-level operational specifications [16].

An exception is the work of Burckhardt et al. [10, 12], who proposed an axiomatic model of causal eventual consistency based on visibility and arbitration relations and an operational model based on revision diagrams. Their store specification does not provide customizable consistency guarantees, and their data type specifications are limited to the sequential \mathcal{S} construction from §3.2, which cannot express advanced conflict resolution used by the multi-value register, the OR-set and many other data types [31]. More significantly, their operational model does not support general op- or state-based implementations, and is thus not suited for studying the correctness or optimality of these commonly used patterns.

Simulation relations have been applied to verify the correctness of sequential [24] and shared-memory concurrent data type implementations [23]. We take this approach to the more complex setting of a replicated store, where the simulation needs to take into account multiple object copies and messages and associate them with structures on events, rather than single abstract states. This poses technical challenges not considered by prior work, which we address by our novel notion of replication-aware simulations.

The distributed computing community has established a number of asymptotic lower bounds on the complexity of implementing certain distributed or concurrent abstractions, including one-

shot timestamp objects [19] and counting protocols [14, 29]. These works have considered either programming models or metrics significantly different from ours. An exception is the work of Charron-Bost [13], who proved that the size of vector clocks [28] is optimal to represent the happens-before relation of a computation (similar to the visibility relation in our model). Specifications of `mvr` and `orset` rely on visibility; however, Charron-Bost’s result does not directly translate into a lower bound on their implementation complexity, since a specification may not require complete knowledge about the relation and an implementation may represent it in an arbitrary manner, not necessarily using a vector.

9. Conclusion and Future Work

We have presented a comprehensive theoretical toolkit to advance the study of replicated eventually consistent stores, by proposing methods for (1) specifying the semantics of replicated data types and stores abstractly, (2) verifying implementations of replicated data types, and (3) proving that such implementations have optimal metadata overhead. By proving both correctness and optimality of four nontrivial data type implementations, we have demonstrated that our methods can indeed be productively applied to the kinds of patterns used by practitioners and researchers in this area.

Although our work marks a big step forward, it is only a beginning, and creates plenty of opportunities for future research. We have already made the first steps in extending our specification framework with more features, such as mixtures of consistency levels [22] and transactions [33, 36]; see [11]. In the future we would also like to study more data types, such as lists used for collaborative editing [31], and to investigate metadata bounds for data type implementations other than state-based ones, including more detailed overhead metrics capturing optimizations invisible to the worst-case overhead analysis. Even though our execution model for replicated stores follows the one used by replicated data type designers [32], there are opportunities for bringing it closer to actual implementations. Thus, we would like to verify the algorithms used by store implementations [26, 33, 36] that our semantics abstracts from. This includes fail-over and session migration protocols, which permit clients to interact with multiple physical replicas, while being provided the illusion of a single virtual replica.

Finally, by bringing together prior work on shared-memory models and data replication, we wish to promote an exchange of ideas and results between the research communities of programming languages and verification on one side and distributed systems on the other.

Acknowledgements. We thank Hagit Attiya, Anindya Banerjee, Carlos Baquero, Lindsey Kuper and Marc Shapiro for comments that helped improve the paper. Gotsman was supported by the EU FET project ADVENT, and Yang by EPSRC.

References

- [1] Riak key-value store. <http://basho.com/products/riak-overview/>.
- [2] S. V. Adve and K. Gharachorloo. Shared memory consistency models: A tutorial. *Computer*, 29(12), 1996.
- [3] J. Alglave, D. Kroening, V. Nimal, and M. Tautschnig. Software verification for weak memory via program transformation. In *ESOP*, 2013.
- [4] M. Batty, M. Dodds, and A. Gotsman. Library abstraction for C/C++ concurrency. In *POPL*, 2013.
- [5] M. Batty, S. Owens, S. Sarkar, P. Sewell, and T. Weber. Mathematizing C++ concurrency. In *POPL*, 2011.
- [6] A. Bieniusa, M. Zawirski, N. Preguiça, M. Shapiro, C. Baquero, V. Balegas, and S. Duarte. An optimized conflict-free replicated set. Technical Report 8083, INRIA, 2012.
- [7] A. Bieniusa, M. Zawirski, N. M. Preguiça, M. Shapiro, C. Baquero, V. Balegas, and S. Duarte. Brief announcement: Semantics of eventually consistent replicated sets. In *DISC*, 2012.
- [8] A.-M. Bosneag and M. Brockmeyer. A formal model for eventual consistency semantics. In *IASTED PDCS*, 2002.
- [9] S. Burckhardt, R. Alur, and M. M. K. Martin. Checkfence: checking consistency of concurrent data types on relaxed memory models. In *PLDI*, 2007.
- [10] S. Burckhardt, M. Fähndrich, D. Leijen, and B. P. Wood. Cloud types for eventual consistency. In *ECOOP*, 2012.
- [11] S. Burckhardt, A. Gotsman, and H. Yang. Understanding eventual consistency. Technical Report MSR-TR-2013-39, Microsoft Research, 2013.
- [12] S. Burckhardt, D. Leijen, M. Fähndrich, and M. Sagiv. Eventually consistent transactions. In *ESOP*, 2012.
- [13] B. Charron-Bost. Concerning the size of logical clocks in distributed systems. *Information Processing Letters*, 39(1), 1991.
- [14] J.-Y. Chen and G. Pandurangan. Optimal gossip-based aggregate computation. In *SPAA*, 2010.
- [15] N. Conway, R. Marczak, P. Alvaro, J. M. Hellerstein, and D. Maier. Logic and lattices for distributed programming. In *SOCC*, 2012.
- [16] A. Fekete, D. Gupta, V. Luchangco, N. Lynch, and A. Shvartsman. Eventually-serializable data services. In *PODC*, 1996.
- [17] G. DeCandia et al. Dynamo: Amazon’s highly available key-value store. In *SOSP*, 2007.
- [18] S. Gilbert and N. Lynch. Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News*, 33(2), 2002.
- [19] M. Helmi, L. Higham, E. Pacheco, and P. Woelfel. The space complexity of long-lived and one-shot timestamp implementations. In *PODC*, 2011.
- [20] C. A. R. Hoare. Proof of correctness of data representations. *Acta Inf.*, 1, 1972.
- [21] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7), 1978.
- [22] C. Li, D. Porto, A. Clement, R. Rodrigues, N. Preguiça, and J. Gehrke. Making geo-replicated systems fast if possible, consistent when necessary. In *OSDI*, 2012.
- [23] H. Liang, X. Feng, and M. Fu. A rely-guarantee-based simulation for verifying concurrent program transformations. In *POPL*, 2012.
- [24] B. Liskov and S. Zilles. Programming with abstract data types. In *ACM Symposium on Very High Level Languages*, 1974.
- [25] F. Liu, N. Nedeve, N. Prasadnikov, M. T. Vechev, and E. Yahav. Dynamic synthesis for relaxed memory models. In *PLDI*, 2012.
- [26] W. Lloyd, M. J. Freedman, M. Kaminsky, and D. G. Andersen. Don’t settle for eventual: scalable causal consistency for wide-area storage with COPS. In *SOSP*, 2011.
- [27] J. Manson, W. Pugh, and S. V. Adve. The Java memory model. In *POPL*, 2005.
- [28] F. Mattern. Virtual time and global states of distributed systems. *Parallel and Distributed Algorithms*, 1989.
- [29] S. Moran, G. Taubenfeld, and I. Yadin. Concurrent counting. In *PODC*, 1992.
- [30] H.-G. Roh, M. Jeon, J.-S. Kim, and J. Lee. Replicated abstract data types: Building blocks for collaborative applications. *J. Parallel Distrib. Comput.*, 71(3), 2011.
- [31] M. Shapiro, N. Preguiça, C. Baquero, and M. Zawirski. A comprehensive study of Convergent and Commutative Replicated Data Types. Technical Report 7506, INRIA, 2011.
- [32] M. Shapiro, N. M. Preguiça, C. Baquero, and M. Zawirski. Conflict-free replicated data types. In *SSS*, 2011.
- [33] Y. Sovran, R. Power, M. K. Aguilera, and J. Li. Transactional storage for geo-replicated systems. In *SOSP*, 2011.
- [34] D. B. Terry, A. J. Demers, K. Petersen, M. Spreitzer, M. Theimer, and B. W. Welch. Session guarantees for weakly consistent replicated data. In *PDIS*, 1994.
- [35] D. B. Terry, M. M. Theimer, K. Petersen, A. J. Demers, M. J. Spreitzer, and C. H. Hauser. Managing update conflicts in Bayou, a weakly connected replicated storage system. In *SOSP*, 1995.
- [36] M. Zawirski, A. Bieniusa, V. Balegas, S. Duarte, C. Baquero, M. Shapiro, and N. Preguiça. SwiftCloud: Fault-tolerant geo-replication integrated all the way to the client machine. Technical Report 8347, INRIA, 2013.

A. Additional proofs of data type correctness

Optimized OR-set (orset). We now discharge the obligations for the simulation for the optimized OR-set we presented in §5.3. The only interesting obligation is RECEIVE. Let

$$\text{receive}(\langle r, V, w \rangle, \langle V', w' \rangle) = \langle r, V'', w'' \rangle,$$

so that $w'' = \lambda s. \max\{w(s), w'(s)\}$ and

$$V'' = (\lambda a, s. \text{if } (V(a, s) = 0 \wedge w(s) \geq V'(a, s)) \vee (V'(a, s) = 0 \wedge w'(s) \geq V(a, s)) \text{ then } 0 \text{ else } \max\{V(a, s), V'(a, s)\}).$$

Assume $\langle r, V, w \rangle [\mathcal{R}_r] I, \langle V', w' \rangle [\mathcal{M}] J$ and

$$\begin{aligned} I &= ((E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar}), \text{info}); \\ J &= ((E', \text{repl}', \text{obj}', \text{oper}', \text{rval}', \text{ro}', \text{vis}', \text{ar}'), \text{info}'); \\ I \sqcup J &= ((E'', \text{repl}'', \text{obj}'', \text{oper}'', \text{rval}'', \text{ro}'', \text{vis}'', \text{ar}''), \text{info}''). \end{aligned}$$

By agree we have $I \sqcup J \in \text{IEx}$. Then

$$\begin{aligned} &\exists \text{ distinct } e_{s,k}. (\{e_{s,k} \mid s \in \text{ReplicaID} \wedge 1 \leq k \leq w(s)\} = \\ &\quad \{e \in E \mid \exists a. \text{oper}''(e) = \text{add}(a)\}) \wedge \\ &(\forall s, j, k. (\text{repl}''(e_{s,k}) = s) \wedge (e_{s,j} \xrightarrow{\text{ro}} e_{s,k} \iff j < k)) \wedge \\ &(\forall a, s. (V(a, s) \leq w(s)) \wedge (V(a, s) \neq 0 \implies \\ &\quad \text{oper}''(e_{s,V(a,s)}) = \text{add}(a) \wedge \\ &\quad (\neg \exists k. V(a, s) < k \leq w(s) \wedge \text{oper}''(e_{s,k}) = \text{add}(a)) \wedge \\ &\quad (\neg \exists f \in E. \text{oper}''(f) = \text{remove}(a) \wedge e_{s,V(a,s)} \xrightarrow{\text{vis}} f))) \wedge \\ &(\forall a, s, k. e_{s,k} \in E \wedge \text{oper}''(e_{s,k}) = \text{add}(a) \implies \\ &\quad (k \leq V(a, s) \vee \exists f \in E. \text{oper}''(f) = \text{remove}(a) \wedge e_{s,k} \xrightarrow{\text{vis}} f)) \end{aligned}$$

and

$$\begin{aligned} &\exists \text{ distinct } e'_{s,k}. (\{e'_{s,k} \mid s \in \text{ReplicaID} \wedge 1 \leq k \leq w'(s)\} = \\ &\quad \{e \in E' \mid \exists a. \text{oper}''(e) = \text{add}(a)\}) \wedge \\ &(\forall s, j, k. (\text{repl}''(e'_{s,k}) = s) \wedge (e'_{s,j} \xrightarrow{\text{ro}'} e'_{s,k} \iff j < k)) \wedge \\ &(\forall a, s. (V'(a, s) \leq w'(s)) \wedge (V'(a, s) \neq 0 \implies \\ &\quad \text{oper}''(e'_{s,V'(a,s)}) = \text{add}(a) \wedge \\ &\quad (\neg \exists k. V'(a, s) < k \leq w'(s) \wedge \text{oper}''(e'_{s,k}) = \text{add}(a)) \wedge \\ &\quad (\neg \exists f \in E'. \text{oper}''(f) = \text{remove}(a) \wedge e'_{s,V'(a,s)} \xrightarrow{\text{vis}'} f))) \wedge \\ &(\forall a, s, k. e'_{s,k} \in E' \wedge \text{oper}''(e'_{s,k}) = \text{add}(a) \implies \\ &\quad (k \leq V'(a, s) \vee \exists f \in E'. \text{oper}''(f) = \text{remove}(a) \wedge e'_{s,k} \xrightarrow{\text{vis}'} f)). \end{aligned}$$

The agree property also implies that

$$\forall s, k. 1 \leq k \leq \min\{w(s), w'(s)\} \implies e_{s,k} = e'_{s,k}.$$

Hence, there exist distinct $e''_{s,k}$ for $s \in \text{ReplicaID}, k = 1..w''(s)$ such that

$$\begin{aligned} (\forall s, k. 1 \leq k \leq w(s) \implies e''_{s,k} = e_{s,k}) \wedge \\ (\forall s, k. 1 \leq k \leq w'(s) \implies e''_{s,k} = e'_{s,k}) \end{aligned}$$

and

$$\begin{aligned} &(\{e''_{s,k} \mid s \in \text{ReplicaID} \wedge 1 \leq k \leq w''(s)\} = \\ &\quad \{e \in E \cup E' \mid \exists a. \text{oper}''(e) = \text{add}(a)\}) \wedge \\ &(\forall s, j, k. (\text{repl}(e''_{s,k}) = s) \wedge (e''_{s,j} \xrightarrow{\text{ro}''} e''_{s,k} \iff j < k)). \end{aligned}$$

Since

$$\forall a, s. V(a, s) \leq w(s) \wedge V'(a, s) \leq w'(s),$$

by the definition of V'' we have

$$\begin{aligned} V''(a, s) &\leq \max\{V(a, s), V'(a, s)\} \leq \\ &\quad \max\{w(s), w'(s)\} = w''(s). \end{aligned}$$

Consider a and s such that $V''(a, s) \neq 0$. Then $V''(a, s) = V(a, s)$ or $V''(a, s) = V'(a, s)$. Hence,

$$\text{oper}''(e'_{s,V''(a,s)}) = \text{oper}''(e_{s,V(a,s)}) = \text{add}(a)$$

or

$$\text{oper}''(e'_{s,V''(a,s)}) = \text{oper}''(e'_{s,V'(a,s)}) = \text{add}(a),$$

so that in any case

$$\text{oper}''(e'_{s,V''(a,s)}) = \text{add}(a).$$

We also have:

$$\begin{aligned} \forall k. (V(a, s) < k \leq w(s) \vee V'(a, s) < k \leq w'(s)) \\ \implies \text{oper}''(e''_{s,k}) \neq \text{add}(a). \end{aligned}$$

Hence,

$$\begin{aligned} \forall k. (\max\{V(a, s), V'(a, s)\} < k \leq \max\{w(s), w'(s)\}) \\ \implies \text{oper}''(e''_{s,k}) \neq \text{add}(a), \end{aligned}$$

and

$$\forall k. V''(a, s) < k \leq w''(s) \implies \text{oper}''(e''_{s,k}) \neq \text{add}(a).$$

We now show that

$$\neg \exists f \in E''. \text{oper}''(f) = \text{remove}(a) \wedge e_{s,V''(a,s)} \xrightarrow{\text{vis}''} f. \quad (7)$$

Since $V''(a, s) \neq 0$, we get

$$\begin{aligned} (V(a, s) \neq 0 \vee (V'(a, s) \neq 0 \wedge w(s) < V'(a, s))) \wedge \\ (V'(a, s) \neq 0 \vee (V(a, s) \neq 0 \wedge w'(s) < V(a, s))), \end{aligned}$$

which is equivalent to

$$\begin{aligned} (V(a, s) \neq 0 \wedge V'(a, s) \neq 0) \vee \\ (V(a, s) = 0 \wedge V'(a, s) \neq 0 \wedge w(s) < V'(a, s)) \vee \\ (V'(a, s) = 0 \wedge V(a, s) \neq 0 \wedge w'(s) < V(a, s)). \end{aligned}$$

Consider every disjunct separately. In the first case, we have

$$\begin{aligned} (\neg \exists f \in E. \text{oper}''(f) = \text{remove}(a) \wedge e_{s,V(a,s)} \xrightarrow{\text{vis}} f) \wedge \\ (\neg \exists f \in E'. \text{oper}''(f) = \text{remove}(a) \wedge e_{s,V'(a,s)} \xrightarrow{\text{vis}'} f) \end{aligned}$$

and $V''(a, s) = \max\{V(a, s), V'(a, s)\}$. If $V(a, s) = V'(a, s)$, then this straightforwardly implies (7). If $V(a, s) < V'(a, s)$, then $V''(a, s) = V'(a, s)$ and, since

$$\neg \exists k. V(a, s) < k \leq w(s) \wedge \text{oper}(e_{s,k}) = \text{add}(a),$$

we have $e_{s,V'(a,s)} \notin E$. This implies (7). The case of $V(a, s) > V'(a, s)$ is analogous.

In the case of the second disjunct, we have $V''(a, s) = V'(a, s)$. From $V'(a, s) \neq 0$ we get

$$\neg \exists f \in E'. \text{oper}''(f) = \text{remove}(a) \wedge e_{s,V'(a,s)} \xrightarrow{\text{vis}'} f,$$

and from $w(s) < V'(a, s)$ we get $e_{s,V'(a,s)} \notin E$. This implies (7). The third disjunct is handled analogously. Thus, (7) holds.

Finally, we show that

$$\begin{aligned} \forall a, s, k. e''_{s,k} \in E'' \wedge \text{oper}''(e''_{s,k}) = \text{add}(a) \implies \\ (k \leq V''(a, s) \vee \exists f \in E''. \text{oper}''(f) = \text{remove}(a) \wedge e''_{s,k} \xrightarrow{\text{vis}''} f). \quad (8) \end{aligned}$$

Take a, s, k such that $e''_{s,k} \in E''$ and $\text{oper}''(e''_{s,k}) = \text{add}(a)$ and assume that

$$\neg \exists f \in E''. \text{oper}''(f) = \text{remove}(a) \wedge e''_{s,k} \xrightarrow{\text{vis}''} f.$$

Then

$$\begin{aligned} (k \leq w(s) \implies \neg \exists f \in E. \text{oper}''(f) = \text{remove}(a) \wedge e_{s,k} \xrightarrow{\text{vis}} f), \\ (k \leq w'(s) \implies \neg \exists f \in E'. \text{oper}''(f) = \text{remove}(a) \wedge e'_{s,k} \xrightarrow{\text{vis}'} f), \end{aligned}$$

Figure 15. Original observed-remove set implementation [6, 31]

$\Sigma = \mathcal{P}(\mathbb{Z} \times \text{AlInfo}) \times \mathcal{P}(\text{AlInfo})$	
$\vec{\sigma}_0 = \lambda r. \langle \emptyset, \emptyset \rangle$	
$M = \Sigma$	
$\text{do}(\text{add}(a), \langle S, D \rangle, t) = (\langle S \cup \{(a, t)\}, D \rangle, \perp)$	
$\text{do}(\text{remove}(a), \langle S, D \rangle, t) =$ $(\langle S \setminus \{(a, t') \in S\}, D \cup \{t' \mid (a, t') \in S\} \rangle, \perp)$	
$\text{do}(\text{rd}, \langle S, D \rangle, t) = (\langle S, D \rangle, \{a \mid \exists t. (a, t) \in S\})$	
$\text{send}(\langle S, D \rangle) = (\langle S, D \rangle, \langle S, D \rangle)$	
$\text{receive}(\langle S, D \rangle, \langle S', D' \rangle) =$ $(\langle (S \cup S') \setminus \{(a, t) \mid t \in D \cup D'\} \rangle, D \cup D')$	

which implies

$$(k \leq w(s) \implies k \leq V(a, s)) \wedge (k \leq w'(s) \implies k \leq V'(a, s)).$$

Since $e''_{s,k} \in E''$, we have

$$1 \leq k \leq \max\{w(s), w'(s)\}$$

and, hence,

$$1 \leq k \leq \max\{V(a, s), V'(a, s)\}.$$

If $V''(a, s) = 0$, then either $V(a, s) = 0$ or $V'(a, s) = 0$, which contradicts $k \geq 1$. Hence, $V''(a, s) \neq 0$ and therefore

$$k \leq \max\{V(a, s), V'(a, s)\} = V''(a, s),$$

which implies (8).

Thus, we have established $\langle r, V'', w'' \rangle [\mathcal{R}_r] (I \sqcup J)$, as required.

Original state-based OR-set (orset). Fig. 15 shows an unoptimized implementation of the OR-set. Like the optimized implementation in Fig. 7, this implementation generates a unique instance identifier for each add operation, in this case using its timestamp. Unlike the optimized implementation, it does not summarize the information about instances. Instead, a state $\langle S, D \rangle$ explicitly includes the set of instances currently in the set and the set of identifiers of instance that have been removed. The implementation can be easily proved correct with respect to the the specification $\mathcal{F}_{\text{orset}}$ defined by (5) using the following simulation relation: $\forall r. \mathcal{R}_r = \mathcal{M}$, where

$$\begin{aligned} \langle S, D \rangle [\mathcal{M}] ((E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar}), \text{info}) \iff \\ (D = \{t \mid \exists e, f, a. e, f \in E \wedge \text{oper}(e) = \text{add}(a) \wedge \\ \text{info}(e) = t \wedge \text{oper}(e) = \text{remove}(a) \wedge e \xrightarrow{\text{vis}} f\}) \wedge \\ (S = \{(a, t) \mid \exists e. e \in E \wedge \text{oper}(e) = \text{add}(a) \wedge \text{info}(e) = t\} \setminus \\ \{(a, t) \mid t \in D\}). \end{aligned}$$

Multi-value register (mvr). We prove the correctness of the implementation \mathcal{D}_{mvr} in Fig. 16 with respect to the specification \mathcal{F}_{mvr} defined in §3.2. We let $v \sqsubseteq v' \iff \forall r. v(r) \leq v'(r)$ and $v \sqsubset v' \iff v \neq v' \wedge v \sqsubseteq v'$. To detect conflicts, the implementation in Fig. 16 associates each value stored in the register with a version vector [28]. In $\text{do}(\text{wr}(a))$, the new value gets a version vector that dominates all the previous ones. In receive , we take the union of every element in each input set whose version vector is not dominated by that of an element in the other input set.

The simulation relation is also shown in Fig. 16. The only interesting cases in the proof are DO for rd and RECEIVE. Consider the former. The only non-trivial obligation is to show that if

$$V [\mathcal{M}] ((E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}), \text{info}),$$

Figure 16. State-based multi-value register implementation [17, 31] and its simulation

Σ	$= \text{ReplicaID} \times \mathcal{P}(\mathbb{Z} \times (\text{ReplicaID} \rightarrow \mathbb{N}_0))$
$\vec{\sigma}_0$	$= \langle r, \emptyset \rangle$
M	$= \mathcal{P}(\mathbb{Z} \times (\text{ReplicaID} \rightarrow \mathbb{N}_0))$
$\text{do}(\text{wr}(a), \langle r, V \rangle, t) =$ $(\langle r, \{(a, (\lambda s. \text{if } s \neq r \text{ then } \max\{v(s) \mid (-, v) \in V\} \\ \text{else } \max\{v(s) \mid (-, v) \in V\} + 1)\} \rangle, \perp)$	
$\text{do}(\text{rd}, \langle r, V \rangle, t) = (\langle r, V \rangle, \{a \mid (a, -) \in V\})$	
$\text{send}(\langle r, V \rangle) = (\langle r, V \rangle, V)$	
$\text{receive}(\langle r, V \rangle, V') = \langle r, \{(a, v) \in V \mid \neg \exists (a', v') \in V'. v \sqsubseteq v'\} \\ \cup \{(a', v') \in V' \mid \neg \exists (a, v) \in V. v' \sqsubseteq v\} \rangle$	

$$\langle s, V \rangle [\mathcal{R}_r] I \iff (r = s) \wedge (V [\mathcal{M}] I)$$

$$\begin{aligned} V [\mathcal{M}] ((E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar}), \text{info}) \iff \\ (\forall (a, v), (a', v') \in V. (a = a' \wedge v = v') \vee \neg (v \sqsubseteq v')) \wedge \\ \exists \text{distinct } e_{s,k}. \exists h : V \rightarrow \text{ReplicaID}. \\ (\{e \in E \mid \exists a. \text{oper}(e) = \text{wr}(a)\} = \{e_{s,k} \mid s \in \text{ReplicaID} \wedge \\ 1 \leq k \leq \max\{v(s) \mid \exists a. (a, v) \in V\}\}) \wedge \\ (\forall s, j, k. (\text{repl}(e_{s,k}) = s) \wedge (e_{s,j} \xrightarrow{\text{ro}} e_{s,k} \iff j < k)) \wedge \\ (\forall (a, v) \in V. \text{oper}(e_{h(a,v), v(h(a,v))}) = \text{wr}(a)) \wedge \\ (\forall (a, v) \in V. \forall f \in \{e \in E \mid \exists b. \text{oper}(e) = \text{wr}(b)\}. \\ f \xrightarrow{\text{vis}} e_{h(a,v), v(h(a,v))} \iff \\ \exists s, k. f = e_{s,k} \wedge ((s \neq h(a, v) \wedge 1 \leq k \leq v(s)) \vee \\ (s = h(a, v) \wedge 1 \leq k \leq v(s) - 1)))) \end{aligned}$$

then

$$\begin{aligned} \{a \mid (a, -) \in V\} = \{a \mid \exists e \in E. \text{oper}(e) = \text{wr}(a) \wedge \\ \neg \exists f \in E. \exists a'. \text{oper}(e) = \text{wr}(a') \wedge e \xrightarrow{\text{vis}} f\}. \quad (9) \end{aligned}$$

Let us first show that

$$\forall (a, v), (a', v') \in V. \neg (e_{h(a,v), v(h(a,v))} \xrightarrow{\text{vis}} e_{h(a',v'), v'(h(a',v'))}). \quad (10)$$

Assume this is not the case: for some $(a, v), (a', v') \in V$ we have

$$e_{h(a,v), v(h(a,v))} \xrightarrow{\text{vis}} e_{h(a',v'), v'(h(a',v'))}.$$

We can assume that $(a, v) \neq (a', v')$. Then

$$v(h(a, v)) \leq v'(h(a, v)).$$

Furthermore, inv implies that vis is transitive, so that

$$\forall s, k. e_{s,k} \xrightarrow{\text{vis}} e_{h(a,v), v(h(a,v))} \implies e_{s,k} \xrightarrow{\text{vis}} e_{h(a',v'), v'(h(a',v'))},$$

which is equivalent to

$$\begin{aligned} \forall s, k. ((s \neq h(a, v) \wedge 1 \leq k \leq v(s)) \vee \\ (s = h(a, v) \wedge 1 \leq k \leq v(s) - 1)) \implies \\ ((s \neq h(a', v') \wedge 1 \leq k \leq v'(s)) \vee \\ (s = h(a', v') \wedge 1 \leq k \leq v'(s) - 1)). \end{aligned}$$

This implies

$$\forall s. s \neq h(a, v) \implies v(s) \leq v'(s).$$

Hence, $v \sqsubseteq v'$, contradicting \mathcal{M} . Therefore, (10) holds.

Consider any $f \in E$ such that $\exists b. \text{oper}(f) = \text{wr}(b)$. Then

$$\exists (a, v) \in V. \exists s, k. f = e_{s,k} \wedge 1 \leq k \leq v(s).$$

This implies

$$e_{s,k} = e_{h(a,v), v(h(a,v))} \vee e_{s,k} \xrightarrow{\text{vis}} e_{h(a,v), v(h(a,v))}.$$

Hence,

$$\forall f \in \{e \in E \mid \exists b. \text{oper}(e) = \mathbf{wr}(b)\}. \exists (a, v) \in V. \\ f = e_{h(a,v),v(h(a,v))} \vee f \xrightarrow{\text{vis}} e_{h(a,v),v(h(a,v))}. \quad (11)$$

The inv property implies that vis is transitive, so that the above and (10) imply

$$\forall (a, v) \in V. \neg \exists f \in \{e \in E \mid \exists b. \text{oper}(e) = \mathbf{wr}(b)\}. \\ e_{h(a,v),v(h(a,v))} \xrightarrow{\text{vis}} f.$$

This and (11) imply

$$\{e_{h(a,v),v(h(a,v))} \mid (a, v) \in V\} = \\ \{e \mid \exists a \in E. \text{oper}(e) = \mathbf{wr}(a) \wedge \\ \neg \exists f \in E. \exists a'. \text{oper}(e) = \mathbf{wr}(a') \wedge e \xrightarrow{\text{vis}} f\}.$$

Since

$$\forall (a, v) \in V. \text{oper}(e_{h(a,v),v(h(a,v))}) = \mathbf{wr}(a),$$

this implies (9).

Let us now discharge RECEIVE. Let $\text{receive}(\langle r, V \rangle, V') = \langle r, V'' \rangle$, where

$$V'' = \{(a, v) \in V \mid \neg \exists (a', v') \in V'. v \sqsubseteq v'\} \cup \\ \{(a', v') \in V' \mid \neg \exists (a, v) \in V. v' \sqsubseteq v\}$$

Assume $\langle r, V \rangle [\mathcal{R}_r] I, V' [\mathcal{M}] J$ and

$$I = ((E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar}), \text{info}); \\ J = ((E', \text{repl}', \text{obj}', \text{oper}', \text{rval}', \text{ro}', \text{vis}', \text{ar}'), \text{info}'); \\ I \sqcup J = ((E'', \text{repl}'', \text{obj}'', \text{oper}'', \text{rval}'', \text{ro}'', \text{vis}'', \text{ar}''), \text{info}'').$$

By agree we have $I \sqcup J \in \text{IEx}$. Then

$$(\forall (a, v), (a', v') \in V. (a = a' \wedge v = v') \vee \neg(v \sqsubseteq v')) \wedge \\ \exists \text{ distinct } e_{s,k}. \exists h : V \rightarrow \text{ReplicaID}. \\ (\{e \in E \mid \exists a. \text{oper}''(e) = \mathbf{wr}(a)\} = \{e_{s,k} \mid s \in \text{ReplicaID} \wedge \\ 1 \leq k \leq \max\{v(s) \mid \exists a. (a, v) \in V\}\}) \wedge \\ (\forall s, j, k. (\text{repl}''(e_{s,k}) = s) \wedge (e_{s,j} \xrightarrow{\text{ro}''} e_{s,k} \iff j < k)) \wedge \\ (\forall (a, v) \in V. \text{oper}''(e_{h(a,v),v(h(a,v))}) = \mathbf{wr}(a)) \wedge \\ (\forall (a, v) \in V. \forall f \in \{e \in E \mid \exists b. \text{oper}(e) = \mathbf{wr}(b)\}. \\ f \xrightarrow{\text{vis}} e_{h(a,v),v(h(a,v))} \iff \\ \exists s, k. f = e_{s,k} \wedge ((s \neq h(a, v) \wedge 1 \leq k \leq v(s)) \vee \\ (s = h(a, v) \wedge 1 \leq k \leq v(s) - 1))))$$

and

$$(\forall (a, v), (a', v') \in V'. (a = a' \wedge v = v') \vee \neg(v \sqsubseteq v')) \wedge \\ \exists \text{ distinct } e'_{s,k}. \exists h' : V' \rightarrow \text{ReplicaID}. \\ (\{e \in E' \mid \exists a. \text{oper}'(e) = \mathbf{wr}(a)\} = \{e'_{s,k} \mid s \in \text{ReplicaID} \wedge \\ 1 \leq k \leq \max\{v(s) \mid \exists a. (a, v) \in V'\}\}) \wedge \\ (\forall s, j, k. (\text{repl}'(e'_{s,k}) = s) \wedge (e'_{s,j} \xrightarrow{\text{ro}' } e'_{s,k} \iff j < k)) \wedge \\ (\forall (a, v) \in V'. \text{oper}'(e'_{h'(a,v),v(h'(a,v))}) = \mathbf{wr}(a)) \wedge \\ (\forall (a, v) \in V'. \forall f \in \{e \in E' \mid \exists b. \text{oper}(e) = \mathbf{wr}(b)\}. \\ f \xrightarrow{\text{vis}'} e'_{h'(a,v),v(h'(a,v))} \iff \\ \exists s, k. f = e'_{s,k} \wedge ((s \neq h'(a, v) \wedge 1 \leq k \leq v(s)) \vee \\ (s = h'(a, v) \wedge 1 \leq k \leq v(s) - 1))))$$

The agree property also implies

$$\forall s, k. 1 \leq k \leq \min \{ \max\{v(s) \mid \exists a. (a, v) \in V\}, \\ \max\{v(s) \mid \exists a. (a, v) \in V'\} \} \implies e_{s,k} = e'_{s,k}.$$

Hence, there exist distinct

$$e''_{s,k} \text{ for } s \in \text{ReplicaID}, k = 1..(\max\{v(s) \mid \exists a. (a, v) \in V''\}),$$

such that

$$(\forall s, k. 1 \leq k \leq \max\{v(s) \mid \exists a. (a, v) \in V\} \implies e''_{s,k} = e_{s,k}) \wedge \\ (\forall s, k. 1 \leq k \leq \max\{v(s) \mid \exists a. (a, v) \in V'\} \implies e''_{s,k} = e'_{s,k})$$

and

$$(\{e \in E \cup E' \mid \exists a. \text{oper}''(e) = \mathbf{wr}(a)\} = \\ \{e''_{s,k} \mid s \in \text{ReplicaID} \wedge 1 \leq k \leq \max\{v(s) \mid \exists a. (a, v) \in V''\}\}) \wedge \\ (\forall s, j, k. (\text{repl}(e''_{s,k}) = s) \wedge (e''_{s,j} \xrightarrow{\text{ro}''} e''_{s,k} \iff j < k)).$$

Assume $(a, v) \in V, (a', v) \in V'$ and $h(a, v) \neq h'(a', v)$. Then

$$e''_{h(a,v),v(h(a,v))} \xrightarrow{\text{vis}''} e''_{h'(a',v),v(h'(a',v))} \xrightarrow{\text{vis}''} e''_{h(a,v),v(h(a,v))},$$

contradicting $I \sqcup J \in \text{IEx}$. Hence, $h(a, v) = h'(a', v)$ and thus

$$\mathbf{wr}(a) = \text{oper}(e_{h(a,v),v(h(a,v))}) = \\ \text{oper}(e'_{h'(a',v),v(h'(a',v))}) = \mathbf{wr}(a').$$

We have thus shown that

$$(a, v) \in V, (a', v) \in V' \implies h(a, v) = h'(a', v) \wedge a = a'. \quad (12)$$

From this, the definition of V'' and

$$\forall a, v, a', v'. ((a, v), (a', v') \in V \vee (a, v), (a', v') \in V') \\ \implies ((a = a' \wedge v = v') \vee \neg(v \sqsubseteq v')),$$

we get

$$\forall (a, v), (a', v') \in V''. (a = a' \wedge v = v') \vee \neg(v \sqsubseteq v').$$

From (12) we get

$$\forall (a, v) \in V \cap V'. h(a, v) = h'(a, v).$$

Hence, there exists $h'' : V'' \rightarrow \text{ReplicaID}$ such that

$$(\forall (a, v) \in V'' \cap V. h''(a, v) = h(a, v)) \wedge \\ (\forall (a, v) \in V'' \cap V'. h''(a, v) = h'(a, v)).$$

Finally, we straightforwardly get

$$(\forall (a, v) \in V''. \forall f \in \{e \in E \cup E' \mid \exists b. \text{oper}(e) = \mathbf{wr}(b)\}. \\ f \xrightarrow{\text{vis}''} e_{h''(a,v),v(h''(a,v))} \iff \\ \exists s, k. f = e''_{s,k} \wedge ((s \neq h''(a, v) \wedge 1 \leq k \leq v(s)) \vee \\ (s = h''(a, v) \wedge 1 \leq k \leq v(s) - 1))).$$

All of the above implies $\langle r, V'' \rangle [\mathcal{R}_r] (I \sqcup J)$, as required.

Optimized multi-value register (mvr). We prove the correctness of the implementation \mathcal{D}_{mvr} in Fig. 17 with respect to the specification \mathcal{F}_{mvr} defined in §3.2. We let least upper bound operator on version vectors be: $\sqcup V = \lambda s. \max\{v(s) \mid v \in V\}$.

The original implementation of multi-value register stores a version vector with each write. If a single value was written concurrently, vectors corresponding to these writes are incomparable, i.e., they do not eliminate one another during receive. In the result, multiple vectors may be stored in the state for a single value.

The optimized implementation avoids this unnecessary overhead with a new definition of receive that coalesces version vectors for the same value. Note that in the original implementation a value a is exposed in the register's observable state, or equivalently, appears as at least one entry $(a, \cdot) \in V$, if there is at least one entry for write of a that has not been dominated by another write of $a' \neq a$. We observe this condition is equivalent to: the least upper bound of all vectors of writes of a is not dominated by the least upper bound of all vectors of writes of $a' \neq a$. Consequently, implementation stores only the least upper bound version vector per element.

The simulation relation is also shown in Fig. 17. The only interesting cases in the proof are DO for rd and RECEIVE. Consider

Figure 17. Optimized state-based multi-value register and its simulation

Σ	$= \text{ReplicaID} \times \mathcal{P}(\mathbb{Z} \times (\text{ReplicaID} \rightarrow \mathbb{N}_0))$
$\vec{\sigma}_0$	$= \langle r, \emptyset \rangle$
M	$= \mathcal{P}(\mathbb{Z} \times (\text{ReplicaID} \rightarrow \mathbb{N}_0))$
$\text{do}(\text{wr}(a), \langle r, V \rangle, t) =$	$\langle r, \{ (a, (\lambda s. \text{if } s \neq r \text{ then } \max\{v(s) \mid (-, v) \in V\} \text{ else } \max\{v(s) \mid (-, v) \in V\} + 1)) \}, \perp \rangle$
$\text{do}(\text{rd}, \langle r, V \rangle, t) =$	$\langle \langle r, V \rangle, \{a \mid (a, -) \in V\} \rangle$
$\text{send}(\langle r, V \rangle) =$	$\langle \langle r, V \rangle, V \rangle$
$\text{receive}(\langle r, V \rangle, V') =$	$\langle r, \{ (a, v) \in V'' \mid v \not\subseteq \bigsqcup \{v' \mid \exists a'. (a', v') \in V'' \wedge a \neq a'\} \} \rangle$,
where $V'' =$	$\{ (a, \bigsqcup \{v' \mid (a, v') \in V \cup V'\}) \mid (a, -) \in V \cup V' \}$
<hr/>	
$\langle s, V \rangle [\mathcal{R}_r] I \iff$	$(r = s) \wedge (V [\mathcal{M}] I)$
<hr/>	
$V [\mathcal{M}] ((E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar}), \text{info}) \iff$	$(\forall (a, v), (a', v') \in V. (a = a' \implies v = v')) \wedge$
	$(\forall (a, v) \in V. \exists s. v(s) > 0) \wedge$
	$(\forall (a, v) \in V. v \not\subseteq \bigsqcup \{v' \mid \exists a'. (a', v') \in V \wedge a \neq a'\}) \wedge$
	$\exists \text{distinct } e_{s,k}.$
	$(\{e \in E \mid \exists a. \text{oper}(e) = \text{wr}(a)\} = \{e_{s,k} \mid s \in \text{ReplicaID} \wedge$
	$1 \leq k \leq \max\{v(s) \mid \exists a. (a, v) \in V\}) \wedge$
	$(\forall s, j, k. (\text{repl}(e_{s,k}) = s) \wedge (e_{s,j} \xrightarrow{\text{ro}} e_{s,k} \iff j < k)) \wedge$
	$(\forall (a, v) \in V. \forall q. \{j \mid \text{oper}(e_{q,j}) = \text{wr}(a)\} \cup$
	$\{j \mid \exists s, k. e_{q,j} \xrightarrow{\text{vis}} e_{s,k} \wedge \text{oper}(e_{s,k}) = \text{wr}(a)\} =$
	$\{j \mid 1 \leq j \leq v(q)\} \wedge$
	$(\forall e \in E. (\text{oper}(e) = \text{wr}(a) \wedge$
	$\neg \exists f \in E. \text{oper}(f) = \text{wr}(-) \wedge e \xrightarrow{\text{vis}} f) \implies (a, -) \in V)$

the former. The only non-trivial obligation is to show that if

$$V [\mathcal{M}] ((E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}), \text{info}),$$

then

$$\{a \mid (a, -) \in V\} \subseteq \{a \mid \exists e \in E. \text{oper}(e) = \text{wr}(a) \wedge \neg \exists f \in E. \exists a'. \text{oper}(e) = \text{wr}(a') \wedge e \xrightarrow{\text{vis}} f\} \quad (13)$$

(the reverse inclusion is straightforwardly implied by \mathcal{R}_r).

Take $(a, v) \in V$. We have $\forall (a, v) \in V. \exists s. v(s) > 0$,

$$v \not\subseteq \bigsqcup \{v' \mid \exists a'. (a', v') \in V \wedge a \neq a'\}$$

and

$$\begin{aligned} & \forall (a, v) \in V. \forall q. \{j \mid \text{oper}(e_{q,j}) = \text{wr}(a)\} \cup \\ & \{j \mid \exists s, k. e_{q,j} \xrightarrow{\text{vis}} e_{s,k} \wedge \text{oper}(e_{s,k}) = \text{wr}(a)\} = \\ & \{j \mid 1 \leq j \leq v(q)\}. \end{aligned}$$

From this we get that for some $e \in E$

$$\begin{aligned} \text{oper}(e) = \text{wr}(a) \wedge \neg \exists f \in E. \exists a'. a' \neq a \wedge \\ \text{oper}(e) = \text{wr}(a') \wedge e \xrightarrow{\text{vis}} f. \end{aligned}$$

Since vis is acyclic, this implies that for some $e' \in E$

$$\text{oper}(e') = \text{wr}(a) \wedge \neg \exists f \in E. \text{oper}(e') = \text{wr}(-) \wedge e' \xrightarrow{\text{vis}} f,$$

which establishes (13).

Let us now discharge RECEIVE. Let $\text{receive}(\langle r, V \rangle, V') = \langle r, V''' \rangle$, where

$$\begin{aligned} V'' &= \{ (a, \bigsqcup \{v' \mid (a, v') \in V \cup V'\}) \mid (a, -) \in V \cup V' \}; \\ V''' &= \{ (a, v) \in V'' \mid v \not\subseteq \bigsqcup \{(a', v') \in V'' \mid a \neq a'\} \}. \end{aligned}$$

Assume $\langle r, V \rangle [\mathcal{R}_r] I, V' [\mathcal{M}] J$ and

$$\begin{aligned} I &= ((E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar}), \text{info}); \\ J &= ((E', \text{repl}', \text{obj}', \text{oper}', \text{rval}', \text{ro}', \text{vis}', \text{ar}'), \text{info}'); \\ I \sqcup J &= ((E'', \text{repl}'', \text{obj}'', \text{oper}'', \text{rval}'', \text{ro}'', \text{vis}'', \text{ar}''), \text{info}''). \end{aligned}$$

By agree we have $I \sqcup J \in \text{IEx}$. Then

$$\begin{aligned} & (\forall (a, v), (a', v') \in V. (a = a' \implies v = v')) \wedge \\ & (\forall (a, v) \in V. \exists s. v(s) > 0) \wedge \\ & (\forall (a, v) \in V. v \not\subseteq \bigsqcup \{v' \mid \exists a'. (a', v') \in V \wedge a \neq a'\}) \wedge \\ & \exists \text{distinct } e_{s,k}. \\ & (\{e \in E \mid \exists a. \text{oper}''(e) = \text{wr}(a)\} = \{e_{s,k} \mid s \in \text{ReplicaID} \wedge \\ & 1 \leq k \leq \max\{v(s) \mid \exists a. (a, v) \in V\}) \wedge \\ & (\forall s, j, k. (\text{repl}''(e_{s,k}) = s) \wedge (e_{s,j} \xrightarrow{\text{ro}} e_{s,k} \iff j < k)) \wedge \\ & (\forall (a, v) \in V. \forall q. \{j \mid \text{oper}''(e_{q,j}) = \text{wr}(a)\} \cup \\ & \{j \mid \exists s, k. e_{q,j} \xrightarrow{\text{vis}} e_{s,k} \wedge \text{oper}''(e_{s,k}) = \text{wr}(a)\} = \\ & \{j \mid 1 \leq j \leq v(q)\} \wedge \\ & (\forall e \in E. (\text{oper}''(e) = \text{wr}(a) \wedge \\ & \neg \exists f \in E. \text{oper}''(f) = \text{wr}(-) \wedge e \xrightarrow{\text{vis}} f) \implies (a, -) \in V) \end{aligned}$$

and

$$\begin{aligned} & (\forall (a, v), (a', v') \in V'. (a = a' \implies v = v')) \wedge \\ & (\forall (a, v) \in V'. \exists s. v(s) > 0) \wedge \\ & (\forall (a, v) \in V'. v \not\subseteq \bigsqcup \{v' \mid \exists a'. (a', v') \in V' \wedge a \neq a'\}) \wedge \\ & \exists \text{distinct } e_{s,k}. \\ & (\{e \in E' \mid \exists a. \text{oper}''(e) = \text{wr}(a)\} = \{e_{s,k} \mid s \in \text{ReplicaID} \wedge \\ & 1 \leq k \leq \max\{v(s) \mid \exists a. (a, v) \in V'\}) \wedge \\ & (\forall s, j, k. (\text{repl}''(e_{s,k}) = s) \wedge (e_{s,j} \xrightarrow{\text{ro}'} e_{s,k} \iff j < k)) \wedge \\ & (\forall (a, v) \in V'. \forall q. \{j \mid \text{oper}''(e_{q,j}) = \text{wr}(a)\} \cup \\ & \{j \mid \exists s, k. e_{q,j} \xrightarrow{\text{vis}'} e_{s,k} \wedge \text{oper}''(e_{s,k}) = \text{wr}(a)\} = \\ & \{j \mid 1 \leq j \leq v(q)\} \wedge \\ & (\forall e \in E'. (\text{oper}''(e) = \text{wr}(a) \wedge \\ & \neg \exists f \in E'. \text{oper}''(f) = \text{wr}(-) \wedge e \xrightarrow{\text{vis}'} f) \implies (a, -) \in V'). \end{aligned}$$

The agree property also implies

$$\forall s, k. 1 \leq k \leq \min \{ \max\{v(s) \mid \exists a. (a, v) \in V\}, \max\{v(s) \mid \exists a. (a, v) \in V'\} \} \implies e_{s,k} = e'_{s,k}.$$

Hence, there exist distinct

$$e''_{s,k} \text{ for } s \in \text{ReplicaID}, k = 1..(\max\{v(s) \mid \exists a. (a, v) \in V''\}),$$

such that

$$\begin{aligned} & (\forall s, k. 1 \leq k \leq \max\{v(s) \mid \exists a. (a, v) \in V\} \implies e''_{s,k} = e_{s,k}) \wedge \\ & (\forall s, k. 1 \leq k \leq \max\{v(s) \mid \exists a. (a, v) \in V'\} \implies e''_{s,k} = e'_{s,k}) \end{aligned}$$

and

$$\begin{aligned} & (\{e \in E \cup E' \mid \exists a. \text{oper}''(e) = \text{wr}(a)\} = \\ & \{e''_{s,k} \mid s \in \text{ReplicaID} \wedge 1 \leq k \leq \max\{v(s) \mid \exists a. (a, v) \in V''\}\}) \\ & \wedge (\forall s, j, k. (\text{repl}''(e''_{s,k}) = s) \wedge (e''_{s,j} \xrightarrow{\text{ro}''} e''_{s,k} \iff j < k)). \end{aligned}$$

By the definition of V'' and V''' we have

$$\forall (a, v), (a', v') \in V'''. (a = a' \implies v = v').$$

We also straightforwardly get

$$\forall (a, v) \in V'. \exists s. v(s) > 0$$

and

$$\begin{aligned} & (\forall (a, v) \in V'''. \forall q. \{j \mid \text{oper}''(e''_{q,j}) = \text{wr}(a)\} \cup \\ & \{j \mid \exists s, k. e''_{q,j} \xrightarrow{\text{vis}''} e''_{s,k} \wedge \text{oper}''(e''_{s,k}) = \text{wr}(a)\} = \\ & \{j \mid 1 \leq j \leq v(q)\}). \end{aligned} \quad (14)$$

Consider $e \in E \cup E'$ such that

$$\text{oper}(e) = \text{wr}(a) \wedge \neg \exists f \in E. \text{oper}(f) = \text{wr}(\cdot) \wedge e \xrightarrow{\text{vis}''} f; \quad (15)$$

then $(a, v) \in V''$ for some v . If

$$v \sqsubseteq \bigsqcup \{v' \mid \exists a'. (a', v') \in V'' \wedge a \neq a'\}.$$

then (14) yields a contradiction with (15). Hence, $(a, v) \in V'''$.

By the definition of V''' we have

$$\forall (a, v) \in V''' . v \not\sqsubseteq \bigsqcup \{v' \mid \exists a'. (a', v') \in V'' \wedge a \neq a'\}.$$

Since $V''' \subseteq V''$, this implies

$$\forall (a, v) \in V''' . v \not\sqsubseteq \bigsqcup \{v' \mid \exists a'. (a', v') \in V''' \wedge a \neq a'\}.$$

All of the above implies $\langle r, V''' \rangle [\mathcal{R}_r] (I \sqcup J)$, as required.

B. Additional material on replication-aware simulations

In this appendix, we prove the theorems and claims in §5. Recall that in that section, we fixed the following entities:

Object x of a type τ ; $D_\tau = (\Sigma, \vec{\sigma}_0, M, \text{do}, \text{send}, \text{receive})$;
Visibility witness \mathcal{V} ; Transport layer specification \mathcal{T} .

Also, we considered abstract, instrumented, decomposed and concrete executions that work only on the object x . Accordingly, we write $\text{AEx}[x]$ and $\text{CEX}[x]$ for the sets of such x -restricted abstract and concrete executions, respectively. Finally, we used the notation I_\emptyset to mean the instrumented abstract execution with the empty event set, and D_\emptyset the unique decomposed execution satisfying the condition below:

$$\text{dom}(D_\emptyset) = \text{ReplicaID} \wedge \forall r \in \text{ReplicaID}. D_\emptyset(r) = I_\emptyset.$$

B.1 Proof of Theorem 9

In this subsection, we prove the soundness of our simulation method (Theorem 9). In the proof, we will use the following notations and concepts.

1. For every partial function $f : X_1 \times \dots \times X_n \rightarrow Y$, we write

$$f(x_1, \dots, x_n) \downarrow$$

to mean that $(x_1, \dots, x_n) \in \text{dom}(f)$.

2. We lift our simulation relation from states of a single object to configurations:

$$\mathcal{S} \subseteq \text{Config} \times \text{DEx}$$

which is defined as follows:

$$\begin{aligned} (R, T) [\mathcal{S}] D &\iff \\ \text{dom}(D) &= (\{r \mid (x, r) \in \text{dom}(R)\} \cup \text{dom}(T)) \\ &\wedge (\forall m \in M. \forall \text{mid} \in \text{dom}(D) \cap \text{MessageID}. \\ &\quad T(\text{mid}) = (\cdot, \cdot, m) \implies m [\mathcal{M}] D(\text{mid})) \\ &\wedge (\forall r \in \text{dom}(D) \cap \text{ReplicaID}. R(x, r) [\mathcal{R}_r] D(r)). \end{aligned}$$

3. For a concrete execution C with event set $C.E$, we write

$$C \sim (R_0, T_0) \xrightarrow{e_1} (R_1, T_1) \xrightarrow{e_2} \dots \xrightarrow{e_n} (R_n, T_n)$$

for some $1 \leq n$, to mean that

(a) $e_1 < \dots < e_n$ is the enumeration of all elements in $C.E$ according to C 's eo order;

(b) $\text{pre}(e_i) = (R_{i-1}, T_{i-1})$ and $\text{post}(e_i) = (R_i, T_i)$ for all i .

We extend this notation to concrete executions C with countably infinite sets $C.E$, and write

$$C \sim (R_0, T_0) \xrightarrow{e_1} (R_1, T_1) \xrightarrow{e_2} \dots \xrightarrow{e_i} (R_i, T_i) \xrightarrow{e_{i+1}} \dots$$

to mean the same property except that it now refers to the infinite enumeration e_1, \dots, e_i, \dots of $C.E$ by the eo order. In both cases, for any prefix e_1, \dots, e_i of the enumeration of $C.E$ according to the eo order, we use the notation

$$C|_{e_1, \dots, e_i}$$

to mean the restriction of C to events in $\{e_1, \dots, e_i\}$. This restriction always gives a well-defined concrete execution, because e_1, \dots, e_i is a prefix of the enumeration of $C.E$ by eo.

4. We say that C **starts legitimately** if its starting configuration (R_0, T_0) satisfies the following condition:

$$\begin{aligned} \text{dom}(T_0) &= \emptyset \wedge \text{dom}(R_0) = \{(x, r) \mid r \in \text{ReplicaID}\} \\ &\wedge \forall r \in \text{ReplicaID}. R_0(x, r) = \vec{\sigma}_0(r). \end{aligned}$$

THEOREM 22. For every concrete execution $C \in \text{CEX}[x] \cap \mathcal{T}$, if C starts legitimately and

$$C \sim (R_0, T_0) \xrightarrow{e_1} (R_1, T_1) \xrightarrow{e_2} \dots \xrightarrow{e_i} (R_i, T_i) \xrightarrow{e_{i+1}} \dots$$

for some e_i 's and (R_i, T_i) 's, then there exists a sequence of decomposed executions:

$$D_0, D_1, D_2, \dots, D_i, \dots$$

such that

1. for every $i \geq 1$,

$$\text{step}(C, e_i, D_{i-1}) \downarrow \wedge \text{step}(C, e_i, D_{i-1}) = D_i;$$

2. for every $i \geq 1$, if we let $A_i = \text{abs}(C|_{e_1, \dots, e_i}, \mathcal{V})$, then

$$\forall e \in A_i.E. \forall r \in \text{ReplicaID}.$$

$$A_i.\text{repl}(e) = r \implies \text{ctxt}(D_i(r).A, e) = \text{ctxt}(A_i, e);$$

3. for every $i \geq 0$,

$$(R_i, T_i) [\mathcal{S}] D_i.$$

A similar property also holds for an concrete execution C with finite $C.E$, if C starts legitimately.

PROOF. For all $i \geq 1$, let

$$C_i = C|_{e_1, \dots, e_i}.$$

Then, for every $i \geq 1$,

$$C_i \in \text{CEX}[x] \cap \mathcal{T} \wedge C_i \sim C_{i-1} \xrightarrow{e_i} (R_i, T_i)$$

$$\wedge \forall D. \text{step}(C_i, e_i, D) = \text{step}(C, e_i, D).$$

For each $k \geq 1$ and a sequence of decomposed execution

$$D_0, D_1, D_2, \dots, D_k,$$

we write

$$\Phi(D_0, \dots, D_k)$$

if the following two conditions are met:

1. For all $1 \leq i \leq k$,

$$\text{step}(C, e_i, D_{i-1}) \downarrow \wedge \text{step}(C, e_i, D_{i-1}) = D_i.$$

2. For all $1 \leq i \leq k$,

$$C_i [\mathcal{G}] D_i.$$

Let I_\emptyset be the instrumented abstract execution with the empty event set, and define D_\emptyset as follows:

$$\text{dom}(D_\emptyset) = \text{ReplicaID} \wedge \forall r \in \text{ReplicaID}. D_\emptyset(r) = I_\emptyset.$$

Since C starts legitimately, $\text{dom}(T_0) = \emptyset$. Hence, by the G-INIT condition of the simulation method in Fig. 9,

$$\text{step}(C_1, e, D_\emptyset) \downarrow \wedge C_1 [\mathcal{G}] \text{step}(C_1, e, D_\emptyset).$$

We define

$$D_0 = D_\emptyset \quad \text{and} \quad D_1 = \text{step}(C_1, e, D_\emptyset).$$

Then,

$$\Phi(D_0, D_1) \wedge (R_0, T_0) [\mathcal{S}] D_0 \quad (16)$$

where the second conjunct comes from the INIT condition of our simulation method in Fig 5, the definition of D_0 and the fact that C starts legitimately.

The rest of the proof consists of two parts. Firstly, we prove that for all $k \geq 1$ and sequences of decomposed executions

$$D_0, D_1, D_2, \dots, D_k,$$

if $\Phi(D_0, D_1, \dots, D_k)$ holds, there exists D_{k+1} such that $\Phi(D_0, D_1, \dots, D_k, D_{k+1})$ holds as well. As we have already noticed in (16), $\Phi(D_0, D_1)$ holds. Hence, this first property guarantees the existence of an infinite sequence

$$D_0, D_1, D_2, \dots, D_i, \dots$$

that satisfies the first item of the theorem. In fact, this sequence also satisfies the second item of the theorem, because

$$\forall i \geq 1. C_i [\mathcal{G}] D_i,$$

which together with the G-CTXT condition in Fig. 9 implies the second item. Let's stop considering the implication of the existence of D_{k+1} and move on to its proof. By the assumption and the definition of C_i 's,

$$C_{k+1} \sim (C_k \xrightarrow{e_{k+1}} (R_{k+1}, T_{k+1})) \wedge C_k [\mathcal{G}] D_k.$$

Hence, from the G-STEP condition in Fig. 9 follows that

$$\text{step}(C_{k+1}, e_{k+1}, D_k) \downarrow \wedge C_{k+1} [\mathcal{G}] \text{step}(C_{k+1}, e_{k+1}, D_k).$$

Furthermore, since C_{k+1} is obtained by restricting C to its first $k+1$ events,

$$\forall D. \text{step}(C_{k+1}, e_{k+1}, D) = \text{step}(C, e_{k+1}, D).$$

Let $D_{k+1} = \text{step}(C, e_{k+1}, D_k)$. From what we just proved follows that

$$\Phi(D_0, D_1, \dots, D_k, D_{k+1}).$$

Secondly, we show that the third item of the theorem also holds for the sequence just constructed:

$$D_0, D_1, \dots, D_i, \dots$$

As we explained in the previous paragraph, the sequence already satisfies the first two items of the theorem. Hence, showing that it also satisfies the third can complete the proof of the theorem. Now we have to show that

$$\forall i \geq 0. (R_i, T_i) [\mathcal{S}] D_i.$$

Our proof is by induction on i . The base case of $i = 0$ directly follows from our observation in (16). For the inductive case ($i \geq 1$), we do the case analysis on $C.\text{act}(e_i)$.

1. The first case is $C.\text{act}(e_i) = \text{do}$. Let $r, o, a, t, \sigma, \sigma'$ be entities such that

$$(r, o, a, t) = (C.\text{repl}(e_i), C.\text{oper}(e_i), C.\text{rval}(e_i), C.\text{time}(e_i)) \\ \wedge R_{i-1}(x, r) \downarrow \wedge \sigma = R_{i-1}(x, r) \wedge (\sigma', a) = \mathcal{D}_\tau.\text{do}(o, \sigma, t).$$

These entities exist because of the definition of the transition relation. By induction hypothesis $(R_{i-1}, T_{i-1}) [\mathcal{S}] D_{i-1}$, so

$$D_{i-1}(r) \downarrow \wedge R_{i-1}(x, r) [\mathcal{R}_r] D_{i-1}(r).$$

Let $I = D_{i-1}(r)$. As we have already shown when constructing the sequence $D_0, D_1, \dots, D_i, \dots$, the application of the operation $\text{step}(C, e_i, D_{i-1})$ is defined and equals D_i . Furthermore, since $\text{act}(e_i) = \text{do}$,

$$D_i = D_{i-1}[r \mapsto \text{do}^\#(I, e_i, r, o, a, t)].$$

Hence, by the DO condition of our simulation method in Fig. 5,

$$\sigma' [\mathcal{R}_r] \text{do}^\#(I, e_i, r, o, a, t).$$

By the definition of the transition relation,

$$R_i = R_i[(x, r) \mapsto \sigma'] \wedge T_i = T_{i-1}.$$

By combining all the things that we proved together and the relationship $(R_{i-1}, T_{i-1}) [\mathcal{S}] D_{i-1}$ from the induction hypothesis, we can conclude that

$$(R_i, T_i) [\mathcal{S}] D_i,$$

as desired.

2. The next case is $C.\text{act}(e_i) = \text{send}$. Let $r, \text{mid}, \sigma, \sigma', m$ be entities such that

$$r = C.\text{repl}(e_i) \wedge \text{mid} = C.\text{msg}(e_i) \\ \wedge \sigma = R_{i-1}(x, r) \wedge (\sigma', m) = \mathcal{D}_\tau.\text{send}(\sigma).$$

These entities exist because of the definition of the transition relation. By induction hypothesis $(R_{i-1}, T_{i-1}) [\mathcal{S}] D_{i-1}$, so

$$D_{i-1}(r) \downarrow \wedge R_{i-1}(x, r) [\mathcal{R}_r] D_{i-1}(r).$$

Let $I = D_{i-1}(r)$. As we have already shown when constructing the sequence $D_0, D_1, \dots, D_i, \dots$, the application of the operation $\text{step}(C, e_i, D_{i-1})$ is defined and equals D_i . Furthermore, since $\text{act}(e_i) = \text{send}$, there exist instrumented executions $I', J \in \text{IEx}$ such that

$$\text{send}^\#(I) \downarrow \wedge \text{send}^\#(I) = (I', J) \\ \wedge D_i = D_{i-1}[r \mapsto I', \text{mid} \mapsto J].$$

Hence, by the SEND condition of our simulation method in Fig. 5,

$$\sigma' [\mathcal{R}_r] I' \wedge m [\mathcal{M}] J.$$

By the definition of our transition relation for concrete executions,

$$R_i = R_{i-1}[(x, r) \mapsto \sigma'] \wedge T_i = T_{i-1}[\text{mid} \mapsto (r, x, m)].$$

By combining all the things that we proved together and the relationship $(R_{i-1}, T_{i-1}) [\mathcal{S}] D_{i-1}$ from the induction hypothesis, we can conclude that

$$(R_i, T_i) [\mathcal{S}] D_i,$$

as desired.

3. The last case is $\text{act}(e_i) = \text{receive}$. Let $r, r', \sigma, \sigma', m$ be entities such that

$$C.\text{repl}(e) = r \wedge C.\text{srepl}(e) = r' \\ \wedge R_{i-1}(x, r) \downarrow \wedge \sigma = R_{i-1}(x, r) \wedge \sigma' = \mathcal{D}_\tau.\text{receive}(\sigma, m) \\ \wedge T_{i-1}(\text{mid}) \downarrow \wedge (r', x, m) = T_{i-1}(\text{mid}).$$

These entities exist because of the definition of the transition relation. By induction hypothesis $(R_{i-1}, T_{i-1}) [\mathcal{S}] D_{i-1}$, so

$$D_{i-1}(r) \downarrow \wedge R_{i-1}(x, r) [\mathcal{R}_r] D_{i-1}(r) \\ \wedge D_{i-1}(\text{mid}) \downarrow \wedge T_{i-1}(\text{mid}) [\mathcal{M}] D_{i-1}(\text{mid})$$

Let $I = D_{i-1}(r)$ and $J = D_{i-1}(\text{mid})$. As we have already shown when constructing the sequence $D_0, D_1, \dots, D_i, \dots$, the application of the operation $\text{step}(C, e_i, D_{i-1})$ is defined and equals D_i . Furthermore, since $\text{act}(e_i) = \text{receive}$, there exists an instrumented execution $I' \in \text{IEx}$ such that

$$\text{receive}^\#(I, J) \downarrow \wedge \text{receive}^\#(I, J) = I' \wedge D_i = D_{i-1}[r \mapsto I'].$$

Hence, by the condition RECEIVE of our simulation method in Fig. 5,

$$\sigma' [\mathcal{R}_r] I'.$$

By the definition of the transition relation for concrete executions,

$$R_i = R_{i-1}[(x, r) \mapsto \sigma'] \wedge T_i = T_{i-1}.$$

Thus, by combining all the things that we proved together and the relationship $(R_{i-1}, T_{i-1}) \llbracket \mathcal{S} \rrbracket D_{i-1}$ from the induction hypothesis, we can conclude that

$$(R_i, T_i) \llbracket \mathcal{S} \rrbracket D_i,$$

as desired. \square

PROOF OF THEOREM 9. Consider

$$\text{do}^\sharp, \text{send}^\sharp, \text{receive}^\sharp, \mathcal{V}, \mathcal{T}, \mathcal{G}$$

such that the conditions in Fig. 9 hold. Assume that \mathcal{D}_τ simulates a specification \mathcal{F}_τ with respect to these operations. Pick a concrete execution

$$C \in \llbracket \mathcal{D}_\tau \rrbracket \cap \mathcal{T}.$$

We should show that $\text{abs}(C, \mathcal{V}) \models [\tau \mapsto \mathcal{F}_\tau]$. Let

$$A = \text{abs}(C, \mathcal{V}).$$

If $A.E$ is empty, there is nothing to prove. Otherwise, C starts legitimately. We will focus on the case that $C.E$ is countably infinite, and prove that $A \models [\tau \mapsto \mathcal{F}_\tau]$. The other case of $C.E$ being finite can be proved similarly.

Let e_i 's and (R_i, T_i) 's be events and configurations such that

$$C \sim (R_0, T_0) \xrightarrow{e_1} (R_1, T_1) \xrightarrow{e_2} \dots \xrightarrow{e_n} (R_i, T_i) \dots$$

By Theorem 22, there exists a sequence of decomposed executions:

$$D_0, D_1, D_2, \dots, D_i, \dots$$

such that

1. for every $i \geq 0$,

$$\text{step}(C, e_i, D_{i-1}) \downarrow \wedge \text{step}(C, e_i, D_{i-1}) = D_i;$$
2. for every $i \geq 1$, if we let $A_i = \text{abs}(C|_{e_1, \dots, e_i}, \mathcal{V})$,

$$\forall e \in A_i.E. \forall r \in \text{ReplicaID}. A_i.\text{repl}(e) = r \implies \text{ctxt}(D_i(r).A, e) = \text{ctxt}(A_i, e);$$
3. for every $i \geq 0$,

$$(R_i, T_i) \llbracket \mathcal{S} \rrbracket D_i.$$

Also, for each $i \geq 1$, let

$$C_i = C|_{e_1, \dots, e_i}.$$

Then, $C_i \in \text{CEx}[x] \cap \mathcal{T}$ and for every $e \in \text{abs}(C_i, \mathcal{V}).E$,

$$\text{ctxt}(\text{abs}(C_i, e), e) = \text{ctxt}(\text{abs}(C), e) = \text{ctxt}(A, e).$$

Using what we have just prepared, we prove the theorem. Pick $e \in A.E$. We will prove that

$$\mathcal{F}_\tau(\text{ctxt}(A, e)) = A.\text{rval}(e).$$

Since $\text{abs}(C, \mathcal{V}) = A$,

$$\exists i \geq 1. e_i = e.$$

The $\text{abs}(C, \mathcal{V})$ operation includes only a do-type event. Hence,

$$C.\text{act}(e_i) = \text{do},$$

Let

$$(r, o, a, t) = (C.\text{repl}(e_i), C.\text{oper}(e_i), C.\text{rval}(e_i), C.\text{time}(e_i))$$

Then, by the definition of the transition relation, there exist σ, σ' such that

$$\begin{aligned} R_{i-1}(x, r) \downarrow \wedge \sigma &= R_{i-1}(x, r) \\ \wedge (\sigma', a) &= \mathcal{D}_\tau.\text{do}(o, \sigma, t) \wedge R_i = R_{i-1}[(x, r) \mapsto \sigma']. \end{aligned}$$

Since $\text{step}(C, e_i, D_{i-1})$ is defined and equals D_i , there exist instrumented executions I, J such that

$$\begin{aligned} D_{i-1}(r) \downarrow \wedge D_{i-1}(r) &= I \wedge \text{do}^\sharp(I, e_i, r, o, a, t) \downarrow \\ \wedge \text{do}^\sharp(I, e_i, r, o, a, t) &= J \wedge D_i(r) = J. \end{aligned}$$

Meanwhile, because of $(R_{i-1}, T_{i-1}) \llbracket \mathcal{S} \rrbracket D_{i-1}$,

$$\sigma = R_{i-1}(x, r) \llbracket \mathcal{R}_r \rrbracket D_{i-1}(r) = I.$$

Now the DO condition of our simulation method in Fig. 5 ensures that

$$a = \mathcal{F}_\tau(\text{ctxt}(J.A, e_i)). \quad (17)$$

Since $C.\text{rval}(e') = A.\text{rval}(e')$ for every $e' \in A.E$,

$$a = C.\text{rval}(e_i) = A.\text{rval}(e). \quad (18)$$

Furthermore,

$$\begin{aligned} \text{ctxt}(J.A, e_i) &= \text{ctxt}(D_i(r).A, e) \\ &= \text{ctxt}(\text{abs}(C_i, \mathcal{V}), e) = \text{ctxt}(A, e). \end{aligned} \quad (19)$$

By combining (17), (18) and (19), we can derive

$$\mathcal{F}_\tau(\text{ctxt}(A, e)) = A.\text{rval}(e),$$

as desired. \square

B.2 Proof of Theorem 10

For any abstract, concrete or instrumented execution X , set of events E and replica id r , we write $X|_E$ for the restriction of X to events in E , and $X|_r$ for the restriction of X to events in the replica r . The $\mathcal{G}^{\text{state}}$ relation for state-based implementations in §5.4 is defined as follows:

$$C \llbracket \mathcal{G}^{\text{state}} \rrbracket D$$

if and only if

$$\begin{aligned} &\text{ReplicaID} \subseteq \text{dom}(D) \\ &\wedge \text{dom}(D) \cap \text{MessageID} = \{C.\text{msg}(e) \mid C.\text{act}(e) = \text{send}\} \\ &\wedge (\forall i \in \text{dom}(D). \text{inv}(D(i))) \\ &\wedge (\forall i \in \text{dom}(D). D(i).A = \text{abs}(C, \mathcal{V}^{\text{state}})|_{D(i).E}) \\ &\wedge (\forall i \in \text{dom}(D). \forall e \in D(i).E. \\ &\quad e \in C.E \wedge D(i).\text{info}(e) = C.\text{time}(e)) \\ &\wedge (\forall r \in \text{dom}(D) \cap \text{ReplicaID}. \\ &\quad D(r).E = (((\text{ro}(C) \cup \text{del}(C))^*)^{-1}(C|_r.E))|_{\text{do}}) \\ &\wedge (\forall \text{mid} \in \text{dom}(D) \cap \text{MessageID}. \forall E_0. \\ &\quad E_0 = \{e \in C.E \mid C.\text{msg}(e) = \text{mid} \wedge C.\text{act}(e) = \text{send}\} \\ &\quad \implies D(\text{mid}).E = (((\text{ro}(C) \cup \text{del}(C))^*)^{-1}E_0)|_{\text{do}}) \\ &\wedge (\forall i, j \in \text{dom}(D). i \neq j \implies \text{agree}(D(i), D(j))) \end{aligned}$$

Here $\mathcal{V}^{\text{state}}$ is the visibility witness for the state-based implementations in §4, and inv and agree are predicates in Fig. 6.

THEOREM 10 *There exists \mathcal{G} such that for all \mathcal{D}_τ , the parameters in Fig. 6 satisfy the conditions in Fig. 9 with respect to this \mathcal{G} , $\mathcal{V} = \mathcal{V}^{\text{state}}$, $\mathcal{T} = \text{T-Any}$.*

PROOF. Let $\mathcal{T} = \text{T-Any}$. We prove each of the four conditions separately.

Condition G-CTXT. Assume that

$$C \llbracket \mathcal{G}^{\text{state}} \rrbracket D.$$

Pick e, r such that

$$e \in \text{abs}(C, \mathcal{V}^{\text{state}}).E \wedge \text{abs}(C, \mathcal{V}^{\text{state}}).\text{repl}(e) = r.$$

Since $C \llbracket \mathcal{G}^{\text{state}} \rrbracket D$, we have $r \in \text{dom}(D)$. Let

$$I = D(r).$$

Also from $C \llbracket \mathcal{G}^{\text{state}} \rrbracket D$ follows that

$$\begin{aligned} I.A &= \text{abs}(C, \mathcal{V}^{\text{state}})|_{I.E} \wedge \\ &(\text{abs}(C, \mathcal{V}^{\text{state}})|_{I.E}).\text{vis} = \text{abs}(C, \mathcal{V}^{\text{state}}).\text{vis} \cap (\text{Event} \times I.E). \end{aligned}$$

Furthermore, $e \in (C|_r.E)|_{\text{do}} \subseteq I.E$. Thus,

$$\text{ctxt}(I.A, e) = \text{ctxt}(\text{abs}(C, \mathcal{V}^{\text{state}}), e).$$

Condition G-STEP. Consider

$$C, C' \in \text{CEx}[x] \cap \mathcal{T}, e \in \text{Event}, (R, T) \in \text{Config}, D \in \text{DEx}$$

such that

$$C' \sim C \xrightarrow{e} (R, T) \wedge C \llbracket \mathcal{G}^{\text{state}} \rrbracket D.$$

We should prove that

$$\text{step}(C', e, D) \downarrow \wedge C' \llbracket \mathcal{G}^{\text{state}} \rrbracket \text{step}(C', e, D). \quad (20)$$

We do the case analysis on $C'.\text{act}(e)$.

1. The first case is that $C'.\text{act}(e) = \text{do}$. Let

$$(r, o, a, t) = (C'.\text{repl}(e), C'.\text{oper}(e), C'.\text{rval}(e), C'.\text{time}(e)).$$

Since $C \llbracket \mathcal{G}^{\text{state}} \rrbracket D$, $r \in \text{ReplicaID}$ and $e \notin C.E$,

$$r \in \text{dom}(D) \wedge e \notin D(r).E.$$

Define I and J as follows:

$$\begin{aligned} I &= D(r) \\ \text{ro}' &= I.\text{ro} \cup \{(f, e) \mid f \in I.E \wedge I.\text{repl}(f) = r\} \\ \text{vis}' &= I.\text{vis} \cup \{(f, e) \mid f \in I.E\}, \\ \text{ar}' &= I.\text{ar} \cup \{(f, e) \mid f \in I.E \wedge I.\text{info}(f) < t\} \\ &\quad \cup \{(e, f) \mid f \in I.E \wedge t < I.\text{info}(f)\}. \\ J &= ((I.E \uplus \{e\}, I.\text{repl}[e \mapsto r], I.\text{obj}[e \mapsto x], \\ &\quad I.\text{oper}[e \mapsto o], I.\text{rval}[e \mapsto a], \text{ro}', \text{vis}', \text{ar}'), \\ &\quad I.\text{info}[e \mapsto t]). \end{aligned}$$

Then, J is an instrumented execution, because I is an instrumented execution, $e \notin I.E$, $C'.\text{time}(e') = I.\text{info}(e')$ for every $e' \in I.E$ and $C'.\text{time}$ is injective. Furthermore,

$$\text{inv}(J).$$

Since $C \llbracket \mathcal{G}^{\text{state}} \rrbracket D$ and $D(r) = I$, we have $\text{inv}(I)$. Also, $e \notin I.E$. Hence, by the definitions of do^\sharp and J ,

$$\text{do}^\sharp(I, e, r, o, a, t) \downarrow \wedge \text{do}^\sharp(I, e, r, o, a, t) = J.$$

This in turn implies that

$$\text{step}(C', e, D) \downarrow \wedge \text{step}(C', e, D) = D[r \mapsto J].$$

Let $D' = D[r \mapsto J]$. Then,

$$\begin{aligned} \text{dom}(D) &= \text{dom}(D') \\ \wedge (\forall i \in \text{dom}(D')). D'(i) &= D(i) \vee D'(i) = J). \end{aligned}$$

Hence, the following conjuncts in the definition of $C' \llbracket \mathcal{G}^{\text{state}} \rrbracket D'$ hold.

$$\begin{aligned} \text{ReplicaID} &\subseteq \text{dom}(D') \\ \wedge \text{dom}(D') \cap \text{MessageID} &= \{C'.\text{msg}(e) \mid C'.\text{act}(e) = \text{send}\} \\ \wedge (\forall i \in \text{dom}(D')). \text{inv}(D'(i)) \\ \wedge (\forall i \in \text{dom}(D')). \forall f \in D'(i).E. \\ &f \in C'.E \wedge D'(i).\text{info}(f) = C'.\text{info}(f)). \end{aligned}$$

We also have

$$\forall i \in \text{dom}(D). D'(i).A = \text{abs}(C', \mathcal{V}^{\text{state}})|_{D'(i).E}$$

because, by $C \llbracket \mathcal{G}^{\text{state}} \rrbracket D$,

$$I.E = (((\text{ro}(C) \cup \text{del}(C))^*)^{-1}(C|_r.E))|_{\text{do}}.$$

By the same reason and the assumption $C \llbracket \mathcal{G}^{\text{state}} \rrbracket D$,

$$\forall r' \in \text{dom}(D') \cap \text{ReplicaID}.$$

$$D'(r').E = (((\text{ro}(C') \cup \text{del}(C'))^*)^{-1}(C'|_{r'}.E))|_{\text{do}}.$$

We do not change anything regarding messages when defining D' , so

$$\begin{aligned} \text{dom}(D) \cap \text{MessageID} &= \text{dom}(D') \cap \text{MessageID} \\ \wedge \forall \text{mid} \in \text{dom}(D) \cap \text{MessageID}. D(\text{mid}) &= D'(\text{mid}). \end{aligned}$$

Thus, from the assumption $C \llbracket \mathcal{G}^{\text{state}} \rrbracket D$ it follows that

$$\forall \text{mid} \in \text{dom}(D') \cap \text{MessageID}. \forall E_0.$$

$$\begin{aligned} E_0 = \{e' \in C'.E \mid C'.\text{msg}(e') = \text{mid} \wedge C'.\text{act}(e') = \text{send}\} \\ \implies D'(\text{mid}).E = (((\text{ro}(C') \cup \text{del}(C'))^*)^{-1}E_0)|_{\text{do}}. \end{aligned}$$

It remains to show that

$$\forall i, j \in \text{dom}(D'). i \neq j \implies \text{agree}(D'(i), D'(j)).$$

Pick $i, j \in \text{dom}(D')$ such that $i \neq j$. Since $C \llbracket \mathcal{G}^{\text{state}} \rrbracket D$ and $D(i') = D'(i')$ for all $i' \neq r$, it is sufficient to consider only the case that

$$i = r \wedge j \neq i.$$

Let $I' = D'(j)$, which should be the same as $D(j)$. We should show

$$\text{agree}(J, I').$$

First, we note that $J \sqcup I'$ is a well-defined instrumented execution. Intuitively, this is because both are restrictions of $\text{abs}(C', \mathcal{V}^{\text{state}})$ and their timestamps are given by $C'.\text{time}$. The formal versions of these properties of J and I' are given below:

$$\begin{aligned} (\forall f \in J.\text{info}(f). f \in C'.E \wedge J.\text{info}(f) &= C'.\text{time}(f)) \\ \wedge (\forall f \in I'.\text{info}(f). f \in C'.E \wedge I'.\text{info}(f) &= C'.\text{time}(f)) \\ \wedge J.A &= \text{abs}(C', \mathcal{V}^{\text{state}})|_{J.E} \\ \wedge I'.A &= \text{abs}(C', \mathcal{V}^{\text{state}})|_{I'.E}. \end{aligned}$$

Next, regarding the prefix condition in $\text{agree}(J, I')$, we only need to show that $(I'|_r.E, I'.\text{ro})$ is a prefix of $(J|_r.E, J.\text{ro})$, because

$$J.E = I.E \uplus \{e\} \wedge \text{agree}(I, I').$$

This prefix property follows from the fact that $I'.\text{ro}$ and $J.\text{ro}$ are restrictions of $C'.\text{ro}$ to $I'.E$ and $J.E$, and that

$$\begin{aligned} I'|_r.E &\subseteq \text{abs}(C, \mathcal{V}^{\text{state}})|_{r.E} \\ &\subseteq \text{abs}(C', \mathcal{V}^{\text{state}})|_{r.E} \\ &= J|_r.E. \end{aligned}$$

2. The next case is that $C'.\text{act}(e) = \text{send}$. Let

$$(r, \text{mid}) = (C'.\text{repl}(e), C'.\text{msg}(e)).$$

By the definitions of the transition relation and the concrete execution,

$$C'.\text{msg}(e) \notin \{C'.\text{msg}(f) \mid f \in C.E \wedge C'.\text{act}(f) = \text{send}\}.$$

Hence, from our assumption $C \llbracket \mathcal{G}^{\text{state}} \rrbracket D$, it follows that

$$r \in \text{dom}(D) \wedge \text{mid} \notin \text{dom}(D).$$

Let $I = D(r)$. Then, $\text{inv}(I)$ holds, because $C \llbracket \mathcal{G}^{\text{state}} \rrbracket D$. Thus,

$$\text{send}^\sharp(I) \downarrow \wedge \text{send}^\sharp(I) = (I, I).$$

This in turn implies that

$$\text{step}(C', e, D) \downarrow \wedge \text{step}(C', e, D) = D[r \mapsto I, \text{mid} \mapsto I].$$

Let $D' = D[r \mapsto I, mid \mapsto I]$. We should show that

$$C' [\mathcal{G}^{\text{state}}] D'$$

The following conjuncts in the definition of $C' [\mathcal{G}^{\text{state}}] D'$ follow immediately from the assumption $C [\mathcal{G}^{\text{state}}] D$:

$$\begin{aligned} \text{ReplicaID} &\subseteq \text{dom}(D') \\ \wedge \text{dom}(D') \cap \text{MessageID} &= \{C'.\text{msg}(e) \mid C'.\text{act}(e) = \text{send}\} \\ \wedge (\forall i \in \text{dom}(D'). \text{inv}(D'(i))). \end{aligned}$$

For the next conjunct in the definition of $C' [\mathcal{G}^{\text{state}}] D'$, we notice that $\text{abs}(C, \mathcal{V}^{\text{state}}) = \text{abs}(C', \mathcal{V}^{\text{state}})$ and that for every $i \in \text{dom}(D')$, there exists some $j \in \text{dom}(D)$ satisfying $D'(i) = D(j)$. Hence, for every $i \in \text{dom}(D')$, if we let j be such j ,

$$\begin{aligned} D'(i).A &= D(j).A = \text{abs}(C, \mathcal{V}^{\text{state}})|_{D(j).E} \\ &= \text{abs}(C', \mathcal{V}^{\text{state}})|_{D(j).E} \\ &= \text{abs}(C', \mathcal{V}^{\text{state}})|_{D'(i).E}. \end{aligned}$$

For the remaining conjuncts, we notice that all the instrumented executions in D' are already present in D , C' is an extension of C with e , and e is the send operation and happens at the end according to the $C'.\text{eo}$ order. From these properties and the assumption $C [\mathcal{G}^{\text{state}}] D$ follow all the remaining conjuncts:

$$\begin{aligned} (\forall i \in \text{dom}(D'). \forall e \in D'(i).E. \\ e \in C'.E \wedge D'(i).\text{info}(e) = C'.\text{time}(e)) \\ \wedge (\forall r \in \text{dom}(D') \cap \text{ReplicaID}. \\ D'(r).E = (((\text{ro}(C') \cup \text{del}(C'))^*)^{-1}(C'|_r.E))|_{\text{do}}) \\ \wedge (\forall mid' \in \text{dom}(D') \cap \text{MessageID}. \forall E_0. \\ E_0 = \{e \in C'.E \mid C'.\text{msg}(e) = mid' \wedge C'.\text{act}(e) = \text{send}\} \\ \implies D'(mid').E = (((\text{ro}(C') \cup \text{del}(C'))^*)^{-1}E_0)|_{\text{do}}) \\ \wedge (\forall i, j \in \text{dom}(D'). i \neq j \implies \text{agree}(D'(i), D'(j))). \end{aligned}$$

3. The last case is that $C'.\text{act}(e) = \text{receive}$. Let

$$r = C'.\text{repl}(e), \text{mid} = C'.\text{msg}(e) \text{ and } (R', T') = C'.\text{pre}(e).$$

By the definition of the transition relation, $mid \in \text{dom}(T')$. Then, since C' is a concrete execution starting from a configuration with the empty T component and $C' \sim C \xrightarrow{e} (R, T)$,

$$mid \in \{C.\text{msg}(f) \mid f \in C.E \wedge C.\text{act}(f) = \text{send}\}.$$

Thus, from our assumption $C [\mathcal{G}^{\text{state}}] D$, it follows that

$$r \in \text{dom}(D) \wedge mid \in \text{dom}(D).$$

Let

$$I_1 = D(r) \text{ and } I_2 = D(mid).$$

Since $C [\mathcal{G}^{\text{state}}] D$ and $r \neq mid$,

$$\text{inv}(I_1) \wedge \text{inv}(I_2) \wedge \text{agree}(I_1, I_2).$$

Hence, $\text{receive}^\sharp(I_1, I_2) \downarrow$. Let

$$J = \text{receive}^\sharp(I_1, I_2) \text{ and } D' = D[r \mapsto J].$$

It remains to show that $C' [\mathcal{G}^{\text{state}}] D'$. Since $\text{dom}(D') = \text{dom}(C')$ and $C'.\text{act}(e) = \text{receive}$, the assumption $C [\mathcal{G}^{\text{state}}] D$ gives

$$\begin{aligned} \text{ReplicaID} &\subseteq \text{dom}(D') \\ \wedge \text{dom}(D') \cap \text{MessageID} &= \{C'.\text{msg}(f) \mid C'.\text{act}(f) = \text{send}\} \end{aligned}$$

The next requirements in the definition of $C' [\mathcal{G}^{\text{state}}] D'$ are:

$$\begin{aligned} (\forall i \in \text{dom}(D'). \text{inv}(D'(i))) \\ \wedge (\forall i \in \text{dom}(D'). D'(i).A = \text{abs}(C', \mathcal{V}^{\text{state}})|_{D'(i).E}). \end{aligned}$$

Since $C [\mathcal{G}^{\text{state}}] D$, we just need to show that

$$\text{inv}(J) \wedge (J.A = \text{abs}(C', \mathcal{V}^{\text{state}})|_{J.E}).$$

The part of $\text{inv}(J)$ relating $J.\text{ar}$ and $J.\text{info}$ follows from $\text{inv}(I_1)$ and $\text{inv}(I_2)$ and the definition of J . For the other part of $\text{inv}(J)$ and $J.A = \text{abs}(C', \mathcal{V}^{\text{state}})|_{J.E}$, we will show the following sufficient condition:

$$\begin{aligned} (J.\text{ro} = \text{abs}(C', \mathcal{V}^{\text{state}}).\text{ro} \cap (\text{Event} \times J.E)) \\ \wedge (J.\text{vis} = \text{abs}(C', \mathcal{V}^{\text{state}}).\text{vis} \cap (\text{Event} \times J.E)) \\ \wedge (J.\text{ar} = \text{abs}(C', \mathcal{V}^{\text{state}}).\text{ar} \cap (\text{Event} \times J.E)). \end{aligned}$$

To see why they imply the second part of $\text{inv}(J)$, look at the below derivation:

$$\begin{aligned} (J.\text{vis} \cup J.\text{ro})^+ \\ = ((\text{abs}(C', \mathcal{V}^{\text{state}}).\text{vis} \cap (\text{Event} \times J.E)) \\ \cup (\text{abs}(C', \mathcal{V}^{\text{state}}).\text{ro} \cap (\text{Event} \times J.E)))^+ \\ = ((\text{abs}(C', \mathcal{V}^{\text{state}}).\text{vis} \cup \text{abs}(C', \mathcal{V}^{\text{state}}).\text{ro}) \cap (\text{Event} \times J.E))^+ \\ \subseteq (\text{abs}(C', \mathcal{V}^{\text{state}}).\text{vis} \cap (\text{Event} \times J.E))^+ \\ \subseteq (\text{abs}(C', \mathcal{V}^{\text{state}}).\text{vis})^+ \cap (\text{Event} \times J.E) \\ = J.\text{vis}. \end{aligned}$$

The proofs of the three conjuncts in the sufficient condition above are essentially the same. We present the proof of the first:

$$\begin{aligned} J.\text{ro} &= I_1.\text{ro} \cup I_2.\text{ro} \\ &= (\text{abs}(C, \mathcal{V}^{\text{state}}).\text{ro} \cap (\text{Event} \times I_1.E)) \\ &\quad \cup (\text{abs}(C, \mathcal{V}^{\text{state}}).\text{ro} \cap (\text{Event} \times I_2.E)) \\ &= \text{abs}(C, \mathcal{V}^{\text{state}}).\text{ro} \cap ((\text{Event} \times I_1.E) \cup (\text{Event} \times I_2.E)) \\ &= \text{abs}(C, \mathcal{V}^{\text{state}}).\text{ro} \cap (\text{Event} \times (I_1.E \cup I_2.E)) \\ &= \text{abs}(C, \mathcal{V}^{\text{state}}).\text{ro} \cap (\text{Event} \times J.E) \\ &= \text{abs}(C', \mathcal{V}^{\text{state}}).\text{ro} \cap (\text{Event} \times J.E). \end{aligned}$$

We move on to the following conditions of $C' [\mathcal{G}^{\text{state}}] D'$:

$$\begin{aligned} (\forall i \in \text{dom}(D'). \forall f \in D'(i).E. \\ f \in C'.E \wedge D'(i).\text{info}(f) = C'.\text{time}(f)) \\ \wedge (\forall r' \in \text{dom}(D') \cap \text{ReplicaID}. \\ D'(r').E = (((\text{ro}(C') \cup \text{del}(C'))^*)^{-1}(C'|_{r'.E}))|_{\text{do}}) \\ \wedge (\forall mid' \in \text{dom}(D') \cap \text{MessageID}. \forall E_0. \\ E_0 = \{f \in C'.E \mid C'.\text{msg}(f) = mid' \wedge C'.\text{act}(f) = \text{send}\} \\ \implies D'(mid').E = (((\text{ro}(C') \cup \text{del}(C'))^*)^{-1}E_0)|_{\text{do}}). \end{aligned}$$

Proving all the three conditions uses the assumption $C [\mathcal{G}^{\text{state}}] D$, which we will not mention explicitly. The first condition holds because $J = \text{receive}^\sharp(I_1, I_2)$ and I_1 and I_2 already satisfy this condition for $i = r$ and $i = mid$, respectively. The third condition holds because D' and D are the same when restricted to the subdomain MessageID and the event e is the receive event that is maximal according to both $\text{ro}(C')$ and $\text{del}(C')$. For the second condition, since $C'.\text{act}(e) = \text{receive}$ and e is maximal in $\text{ro}(C')$ and $\text{del}(C')$, it is sufficient to show that

$$J.E = (((\text{ro}(C') \cup \text{del}(C'))^*)^{-1}(C'|_r.E))|_{\text{do}}. \quad (21)$$

Let

$$E_0 = \{f \in C.E \mid C.\text{msg}(f) = mid \wedge C.\text{act}(f) = \text{send}\}.$$

We prove the equality in (21) as follows:

$$\begin{aligned}
J.E &= I_1.E \cup I_2.E \\
&= (((\text{ro}(C) \cup \text{del}(C))^*)^{-1}(C|_r.E))|_{\text{do}} \\
&\quad \cup (((\text{ro}(C) \cup \text{del}(C))^*)^{-1}(E_0))|_{\text{do}} \\
&= (((\text{ro}(C) \cup \text{del}(C))^*)^{-1}(C|_r.E \cup E_0))|_{\text{do}} \\
&= (((\text{ro}(C') \cup \text{del}(C'))^*)^{-1}(C'|_r.E))|_{\text{do}}.
\end{aligned}$$

The remaining condition in the definition of $C' [\mathcal{G}^{\text{state}}] D'$ is:

$$\forall i, j \in \text{dom}(D). i \neq j \implies \text{agree}(D'(i), D'(j)).$$

Since D and D' are identical except at r and the same property as the one above already holds for D , it is sufficient to prove the following property instead:

$$\forall i \in \text{dom}(D). i \neq r \implies \text{agree}(D(i), J).$$

Pick $i \in \text{dom}(D)$ such that $i \neq r$. Let $I = D(i)$. Consider any replica id r' . Then,

- (a) $(I_1|_r.E, I_1|_r.\text{ro})$ is a prefix of $(I|_r.E, I|_r.\text{ro})$ or vice versa; and
- (b) $(I_2|_r.E, I_2|_r.\text{ro})$ is a prefix of $(I|_r.E, I|_r.\text{ro})$ or vice versa; and
- (c) $(I_1|_r.E, I_1|_r.\text{ro})$ is a prefix of $(I_2|_r.E, I_2|_r.\text{ro})$ or vice versa.

By the definition of J ,

$$\begin{aligned}
(J|_r.E, J|_r.\text{ro}) &= (I_1|_r.E, I_1|_r.\text{ro}) \\
\vee (J|_r.E, J|_r.\text{ro}) &= (I_2|_r.E, I_2|_r.\text{ro}).
\end{aligned}$$

From what we have just proved follows that $(J|_r.E, J|_r.\text{ro})$ is a prefix of $(I|_r.E, I|_r.\text{ro})$ or vice versa. It remains to prove that $I \sqcup J$ is a well-defined instrumented execution. Because of the conditions of $C' [\mathcal{G}^{\text{state}}] D'$ that we have proved and the assumption $C [\mathcal{G}^{\text{state}}] D$, we know that

$$\begin{aligned}
I.A &= \text{abs}(C', \mathcal{V}^{\text{state}})|_{I.E} \\
\wedge J.A &= \text{abs}(C', \mathcal{V}^{\text{state}})|_{J.E} \\
\wedge (\forall e \in I.E. I.\text{info}(e) &= C'.\text{time}(e)) \\
\wedge (\forall e \in J.E. J.\text{info}(e) &= C'.\text{time}(e)).
\end{aligned}$$

The last two conjuncts imply that

$$\forall e \in (I \sqcup J).E. (I \sqcup J).\text{info}(e) = C'.\text{time}(e).$$

This derived fact and the other two conjuncts give that

$$(I \sqcup J).A = \text{abs}(C', \mathcal{V}^{\text{state}})|_{(I \sqcup J).E}.$$

Hence, $I \sqcup J$ is a well-defined instrumented execution.

Condition G-INIT. Consider C, e, R, T such that

$$C.E = \{e\} \wedge E.\text{pre}(e) = (R, T) \wedge \text{dom}(T) = \emptyset.$$

We have to prove that

$$\text{step}(C, e, D_\emptyset) \downarrow \wedge C [\mathcal{G}^{\text{state}}] \text{step}(C, e, D_\emptyset).$$

By the definition of the transition relation,

$$C.\text{act}(e) = \text{do} \vee C.\text{act}(e) = \text{send}.$$

We do the case analysis on these two disjuncts.

First, assume that $C.\text{act}(e) = \text{do}$. Let

$$(r, o, a, t) = (C.\text{repl}(e), C.\text{oper}(e), C.\text{rval}(e), C.\text{time}(e)).$$

Also, define I and J as follows:

$$I = D_\emptyset(r)$$

$$J = ((\{e\}, [e \mapsto r], [e \mapsto x], [e \mapsto o], [e \mapsto a], \emptyset, \emptyset, \emptyset), [e \mapsto t]).$$

Then, $I.E = \emptyset$, so $\text{inv}(I)$ holds and $e \notin I$. Also, $J \in \text{IEx}$. Hence,

$$\text{do}^\sharp(I, e, r, o, a, t) \downarrow \wedge \text{do}^\sharp(I, e, r, o, a, t) = J.$$

This in turn implies that

$$\text{step}(C, e, D_\emptyset) \downarrow \wedge \text{step}(C, e, D_\emptyset) = D_\emptyset[r \mapsto J].$$

Let $D' = D_\emptyset[r \mapsto J]$. We will now show that

$$C [\mathcal{G}^{\text{state}}] D'.$$

Since $\text{dom}(D') = \text{ReplicaID}$, $C.E = \{e\}$ and $C.\text{act}(e) = \text{do}$, the first two conditions of $C [\mathcal{G}^{\text{state}}] D'$ hold:

$$\text{ReplicaID} \subseteq \text{dom}(D')$$

$$\wedge \text{dom}(D') \cap \text{MessageID} = \{C.\text{msg}(f) \mid C.\text{act}(f) = \text{send}\}$$

For every $i \in \text{dom}(D')$, $D'(i)$ is I_\emptyset (i.e., an instrumented execution with the empty event set) or J . In particular, $D'(i).E$ is empty or the singleton set $\{e\}$. Hence,

$$(\forall i \in \text{dom}(D'). \text{inv}(D'(i)))$$

$$\wedge (\forall i \in \text{dom}(D'). D'(i).A = \text{abs}(C, \mathcal{V}^{\text{state}})|_{D'(i).E})$$

$$\wedge (\forall i \in \text{dom}(D'). \forall f \in D'(i).E.$$

$$f \in C.E \wedge D'(i).\text{info}(f) = C.\text{time}(f))$$

$$\wedge (\forall r' \in \text{dom}(D') \cap \text{ReplicaID}.$$

$$D'(r').E = (((\text{ro}(C) \cup \text{del}(C))^*)^{-1}(C|_{r'}.E))|_{\text{do}}).$$

Since there are no $\text{mid} \in \text{dom}(D')$,

$\forall \text{mid} \in \text{dom}(D') \cap \text{MessageID}. \forall E_0.$

$$E_0 = \{f \in C.E \mid C.\text{msg}(f) = \text{mid} \wedge C.\text{act}(f) = \text{send}\}$$

$$\implies D'(\text{mid}).E = (((\text{ro}(C) \cup \text{del}(C))^*)^{-1}E_0)|_{\text{do}}.$$

The remaining condition of $C [\mathcal{G}^{\text{state}}] D'$ is:

$$\forall i, j \in \text{dom}(D'). i \neq j \implies \text{agree}(D'(i), D'(j)).$$

For all $i, j \in \text{dom}(D')$ such that $i \neq j$,

$$D'(i).E = \emptyset \vee D'(j).E = \emptyset.$$

Thus, the part of $\text{agree}(D'(i), D'(j))$ involving prefixes holds. Also, $D'(i) \sqcup D'(j)$ is one of $D'(i)$ or $D'(j)$, and it is a well-defined instrumented execution. We have just shown the last condition of $C [\mathcal{G}^{\text{state}}] D'$.

Next, assume that $C.\text{act}(e) = \text{send}$. Let

$$(r, \text{mid}) = (C.\text{repl}(e), C.\text{msg}(e)).$$

Since $\text{dom}(D_\emptyset) = \text{ReplicaID}$,

$$\text{mid} \notin \text{dom}(D_\emptyset) \wedge r \in \text{dom}(D_\emptyset).$$

Let $I = D_\emptyset(r)$. Then, $I.E = \emptyset$ by the definition of D_\emptyset . Thus, $\text{inv}(I)$ holds, from which follows that

$$\text{send}^\sharp(I) \downarrow \wedge \text{send}^\sharp(I) = (I, I).$$

Since $\text{mid} \notin \text{dom}(D_\emptyset)$, this in turn implies that

$$\text{step}(C, e, D_\emptyset) \downarrow \wedge \text{step}(C, e, D_\emptyset) = D_\emptyset[r \mapsto I, \text{mid} \mapsto I].$$

Let $D' = D_\emptyset[r \mapsto I, \text{mid} \mapsto I]$. We should show that

$$C [\mathcal{G}^{\text{state}}] D'.$$

Since $\text{dom}(D') = \text{ReplicaID} \cup \{\text{mid}\}$, the first two conditions of $C [\mathcal{G}^{\text{state}}] D'$ hold:

$$\text{ReplicaID} \subseteq \text{dom}(D')$$

$$\wedge \text{dom}(D') \cap \text{MessageID} = \{C.\text{msg}(f) \mid C.\text{act}(f) = \text{send}\}.$$

For every $i \in \text{dom}(D')$, $D'(i).E = \emptyset$. Also, $\{f \in C.E \mid C.\text{act}(f) = \text{do}\} = \emptyset$. Thus,

$$\begin{aligned} & (\forall i \in \text{dom}(D'). \text{inv}(D'(i))) \\ & \wedge (\forall i \in \text{dom}(D'). D'(i).A = \text{abs}(C, \mathcal{V}^{\text{state}})|_{D'(i).E}) \\ & \wedge (\forall i \in \text{dom}(D'). \forall f \in D'(i).E. \\ & \quad f \in C.E \wedge D'(i).\text{info}(f) = C.\text{time}(f)) \\ & \wedge (\forall r' \in \text{dom}(D') \cap \text{ReplicaID}. \\ & \quad D'(r').E = (((\text{ro}(C) \cup \text{del}(C))^*)^{-1}(C|_{r'}.E))|_{\text{do}}) \\ & \wedge (\forall \text{mid} \in \text{dom}(D') \cap \text{MessageID}. \forall E_0. \\ & \quad E_0 = \{f \in C.E \mid C.\text{msg}(f) = \text{mid} \wedge C.\text{act}(f) = \text{send}\} \\ & \quad \implies D'(\text{mid}).E = (((\text{ro}(C) \cup \text{del}(C))^*)^{-1}E_0)|_{\text{do}}). \end{aligned}$$

The only remaining condition of $C \llbracket \mathcal{G}^{\text{state}} \rrbracket D'$ is:

$$\forall i, j \in \text{dom}(D'). i \neq j \implies \text{agree}(D'(i), D'(j)).$$

This holds because $D'(i).E = \emptyset$ for every $i \in \text{dom}(D')$.

Condition G-VIS. Consider

$$\begin{aligned} C_1, C_2 \in \text{CEx}[x], \quad e \in \text{abs}(C_1, \mathcal{V}^{\text{state}}).E, \\ (R, T) \in \text{Config}, \quad D \in \text{DEx} \end{aligned}$$

such that C_1 is a prefix of C_2 under $C_2.\text{eo}$. We should show that

$$\text{ctxt}(\text{abs}(C_1, \mathcal{V}^{\text{state}}), e) = \text{ctxt}(\text{abs}(C_2, \mathcal{V}^{\text{state}}), e).$$

Let

$$\begin{aligned} (o_1, G_1, \text{oper}_1, \text{vis}_1, \text{ar}_1) &= \text{ctxt}(\text{abs}(C_1, \mathcal{V}^{\text{state}}), e), \\ (o_2, G_2, \text{oper}_2, \text{vis}_2, \text{ar}_2) &= \text{ctxt}(\text{abs}(C_2, \mathcal{V}^{\text{state}}), e). \end{aligned}$$

Then, $o_1 = o_2$. Also, if $G_1 = G_2$, then $\text{oper}_1 = \text{oper}_2$ and $\text{ar}_1 = \text{ar}_2$. Thus, it suffices to prove that

$$G_1 = G_2 \wedge \text{vis}_1 = \text{vis}_2. \quad (22)$$

We recall the definitions of G_1, G_2, vis_1 and vis_2 :

$$\begin{aligned} G_1 &= (\mathcal{V}^{\text{state}}(C_1))^{-1}(e), \quad G_2 = (\mathcal{V}^{\text{state}}(C_2))^{-1}(e), \\ \text{vis}_1 &= \mathcal{V}^{\text{state}}(C_1)|_{G_1}, \quad \text{vis}_2 = \mathcal{V}^{\text{state}}(C_2)|_{G_2}. \end{aligned}$$

Let $E = \{e' \mid (e', e) \in (C_2.\text{eo})^*\}$. By the definition of $\mathcal{V}^{\text{state}}$,

$$(\mathcal{V}^{\text{state}}(C_1) \subseteq C_1.\text{eo} \subseteq C_2.\text{eo}) \wedge (\mathcal{V}^{\text{state}}(C_2) \subseteq C_2.\text{eo}).$$

Hence,

$$\begin{aligned} \mathcal{V}^{\text{state}}(C_1) \cap (\text{Event} \times E) &= \mathcal{V}^{\text{state}}(C_1) \cap (E \times E), \\ \mathcal{V}^{\text{state}}(C_2) \cap (\text{Event} \times E) &= \mathcal{V}^{\text{state}}(C_2) \cap (E \times E). \end{aligned}$$

Then, we can discharge our proof obligation in (22) by showing that

$$\mathcal{V}^{\text{state}}(C_1) \cap (E \times E) = \mathcal{V}^{\text{state}}(C_2) \cap (E \times E). \quad (23)$$

Since $(\text{ro}(C_1) \cup \text{del}(C_1)) \subseteq C_1.\text{eo} \subseteq C_2.\text{eo}$,

$$\begin{aligned} \mathcal{V}^{\text{state}}(C_1) \cap (E \times E) &= (\text{ro}(C_1) \cup \text{del}(C_1))^+|_{\text{do}} \cap (E \times E) \\ &= ((\text{ro}(C_1) \cup \text{del}(C_1)) \cap (E \times E))^+|_{\text{do}}. \end{aligned}$$

Similarly, since $(\text{ro}(C_2) \cup \text{del}(C_2)) \subseteq C_2.\text{eo}$,

$$\begin{aligned} \mathcal{V}^{\text{state}}(C_2) \cap (E \times E) &= (\text{ro}(C_2) \cup \text{del}(C_2))^+|_{\text{do}} \cap (E \times E) \\ &= ((\text{ro}(C_2) \cup \text{del}(C_2)) \cap (E \times E))^+|_{\text{do}}. \end{aligned}$$

This means that we can prove (23) by showing that

$$\begin{aligned} ((\text{ro}(C_1) \cup \text{del}(C_1)) \cap (E \times E)) \\ = ((\text{ro}(C_2) \cup \text{del}(C_2)) \cap (E \times E)). \end{aligned}$$

Figure 18. Instantiation for op-based implementations. $_n$ means the projection of the n -th component of a tuple. $I|_E$ is the restriction of I to events in E . The missing components in I' are as in Figure 6.

$$\begin{aligned} \text{AlInfo} &= \text{Timestamp} \times \{\text{old}, \text{new}\} \times \mathcal{P}(\text{Event}) \times \mathcal{P}(\text{Event}) \\ \text{inv}(I) &\iff (\forall e, f \in I.E. (e, f) \in I.\text{ar} \iff I.\text{info}(e).1 < I.\text{info}(f).1) \\ \text{agree}(I, J) &\iff I.E \cap J.E = \emptyset \\ \text{do}^\sharp(I, e, r, o, a, t) &= I', \quad \text{if } \text{inv}(I) \wedge e \notin I.E \wedge I' \in \text{IEx} \\ & \quad \text{where } g = (t, \text{new}, \{f \in I.E \mid I.\text{repl}(f) = r\}, I.E) \\ & \quad I' = ((I.E \cup \{e\}, I.\text{repl}[e \mapsto r], \dots), \text{info}[e \mapsto g]) \\ \text{send}^\sharp(I) &= (J, J|_E), \quad \text{if } \text{inv}(I) \\ & \quad \text{where } G(e) = (I.\text{info}(e).1, \text{old}, I.\text{info}(e).3, I.\text{info}(e).4) \\ & \quad J = (I.A, G) \\ & \quad E = \{e \in I.E \mid I.\text{info}(e).2 = \text{new}\} \\ \text{receive}^\sharp(I, J) &= J'', \quad \text{if } \text{inv}(I) \wedge \text{inv}(J) \wedge \text{agree}(I, J) \\ & \quad \wedge J'' \in \text{IEx} \\ & \quad \text{where } J' = I \sqcup J \\ & \quad R_r = \{(e, f) \in I.E \times J.E \mid e \in J.\text{info}(f).3\} \\ & \quad R'_r = \{(e, f) \in J.E \times I.E \mid e \in I.\text{info}(f).3\} \\ & \quad R_v = \{(e, f) \in I.E \times J.E \mid e \in J.\text{info}(f).4\} \\ & \quad R'_v = \{(e, f) \in J.E \times I.E \mid e \in I.\text{info}(f).4\} \\ & \quad J'' = (J'.E, J'.\text{repl}, J'.\text{obj}, J'.\text{oper}, J'.\text{rval}, \\ & \quad J'.\text{ro} \cup R_r \cup R'_r, J'.\text{vis} \cup R_v \cup R'_v, J'.\text{ar}), J'.\text{info}) \end{aligned}$$

This equality holds because of the definitions of ro and del , and the fact that

$$C_1.\text{eo} \cap (E \times E) = C_2.\text{eo} \cap (E \times E),$$

which itself is true because C_1 is a prefix of C_2 under $C_2.\text{eo}$. \square

B.3 Instantiation for op-based implementations

Fig. 18 gives an instantiation appropriate for op-based implementations, which ensures $\mathcal{D}_\tau \text{ sat } \mathcal{V}^{\text{op}}, \text{T-Unique} \llbracket \mathcal{F}_\tau \rrbracket$. In the op-based case, a message describes the operations the replica performed since the last send (§2). We keep track of the corresponding events by marking them as new in info . Thus, we relate a state $\langle a, d \rangle$ of the counter in Fig. 2 to executions with a increment operations, out of which d are new:

$$\begin{aligned} \langle a, d \rangle \llbracket \mathcal{R}_\tau \rrbracket I &\iff a = |\{e \in I.E \mid I.\text{oper}(e) = \text{inc}\}| \wedge d \llbracket \mathcal{M} \rrbracket I; \\ d \llbracket \mathcal{M} \rrbracket I &\iff \\ & \quad d = |\{e \in I.E \mid I.\text{oper}(e) = \text{inc} \wedge I.\text{info}(e) = (_, \text{new})\}|. \end{aligned}$$

For each event e , info also keeps track of its predecessors in replica order and visibility relations. This is required to update these relations on a receive^\sharp , similarly to how timestamps are needed to update arbitration. However, these components of info are not used in the case of the counter.

The definitions of abstract methods are similar to Fig. 6, except do^\sharp marks the event e as new in info , and send^\sharp returns the projection of the execution onto new events and marks them as old. This and the T-Unique specification ensure a key agreement property that the arguments of receive^\sharp have disjoint sets of events. Additionally, the methods update the last two components in info and use them to update replica order and visibility relation in receive^\sharp . The obligations for the counter in Fig. 2 can be discharged straightforwardly.

We now prove the soundness of abstract methods for op-based implementations. Our proof uses the notations in the previous subsection, as well as the visibility witness \mathcal{V}^{op} for the op-based implementations in §4, and the predicates inv and agree in Fig. 18.

We also need a few new notations. Let C be a concrete execution in $\text{CEx}[x]$. First, we define two relations on events in $C.E$:

$$\text{roi}(C) \subseteq (C.E \times C.E), \quad \text{rc}(C) \subseteq (C.E \times C.E)$$

Figure 19. Conditions for \mathcal{G}^{op}

$$\begin{aligned}
& \text{RepicalD} \subseteq \text{dom}(D) \\
& \wedge \text{dom}(D) \cap \text{MessageID} = \{C.\text{msg}(e) \mid C.\text{act}(e) = \text{send}\} \\
& \wedge (\forall i \in \text{dom}(D). \text{inv}(D(i))) \\
& \wedge (\forall i \in \text{dom}(D). D(i).A = \text{abs}(C, \mathcal{V}^{\text{op}})|_{D(i).E}) \\
& \wedge (\forall i \in \text{dom}(D). \forall e \in D(i).E. \\
& \quad e \in C.E \wedge D(i).\text{info}(e).1 = C.\text{time}(e)) \\
& \wedge (\forall r \in \text{dom}(D) \cap \text{RepicalD}. \\
& \quad D(r).E \supseteq (\text{abs}(C, \mathcal{V}^{\text{op}}).\text{vis})^{-1}(D(r)|_r.E)) \\
& \wedge (\forall r \in \text{dom}(D) \cap \text{RepicalD}. \\
& \quad D(r).E = (C|_r.E \cup ((\text{roi}(C); \text{del}(C))^{-1}C|_r.E))|_{\text{do}}) \\
& \wedge (\forall \text{mid} \in \text{dom}(D) \cap \text{MessageID}. \forall E_0. \\
& \quad E_0 = \{e \in C.E \mid C.\text{msg}(e) = \text{mid} \wedge C.\text{act}(e) = \text{send}\} \\
& \quad \implies D(\text{mid}).E = (\text{roi}(C)^{-1}E_0)|_{\text{do}}) \\
& \wedge (\forall r \in \text{dom}(D) \cap \text{RepicalD}. D(r).N = C|_r.N) \\
& \wedge (\forall i \in \text{dom}(D) \cap \text{RepicalD}. \forall j \in \text{dom}(D). \\
& \quad i \neq j \implies D(i).N \cap D(j).E = \emptyset) \\
& \wedge (\forall i \in \text{dom}(D) \cap \text{MessageID}. \forall j \in \text{dom}(D). \\
& \quad (i \neq j \wedge j \notin \text{rcv}(C, i)) \implies D(i).E \cap D(j).E = \emptyset). \\
& \wedge (\forall i \in \text{dom}(D) \cap \text{MessageID}. \forall e \in D(i).E. \\
& \quad D(i).\text{info}(e).2 = \text{old}) \\
& \wedge (\forall i \in \text{dom}(D). \forall e \in D(i).E. \\
& \quad D(i).\text{info}(e).3 = \text{abs}(C, \mathcal{V}^{\text{op}}).\text{ro}^{-1}(e)) \\
& \wedge (\forall i \in \text{dom}(D). \forall e \in D(i).E. \\
& \quad D(i).\text{info}(e).4 = \text{abs}(C, \mathcal{V}^{\text{op}}).\text{vis}^{-1}(e))
\end{aligned}$$

as follows:

$$\begin{aligned}
\text{roi}(C) &= \{(e, f) \mid e, f \in C.E \wedge C.\text{act}(e) = \text{do} \\
& \wedge C.\text{act}(f) = \text{send} \wedge e \xrightarrow{\text{ro}(C)} f \\
& \wedge \neg \exists e' \in C.E. e \xrightarrow{\text{ro}(C)} e' \xrightarrow{\text{ro}(C)} f \wedge C.\text{act}(e') = \text{send}\}
\end{aligned}$$

$$\begin{aligned}
\text{rc}(C) &= \{(e, f) \mid e, f \in C.E \wedge C.\text{act}(e) = \text{do} \\
& \wedge C.\text{act}(f) = \text{do} \wedge \exists e', f'. e \xrightarrow{\text{ro}(C)} e' \xrightarrow{\text{del}(C)} f' \xrightarrow{\text{ro}(C)} f \\
& \wedge \neg \exists e''. e \xrightarrow{\text{ro}(C)} e'' \xrightarrow{\text{ro}(C)} e' \wedge C.\text{act}(e'') = \text{send}\}.
\end{aligned}$$

Second, for every $\text{mid} \in \text{MessageID}$, we let

$$\begin{aligned}
\text{rcv}(C, \text{mid}) &= \\
& \{C.\text{repl}(e) \mid (C.\text{act}(e) = \text{receive} \vee C.\text{act}(e) = \text{send}) \\
& \quad \wedge C.\text{msg}(e) = \text{mid}\}
\end{aligned}$$

Fourth, we let

$$\begin{aligned}
C.N &= \{e \in C.E \mid \exists e' \in C.E. \\
& \quad e \xrightarrow{\text{ro}(C)} e' \wedge C.\text{act}(e') = \text{send}\}.
\end{aligned}$$

Fourth, for every instrumented execution I , we write

$$I.N = \{e \in I.E \mid I.\text{info}(e).2 = \text{new}\}.$$

Fifth, for each tuple $t = (x_1, \dots, x_n)$, we write $t.j$ to mean the j -th component x_j of this tuple. Finally, we define the relation \mathcal{G}^{op} between concrete executions and decomposed executions, which forms the core of our soundness argument in this subsection:

$$C [\mathcal{G}^{\text{op}}] D$$

if and only if the conditions in Fig. 19 hold

THEOREM 23. *There exists \mathcal{G} such that for all \mathcal{D}_τ , the parameters in Fig. 18 satisfy the conditions in Fig. 9 with respect to this \mathcal{G} , $\mathcal{V} = \mathcal{V}^{\text{op}}$, $\mathcal{T} = \text{T-Unique}$.*

PROOF. Let $\mathcal{T} = \text{T-Unique}$. We prove each of the four conditions separately.

Condition G-CTXT. Assume that

$$C [\mathcal{G}^{\text{op}}] D.$$

Pick e, r such that

$$e \in \text{abs}(C, \mathcal{V}^{\text{op}}).E \wedge \text{abs}(C, \mathcal{V}^{\text{op}}).\text{repl}(e) = r.$$

Since $C [\mathcal{G}^{\text{op}}] D$, we have $r \in \text{dom}(D)$. Let

$$I = D(r).$$

Also, since $C [\mathcal{G}^{\text{op}}] D$,

$$e \in (C|_r.E)|_{\text{do}} \subseteq I.E.$$

We use $C [\mathcal{G}^{\text{op}}] D$ again and derive that

$$\begin{aligned}
I.A &= \text{abs}(C, \mathcal{V}^{\text{op}})|_{I.E} \wedge \\
I.E &\supseteq (\text{abs}(C, \mathcal{V}^{\text{op}}).\text{vis})^{-1}(e).
\end{aligned}$$

Thus,

$$\begin{aligned}
\text{ctxt}(I.A, e) &= \text{ctxt}(\text{abs}(C, \mathcal{V}^{\text{op}})|_{I.E}, e) \\
&= \text{ctxt}(\text{abs}(C, \mathcal{V}^{\text{op}}), e).
\end{aligned}$$

Condition G-STEP. Consider

$C, C' \in \text{CEX}[x] \cap \mathcal{T}$, $e \in \text{Event}$, $(R, T) \in \text{Config}$, $D \in \text{DEX}$ such that

$$C' \sim C \xrightarrow{e} (R, T) \wedge C [\mathcal{G}^{\text{op}}] D.$$

We should prove that

$$\text{step}(C', e, D) \downarrow \wedge C' [\mathcal{G}^{\text{op}}] \text{step}(C', e, D). \quad (24)$$

We do the case analysis on $C'.\text{act}(e)$.

1. The first case is that $C'.\text{act}(e) = \text{do}$. Let

$$(r, o, a, t) = (C'.\text{repl}(e), C'.\text{oper}(e), C'.\text{rval}(e), C'.\text{time}(e)).$$

Since $C [\mathcal{G}^{\text{op}}] D$, $r \in \text{RepicalD}$ and $e \notin C.E$,

$$r \in \text{dom}(D) \wedge e \notin D(r).E.$$

Define I and J as follows:

$$\begin{aligned}
I &= D(r) \\
\text{ro}' &= I.\text{ro} \cup \{(f, e) \mid f \in I.E \wedge I.\text{repl}(f) = r\} \\
\text{vis}' &= I.\text{vis} \cup \{(f, e) \mid f \in I.E\}, \\
\text{ar}' &= I.\text{ar} \cup \{(f, e) \mid f \in I.E \wedge I.\text{info}(f).1 < t \\
& \quad \cup \{(e, f) \mid f \in I.E \wedge t < I.\text{info}(f).1\} \\
E_r &= \{f \in I.E \mid I.\text{repl}(f) = r\} \\
E_v &= I.E \\
J &= ((I.E \uplus \{e\}, I.\text{repl}[e \mapsto r], I.\text{obj}[e \mapsto x], \\
& \quad I.\text{oper}[e \mapsto o], I.\text{rval}[e \mapsto a], \text{ro}', \text{vis}', \text{ar}'), \\
& \quad I.\text{info}[e \mapsto (t, \text{new}, E_r, E_v)]).
\end{aligned}$$

Then, J is an instrumented execution, because I is an instrumented execution, $e \notin I.E$, $C'.\text{time}(e') = I.\text{info}(e').1$ for every $e' \in I.E$ and $C'.\text{time}$ is injective. Furthermore,

$$\text{inv}(J).$$

Since $C [\mathcal{G}^{\text{op}}] D$ and $D(r) = I$, we have $\text{inv}(I)$. Also, $e \notin I.E$. Hence, by the definitions of do^\sharp and J ,

$$\text{do}^\sharp(I, e, r, o, a, t) \downarrow \wedge \text{do}^\sharp(I, e, r, o, a, t) = J.$$

This in turn implies that

$$\text{step}(C', e, D) \downarrow \wedge \text{step}(C', e, D) = D[r \mapsto J].$$

Let $D' = D[r \mapsto J]$. Then,

$$\begin{aligned} \text{dom}(D') &= \text{dom}(D) \\ \wedge \forall i \in \text{dom}(D'). (D'(i) &= D(i) \vee D'(i) = J). \end{aligned}$$

Hence, the following conjuncts in the definition of $C' [\mathcal{G}^{\text{op}}] D'$ hold.

$$\begin{aligned} \text{ReplicaID} &\subseteq \text{dom}(D') \\ \wedge \text{dom}(D') \cap \text{MessageID} &= \{C'.\text{msg}(e) \mid C'.\text{act}(e) = \text{send}\} \\ \wedge (\forall i \in \text{dom}(D'). \text{inv}(D'(i))) \\ \wedge (\forall i \in \text{dom}(D'). \forall f \in D'(i).E. \\ &f \in C'.E \wedge D'(i).\text{info}(f).1 = C'.\text{time}(f)). \end{aligned}$$

We also have

$$\forall i \in \text{dom}(D). D'(i).A = \text{abs}(C', \mathcal{V}^{\text{op}})|_{D'(i).E}$$

for the following two reasons. First, for all $i \in \text{dom}(D)$ such that $i \neq r$,

$$\begin{aligned} D'(i).A &= D(i).A = \text{abs}(C, \mathcal{V}^{\text{op}})|_{D(i).E} \\ &= \text{abs}(C', \mathcal{V}^{\text{op}})|_{D(i).E} \\ &= \text{abs}(C', \mathcal{V}^{\text{op}})|_{D'(i).E}. \end{aligned}$$

Second, by $C [\mathcal{G}^{\text{op}}] D$,

$$\begin{aligned} (J.\text{vis})^{-1}(e) &= I.E \\ &= (C|_r.E \cup ((\text{roi}(C); \text{del}(C))^{-1}C|_r.E))|_{\text{do}} \\ &= (\text{abs}(C', \mathcal{V}^{\text{op}}).\text{vis})^{-1}(e). \end{aligned}$$

By the same reasons and the assumption $C [\mathcal{G}^{\text{op}}] D$,

$$\begin{aligned} \forall r' \in \text{dom}(D') \cap \text{ReplicaID}. \\ D'(r').E &\supseteq (\text{abs}(C', \mathcal{V}^{\text{op}}).\text{vis})^{-1}(D(r')|_{r'.E}) \wedge \\ D'(r').E &= (C'|_{r'.E} \cup ((\text{roi}(C'); \text{del}(C'))^{-1}C'|_{r'.E}))|_{\text{do}}. \end{aligned}$$

We do not change anything regarding messages when defining D' , so

$$\begin{aligned} \text{dom}(D) \cap \text{MessageID} &= \text{dom}(D') \cap \text{MessageID} \\ \wedge \forall \text{mid} \in \text{dom}(D) \cap \text{MessageID}. D(\text{mid}) &= D'(\text{mid}). \end{aligned}$$

Thus, from the assumption $C [\mathcal{G}^{\text{op}}] D$ it follows that

$$\begin{aligned} \forall \text{mid} \in \text{dom}(D') \cap \text{MessageID}. \forall E_0. \\ E_0 = \{f \in C'.E \mid C'.\text{msg}(f) = \text{mid} \wedge C'.\text{act}(f) = \text{send}\} \\ \implies D'(\text{mid}).E = (\text{roi}(C')^{-1}E_0)|_{\text{do}}. \end{aligned}$$

We will now show that

$$\forall r' \in \text{dom}(D') \cap \text{ReplicaID}. D'(r').N = C'|_{r'.N}.$$

Pick $r' \in \text{dom}(D') \cap \text{ReplicaID}$. If $r' \neq r$, the desired equality follows from our assumption that $C [\mathcal{G}^{\text{op}}] D$. Otherwise,

$$D'(r').N = D(r).N \cup \{e\} = C|_r.N \cup \{e\} = C'|_{r'.N}.$$

The next property to show is that

$$\begin{aligned} \forall i \in \text{dom}(D') \cap \text{ReplicaID}. \forall j \in \text{dom}(D'). \\ i \neq j \implies D'(i).N \cap D'(j) = \emptyset. \end{aligned}$$

Pick $i \in \text{dom}(D') \cap \text{ReplicaID}$ and $j \in \text{dom}(D')$ such that $i \neq j$. Since $C [\mathcal{G}^{\text{op}}] D$, it suffices to check the cases that $i = r$ or $j = r$. If $i = r$,

$$D'(i).N \cap D'(j) = (D(i).N \cup \{e\}) \cap D(j) = \emptyset.$$

If $j = r$,

$$D'(i).N \cap D'(j) = D(i).N \cap (D(j) \cup \{e\}) = \emptyset.$$

Now consider the following requirements of $C' [\mathcal{G}^{\text{op}}] D'$.

$$\begin{aligned} \forall i \in \text{dom}(D') \cap \text{MessageID}. \forall j \in \text{dom}(D'). \\ (i \neq j \wedge j \notin \text{rcv}(C', i)) \implies D'(i).E \cap D'(j).E = \emptyset. \end{aligned}$$

Pick $i \in \text{dom}(D') \cap \text{MessageID}$ and $j \in \text{dom}(D')$ such that

$$i \neq j \wedge j \notin \text{rcv}(C', i).$$

By the definitions of D' and C' ,

$$i, j \in \text{dom}(D) \wedge j \notin \text{rcv}(C, i) \wedge D'(i).E = D(i).E.$$

Hence, from our assumption $C [\mathcal{G}^{\text{op}}] D$ follows that

$$D(i).E \cap D(j).E = \emptyset.$$

Since $e \notin (D(i).E \cup D(j).E)$ and $D'(j).E \subseteq D(j).E \cup \{e\}$,

$$D'(i).E \cap D'(j).E = \emptyset.$$

The following requirement of $C' [\mathcal{G}^{\text{op}}] D'$ holds because D' and D are the same when their domains are restricted to MessageID and $C [\mathcal{G}^{\text{op}}] D$:

$$\begin{aligned} \forall i \in \text{dom}(D') \cap \text{MessageID}. \forall f \in D'(i).E. \\ D'(i).\text{info}(f).2 = \text{old}. \end{aligned}$$

It remains to show the following requirements:

$$\begin{aligned} (\forall i \in \text{dom}(D'). \forall f \in D'(i).E. \\ D'(i).\text{info}(f).3 = \text{abs}(C', \mathcal{V}^{\text{op}}).\text{ro}^{-1}(f)) \\ \wedge (\forall i \in \text{dom}(D'). \forall f \in D'(i).E. \\ D'(i).\text{info}(f).4 = \text{abs}(C', \mathcal{V}^{\text{op}}).\text{vis}^{-1}(f)) \end{aligned}$$

These requirements mostly follow from the assumption $C [\mathcal{G}^{\text{op}}] D$ and the way that D' and C' are defined from D and C , respectively. The only nontrivial things to check are:

$$\begin{aligned} D'(r).\text{info}(e).3 = I|_r.E = \text{abs}(C', \mathcal{V}^{\text{op}}).\text{ro}^{-1}(e), \\ D'(r).\text{info}(e).4 = I.E = \text{abs}(C', \mathcal{V}^{\text{op}}).\text{vis}^{-1}(e). \end{aligned}$$

Both equalities follow from the assumption that $C [\mathcal{G}^{\text{op}}] D$.

2. The next case is that $C'.\text{act}(e) = \text{send}$. Let

$$(r, \text{mid}) = (C'.\text{repl}(e), C'.\text{msg}(e)).$$

By the definitions of the transition relation and the concrete execution,

$$C'.\text{msg}(e) \notin \{C.\text{msg}(f) \mid f \in C.E \wedge C.\text{act}(f) = \text{send}\}.$$

Hence, from our assumption $C [\mathcal{G}^{\text{op}}] D$, it follows that

$$r \in \text{dom}(D) \wedge \text{mid} \notin \text{dom}(D).$$

Let $I = D(r)$. Then, $\text{inv}(I)$ holds, because $C [\mathcal{G}^{\text{op}}] D$. Thus,

$$\text{send}^\sharp(I) \downarrow \wedge \text{send}^\sharp(I) = (J, I')$$

where

$$\begin{aligned} J &= (I.A, \lambda e. (I.\text{info}(e).1, \text{old}, I.\text{info}(e).3, I.\text{info}(e).4)), \\ E &= \{e \in I.E \mid I.\text{info}(e).2 = \text{new}\}, \\ I' &= J|_E. \end{aligned}$$

This in turn implies that

$$\text{step}(C', e, D) \downarrow \wedge \text{step}(C', e, D) = D[r \mapsto J, \text{mid} \mapsto I'].$$

Let $D' = D[r \mapsto J, \text{mid} \mapsto I']$. We should show that

$$C' [\mathcal{G}^{\text{op}}] D'.$$

The following conjuncts in the definition of $C' [\mathcal{G}^{\text{op}}] D'$ follow immediately from the assumption $C [\mathcal{G}^{\text{op}}] D$ and the definitions

of J and I' :

$\text{ReplicaID} \subseteq \text{dom}(D')$

$\wedge \text{dom}(D') \cap \text{MessageID} = \{C'.\text{msg}(e) \mid C'.\text{act}(e) = \text{send}\}$

$\wedge (\forall i \in \text{dom}(D'). \text{inv}(D'(i)))$.

For the next conjunct in the definition of $C' [\mathcal{G}^{\text{op}}] D'$, we notice that $\text{abs}(C, \mathcal{V}^{\text{op}}) = \text{abs}(C', \mathcal{V}^{\text{op}})$. Hence, for every $i \in \text{dom}(D')$, there exists some $j \in \text{dom}(D)$ satisfying

$$\begin{aligned} D'(i).E &\subseteq D(j).E \wedge \\ D'(i).A &= D(j).A|_{D'(i).E}. \end{aligned}$$

Hence, for every $i \in \text{dom}(D')$, if we let j be such j ,

$$\begin{aligned} D'(i).A &= D(j).A|_{D'(i).E} \\ &= (\text{abs}(C, \mathcal{V}^{\text{op}})|_{D(j).E})|_{D'(i).E} \\ &= (\text{abs}(C', \mathcal{V}^{\text{op}})|_{D(j).E})|_{D'(i).E} \\ &= \text{abs}(C', \mathcal{V}^{\text{op}})|_{D'(j).E}. \end{aligned}$$

Now, we notice that

- (a) for all $r \in \text{ReplicaID}$, $D'(r)$ is the same as $D(r)$ except that the $\{\text{old}, \text{new}\}$ tag of each event;
- (b) for every $i \in \text{dom}(D')$, there exists some $j \in \text{dom}(D)$ such that $D'(i)$ is the restriction of $D(j)$ to $D'(i).E$ except the $\{\text{old}, \text{new}\}$ tags;
- (c) C' is an extension of C with e ; and
- (d) e is the send operation and happens at the end according to the $C'.\text{eo}$ order.

From these properties and the assumption $C [\mathcal{G}^{\text{op}}] D$ follow the following conjuncts in the definition of $C' [\mathcal{G}^{\text{op}}] D'$:

$$\begin{aligned} (\forall i \in \text{dom}(D'). \forall f \in D'(i).E. \\ f \in C'.E \wedge D'(i).\text{info}(f).1 = C'.\text{time}(f)) \\ \wedge (\forall r' \in \text{dom}(D') \cap \text{ReplicaID}. \\ D'(r').E \supseteq (\text{abs}(C', \mathcal{V}^{\text{op}}).\text{vis})^{-1}(D'(r')|_{r'}.E)) \\ \wedge (\forall r' \in \text{dom}(D') \cap \text{ReplicaID}. \\ D'(r').E = (C'|_{r'}.E \cup ((\text{roi}(C'); \text{del}(C'))^{-1}C'|_{r'}.E))|_{\text{do}}). \end{aligned}$$

The next requirement to check is:

$\forall \text{mid}' \in \text{dom}(D') \cap \text{MessageID}. \forall E_0$.

$$\begin{aligned} E_0 = \{f \in C'.E \mid C'.\text{msg}(f) = \text{mid}' \wedge C'.\text{act}(f) = \text{send}\} \\ \implies D'(\text{mid}').E = (\text{roi}(C')^{-1}E_0)|_{\text{do}}. \end{aligned}$$

Pick $\text{mid}' \in \text{dom}(D')$ and E_0 satisfying the condition of the above formula. Then, $\text{mid}' = \text{mid}$ or $\text{mid}' \in \text{dom}(D)$. If $\text{mid}' \neq \text{mid}$, we have

$$\begin{aligned} E_0 = \{f \in C.E \mid C.\text{msg}(f) = \text{mid}' \wedge C.\text{act}(f) = \text{send}\} \\ \wedge D'(\text{mid}') = D(\text{mid}') \wedge \text{roi}(C')^{-1}E_0 = \text{roi}(C)^{-1}E_0. \end{aligned}$$

The desired equality in this case follows from our assumption that $C [\mathcal{G}^{\text{op}}] D$. If $\text{mid} = \text{mid}'$, by the definition of the transition relation, we have

$$E_0 = \{\text{mid}\}.$$

We should show that

$$D'(\text{mid}).E = D(\text{roi}(C')^{-1}(\{\text{mid}\})).$$

This holds because by our assumption $C [\mathcal{G}^{\text{op}}] D$,

$$D(r).N = C|_{r}.N.$$

We move on to the following requirement in the definition of $C [\mathcal{G}^{\text{op}}] D$:

$$\forall r' \in \text{dom}(D') \cap \text{ReplicaID}. D'(r').N = C'|_{r'}.N.$$

Pick $r' \in \text{dom}(D') \cap \text{ReplicaID}$. If $r' \neq r$,

$$D'(r').N = D(r').N = C|_{r'}.N = C'|_{r'}.N.$$

If $r = r'$,

$$D'(r').N = \emptyset = C'|_{r'}.N.$$

Next we prove the following two properties:

$(\forall i \in \text{dom}(D') \cap \text{ReplicaID}. \forall j \in \text{dom}(D').$

$$i \neq j \implies D'(i).N \cap D'(j).E = \emptyset) \wedge$$

$(\forall i \in \text{dom}(D') \cap \text{MessageID}. \forall j \in \text{dom}(D').$

$$(i \neq j \wedge j \notin \text{rcv}(C', i)) \implies D'(i).E \cap D'(j).E = \emptyset).$$

To prove the first property, consider i, j such that

$$i \in \text{dom}(D') \cap \text{ReplicaID} \wedge j \in \text{dom}(D') \wedge i \neq j.$$

If $i = r$, the desired equality follows from the fact that $D'(i).N = \emptyset$. If $i \neq r$ but $j \neq \text{mid}$, we can use the assumption $C [\mathcal{G}^{\text{op}}] D$ and derive $D'(i).N \cap D'(j).E = \emptyset$. Finally, if $i \neq r$ and $j = \text{mid}$,

$$\begin{aligned} D'(i).N &= D(i).N \wedge D'(j).E \subseteq D(r).E \\ &\wedge D(i).N \cap D(r).E = \emptyset. \end{aligned}$$

From these follows the desired equality. Let's move on to the second property. Pick i, j such that

$$\begin{aligned} i \in \text{dom}(D') \cap \text{MessageID} \wedge j \in \text{dom}(D') \\ \wedge i \neq j \wedge j \notin \text{rcv}(C', i). \end{aligned}$$

If $i = \text{mid}$,

$$\begin{aligned} \text{rcv}(C', i) &= \text{rcv}(C, i) \cup \{e\} \\ \wedge j \neq r \wedge j \neq i \wedge D'(i).E &\subseteq D(r).N. \end{aligned}$$

Hence, $D'(j).E = D(j).E$. Also, by our assumption $C [\mathcal{G}^{\text{op}}] D$,

$$D(r).N \cap D(j).E = \emptyset.$$

From what we have just shown follows that $D'(i).E \cap D'(j).E = \emptyset$. Now assume that $i \neq \text{mid}$. Then,

$$\text{rcv}(C', i) = \text{rcv}(C, i).$$

We complete the proof of this case ($i \neq \text{mid}$) by doing the further case analysis on j . If $j = \text{mid}$, the desired equality follows by the below reasoning:

$$D'(i).E \cap D'(j).E \subseteq D(i).E \cap D(r).N = \emptyset.$$

If $j \neq \text{mid}$,

$$D'(i).E \cap D'(j).E = D(i).E \cap D(j).E = \emptyset.$$

The following requirement in the definition of $C' [\mathcal{G}^{\text{op}}] D'$ holds because of the assumption $C [\mathcal{G}^{\text{op}}] D$ and the definition of $D'(\text{mid})$ (i.e., I'):

$$\begin{aligned} \forall i \in \text{dom}(D') \cap \text{MessageID}. \forall f \in D'(i).E. \\ D'(i).\text{info}(f).2 = \text{old}. \end{aligned}$$

It remains to prove the following properties:

$$(\forall i \in \text{dom}(D'). \forall f \in D'(i).E.$$

$$D'(i).\text{info}(f).3 = \text{abs}(C', \mathcal{V}^{\text{op}}).\text{ro}^{-1}(f))$$

$$\wedge (\forall i \in \text{dom}(D'). \forall f \in D'(i).E.$$

$$D'(i).\text{info}(f).4 = \text{abs}(C', \mathcal{V}^{\text{op}}).\text{vis}^{-1}(f))$$

These two follow from the assumption $C [\mathcal{G}^{\text{op}}] D$, and the facts that the third and fourth components of info haven't changed from D to D' and

$$\begin{aligned} \text{abs}(C', \mathcal{V}^{\text{op}}).\text{ro} &= \text{abs}(C, \mathcal{V}^{\text{op}}).\text{ro} \wedge \\ \text{abs}(C', \mathcal{V}^{\text{op}}).\text{vis} &= \text{abs}(C, \mathcal{V}^{\text{op}}).\text{vis}. \end{aligned}$$

3. The last case is that $C'.\text{act}(e) = \text{receive}$. Let

$$r = C'.\text{repl}(e), \quad \text{mid} = C'.\text{msg}(e) \text{ and } (R', T') = C'.\text{pre}(e).$$

By the definition of the transition relation, $\text{mid} \in \text{dom}(T')$. Then, since C' is a concrete execution starting from a configuration with the empty T component and $C' \sim C \xrightarrow{e} (R, T)$,

$$\text{mid} \in \{C.\text{msg}(f) \mid f \in C.E \wedge C.\text{act}(f) = \text{send}\}.$$

Thus, from our assumption $C \llbracket \mathcal{G}^{\text{op}} \rrbracket D$, it follows that

$$r \in \text{dom}(D) \wedge \text{mid} \in \text{dom}(D).$$

Let

$$\begin{aligned} I_1 &= D(r), & I_2 &= D(\text{mid}), & J &= I_1 \sqcup I_2, \\ R_r &= \{(e, f) \in I_1.E \times I_2.E \mid e \in I_2.\text{info}(f).3\}, \\ R'_r &= \{(e, f) \in I_2.E \times I_1.E \mid e \in I_1.\text{info}(f).3\}, \\ R_v &= \{(e, f) \in I_1.E \times I_2.E \mid e \in I_2.\text{info}(f).4\}, \\ R'_v &= \{(e, f) \in I_2.E \times I_1.E \mid e \in I_1.\text{info}(f).4\}, \\ J' &= ((J.E, J.\text{repl}, J.\text{obj}, J.\text{oper}, J.\text{rval}, \\ &\quad J.\text{ro} \cup R_r \cup R'_r, J.\text{vis} \cup R_v \cup R'_v, J.\text{ar}), J.\text{info}). \end{aligned}$$

Here $J = I_1 \sqcup I_2$ is defined because

$$\begin{aligned} I_1.A &= \text{abs}(C, \mathcal{V}^{\text{op}})|_{I_1.E} \\ \wedge I_2.A &= \text{abs}(C, \mathcal{V}^{\text{op}})|_{I_2.E} \\ \wedge (\forall i \in \{1, 2\}. \forall f \in I_i.E. I_i.\text{info}(f).1 &= C.\text{time}(f)) \\ \wedge (\forall i \in \{1, 2\}. \forall f \in I_i.E. \\ &\quad I_i.\text{info}(f).3 = \text{abs}(C, \mathcal{V}^{\text{op}}).\text{ro}^{-1}(f)) \\ \wedge (\forall i \in \{1, 2\}. \forall f \in I_i.E. \\ &\quad I_i.\text{info}(f).4 = \text{abs}(C, \mathcal{V}^{\text{op}}).\text{vis}^{-1}(f)). \end{aligned}$$

Also, by the same reason and the definition of J' ,

$$\begin{aligned} J'.A &= \text{abs}(C, \mathcal{V}^{\text{op}})|_{J'.E} \\ \wedge (\forall f \in J'.E. J'.\text{info}(f).1 &= C.\text{time}(f)) \\ \wedge (\forall f \in J'.E. J'.\text{info}(f).3 &= \text{abs}(C, \mathcal{V}^{\text{op}}).\text{ro}^{-1}(f)) \quad (25) \\ \wedge (\forall f \in J'.E. J'.\text{info}(f).4 &= \text{abs}(C, \mathcal{V}^{\text{op}}).\text{vis}^{-1}(f)). \end{aligned}$$

One consequence of this is that J' is a well-defined instrumented execution. Since $C \in \mathcal{T} = \text{T-Unique}$, the message mid is not delivered to the replica r in the concrete execution C :

$$r \notin \text{rcv}(C, \text{mid}).$$

Hence, our assumption $C \llbracket \mathcal{G}^{\text{op}} \rrbracket D$ implies that

$$\text{inv}(I_1) \wedge \text{inv}(I_2) \wedge I_1.E \cap I_2.E = \emptyset.$$

Notice that by showing $J' \in \text{IEx}$ and the above properties of I_1 and I_2 , we have shown that the side condition of $\text{receive}^\#(I_1, I_2)$ is met. Hence $\text{receive}^\#(I_1, I_2) \downarrow$. Let

$$D' = D[r \mapsto J'].$$

It remains to show that $C' \llbracket \mathcal{G}^{\text{op}} \rrbracket D'$. Since $\text{dom}(D') = \text{dom}(C')$ and $C'.\text{act}(e) = \text{receive}$, the assumption $C \llbracket \mathcal{G}^{\text{op}} \rrbracket D$ gives

$\text{ReplicaID} \subseteq \text{dom}(D')$

$$\wedge \text{dom}(D') \cap \text{MessageID} = \{C'.\text{msg}(f) \mid C'.\text{act}(f) = \text{send}\}$$

The next requirements in the definition of $C' \llbracket \mathcal{G}^{\text{op}} \rrbracket D'$ are:

$$\begin{aligned} (\forall i \in \text{dom}(D'). \text{inv}(D'(i))) \\ \wedge (\forall i \in \text{dom}(D'). D'(i).A &= \text{abs}(C', \mathcal{V}^{\text{op}})|_{D'(i).E}) \\ \wedge (\forall i \in \text{dom}(D'). \forall f \in D'(i).E. \\ &\quad f \in C'.E \wedge D'(i).\text{info}(f).1 = C'.\text{time}(f)) \end{aligned}$$

Since $C \llbracket \mathcal{G}^{\text{op}} \rrbracket D$ and $\text{abs}(C, \mathcal{V}^{\text{op}}) = \text{abs}(C', \mathcal{V}^{\text{op}})$, we just need to show that

$$\begin{aligned} \text{inv}(J') \wedge (J'.A &= \text{abs}(C', \mathcal{V}^{\text{op}})|_{J'.E}) \\ \wedge (\forall f \in J'.E. f \in C'.E \wedge J'.\text{info}(f).1 &= C'.\text{time}(f)). \end{aligned}$$

All these follow from (25). We move on to the following conditions of $C' \llbracket \mathcal{G}^{\text{op}} \rrbracket D'$:

$$\begin{aligned} (\forall r' \in \text{dom}(D') \cap \text{ReplicaID}. \\ D'(r').E \supseteq (\text{abs}(C', \mathcal{V}^{\text{op}}).\text{vis})^{-1}(D'(r)|_{r'.E})) \wedge \\ (\forall r' \in \text{dom}(D') \cap \text{ReplicaID}. \\ D'(r').E = (C'|_{r'.E} \cup (\text{roi}(C'); \text{del}(C'))^{-1}(C'|_{r'.E}))|_{\text{do}}) \wedge \\ (\forall \text{mid}' \in \text{dom}(D') \cap \text{MessageID}. \forall E_0. \\ E_0 = \{f \in C'.E \mid C'.\text{msg}(f) = \text{mid}' \wedge C'.\text{act}(f) = \text{send}\} \\ \implies D'(\text{mid}').E = (\text{roi}(C')^{-1}E_0)|_{\text{do}}). \end{aligned}$$

Proving all the three conditions uses the assumption $C \llbracket \mathcal{G}^{\text{op}} \rrbracket D$ and the fact that $\text{abs}(C, \mathcal{V}^{\text{op}}) = \text{abs}(C', \mathcal{V}^{\text{op}})$, which we will not mention explicitly. The first condition holds because $J'.E = I_1.E \cup I_2.E$ and

$$I_2|r.E = \emptyset \wedge I_1.E \supseteq (\text{abs}(C, \mathcal{V}^{\text{op}}).\text{vis})^{-1}(I_1|r.E).$$

The third condition holds because D' and D are the same when restricted to the subdomain MessageID and the event e is the receive event that is maximal according to both $\text{ro}(C')$ and $\text{del}(C')$. For the second condition, since $C'.\text{act}(e) = \text{receive}$ and e is maximal in $\text{ro}(C')$ and $\text{del}(C')$, it is sufficient to show that

$$J'.E = (C'|_{r'.E} \cup (\text{roi}(C'); \text{del}(C'))^{-1}(C'|_{r'.E}))|_{\text{do}}. \quad (26)$$

Let

$$E_0 = \{f \in C.E \mid C.\text{msg}(f) = \text{mid} \wedge C.\text{act}(f) = \text{send}\}.$$

We prove the equality in (26) as follows:

$$\begin{aligned} J'.E &= I_1.E \cup I_2.E \\ &= (C|_{r.E} \cup (\text{roi}(C); \text{del}(C))^{-1}(C|_{r.E}))|_{\text{do}} \\ &\quad \cup (\text{roi}(C)^{-1}E_0)|_{\text{do}} \\ &= (C|_{r.E} \cup (\text{roi}(C); \text{del}(C))^{-1}(C|_{r.E}) \cup \text{roi}(C)^{-1}E_0)|_{\text{do}} \\ &= (C'|_{r'.E} \cup (\text{roi}(C'); \text{del}(C'))^{-1}(C'|_{r'.E} \cup \{e\}))|_{\text{do}} \\ &= (C'|_{r'.E} \cup (\text{roi}(C'); \text{del}(C'))^{-1}(C'|_{r'.E}))|_{\text{do}}. \end{aligned}$$

We now consider the following requirements in the definition of $C' \llbracket \mathcal{G}^{\text{op}} \rrbracket D'$:

$$\begin{aligned} (\forall r' \in \text{dom}(D') \cap \text{ReplicaID}. D'(r').N = C'|_{r'.N}) \wedge \\ (\forall i \in \text{dom}(D') \cap \text{ReplicaID}. \forall j \in \text{dom}(D'). \\ i \neq j \implies D'(i).N \cap D'(j).E = \emptyset) \wedge \\ (\forall i \in \text{dom}(D') \cap \text{MessageID}. \forall j \in \text{dom}(D'). \\ (i \neq j \wedge j \notin \text{rcv}(C', i)) \implies D'(i).E \cap D'(j).E = \emptyset). \end{aligned}$$

The first holds because for every $r' \in \text{dom}(D') \cap \text{ReplicaID}$,

$$r' \in \text{dom}(D) \wedge D'(r').N = D(r').N \wedge C'|_{r'.N} = C|_{r'.N},$$

which itself follows from the assumption $C \llbracket \mathcal{G}^{\text{op}} \rrbracket D$ and the fact that the event e is a receive event. For the other two requirements, since $\text{rcv}(C', \text{mid}) = \text{rcv}(C, \text{mid}) \cup \{r\}$, $D'(r).N = D(r).N$ and similar requirements hold for C and D , it suffices to prove

that

$$\begin{aligned} & (\forall r' \in \text{dom}(D) \cap \text{ReplicaID}. \\ & \quad r' \neq r \implies D(r').N \cap J'.E = \emptyset) \wedge \\ & (\forall \text{mid}' \in \text{dom}(D) \cap \text{ReplicaID}. \\ & \quad \text{mid}' \neq \text{mid} \wedge r \notin \text{rcv}(C, \text{mid}') \\ & \quad \implies D(\text{mid}').E \cap J'.E = \emptyset). \end{aligned}$$

The first holds because for all r' with $r' \neq r$,

$$D(r').N \cap J'.E = D(r').N \cap I_1.E \cup D(r').N \cap I_2.E = \emptyset.$$

Also, the second holds because for all mid' with $\text{mid}' \neq \text{mid}$ and $r \notin \text{rcv}(C, \text{mid}')$,

$$D(\text{mid}').E \cap J'.E = D(\text{mid}').E \cap (I_1.E \cup I_2.E) = \emptyset.$$

The remaining requirements of $C' [\mathcal{G}^{\text{op}}] D'$ are as follows:

$$\begin{aligned} & (\forall i \in \text{dom}(D') \cap \text{MessageID}. \forall f \in D'(i).E. \\ & \quad D'(i).\text{info}(f).2 = \text{old}) \wedge \\ & (\forall i \in \text{dom}(D'). \forall f \in D'(i).E. \\ & \quad D'(i).\text{info}(f).3 = \text{abs}(C', \mathcal{V}^{\text{op}}).\text{ro}^{-1}(f)) \wedge \\ & (\forall i \in \text{dom}(D'). \forall f \in D'(i).E. \\ & \quad D'(i).\text{info}(f).4 = \text{abs}(C', \mathcal{V}^{\text{op}}).\text{vis}^{-1}(f)). \end{aligned}$$

The first requirement holds because a similar property holds for D and D and D' are identical when their domains are restricted to MessageID . Regarding the other two requirements, we notice that for every $i \in \text{dom}(D')$, the info part of $D'(i)$ is obtained from that of $D(j)$ and/or $D(j')$ for some $j, j' \in \text{dom}(D)$. Furthermore, $\text{abs}(C, \mathcal{V}^{\text{op}}) = \text{abs}(C', \mathcal{V}^{\text{op}})$ and

$$\begin{aligned} & (\bigcup \{D'(i).E \mid i \in \text{dom}(D')\}) \\ & = (\bigcup \{D(i).E \mid i \in \text{dom}(D)\}). \end{aligned}$$

From these facts follow the two requirements.

Condition G-INIT. Consider C, e, R, T such that

$$C.E = \{e\} \wedge E.\text{pre}(e) = (R, T) \wedge \text{dom}(T) = \emptyset.$$

We have to prove that

$$\text{step}(C, e, D_0) \downarrow \wedge C [\mathcal{G}^{\text{op}}] \text{step}(C, e, D_0).$$

By the definition of the transition relation,

$$C.\text{act}(e) = \text{do} \vee C.\text{act}(e) = \text{send}.$$

We do the case analysis on these two disjuncts.

First, assume that $C.\text{act}(e) = \text{do}$. Let

$$(r, o, a, t) = (C.\text{repl}(e), C.\text{oper}(e), C.\text{rval}(e), C.\text{time}(e)).$$

Also, define I and J as follows:

$$\begin{aligned} I &= D_0(r) \\ J &= ((\{e\}, [e \mapsto r], [e \mapsto x], [e \mapsto o], [e \mapsto a], \emptyset, \emptyset, \emptyset), \\ & \quad [e \mapsto (t, \text{new}, \emptyset, \emptyset)]). \end{aligned}$$

Then, $I.E = \emptyset$, so $\text{inv}(I)$ holds and $e \notin I$. Also, $J \in \text{IE}x$. Hence,

$$\text{do}^\sharp(I, e, r, o, a, t) \downarrow \wedge \text{do}^\sharp(I, e, r, o, a, t) = J.$$

This in turn implies that

$$\text{step}(C, e, D_0) \downarrow \wedge \text{step}(C, e, D_0) = D_0[r \mapsto J].$$

Let $D' = D_0[r \mapsto J]$. We will now show that

$$C [\mathcal{G}^{\text{op}}] D'.$$

Since $\text{dom}(D') = \text{ReplicaID}$, $C.E = \{e\}$ and $C.\text{act}(e) = \text{do}$, the first two conditions of $C [\mathcal{G}^{\text{op}}] D'$ hold:

$$\begin{aligned} & \text{ReplicaID} \subseteq \text{dom}(D') \wedge \\ & \text{dom}(D') \cap \text{MessageID} = \emptyset = \{C.\text{msg}(f) \mid C.\text{act}(f) = \text{send}\}. \end{aligned}$$

For every $i \in \text{dom}(D')$, $D'(i)$ is I_0 (i.e., an instrumented execution with the empty event set) or J . In particular, $D'(i).E$ is empty or the singleton set $\{e\}$. Hence,

$$\begin{aligned} & (\forall i \in \text{dom}(D'). \text{inv}(D'(i))) \\ & \wedge (\forall i \in \text{dom}(D'). D'(i).A = \text{abs}(C, \mathcal{V}^{\text{op}})|_{D'(i).E}) \\ & \wedge (\forall i \in \text{dom}(D'). \forall f \in D'(i).E. \\ & \quad f \in C.E \wedge D'(i).\text{info}(f).1 = C.\text{time}(f)) \\ & \wedge (\forall r' \in \text{dom}(D') \cap \text{ReplicaID}. \\ & \quad D'(r').E \supseteq (\text{abs}(C, \mathcal{V}^{\text{op}}).\text{vis})^{-1}(D'(r')|_{r'}.E)) \\ & \wedge (\forall r' \in \text{dom}(D') \cap \text{ReplicaID}. \\ & \quad D'(r').E = (C|_{r'}.E \cup ((\text{roi}(C); \text{del}(C))^{-1}C|_{r'}.E))|_{\text{do}}). \end{aligned}$$

Since there are no $\text{mid} \in \text{dom}(D')$,

$$\forall \text{mid} \in \text{dom}(D') \cap \text{MessageID}. \forall E_0.$$

$$\begin{aligned} E_0 &= \{f \in C.E \mid C.\text{msg}(f) = \text{mid} \wedge C.\text{act}(f) = \text{send}\} \\ &\implies D'(\text{mid}).E = (\text{roi}(C)^{-1}E_0)|_{\text{do}}. \end{aligned}$$

Since only $D'(r)$ has one event e marked new by its info and all of the other $D'(i)$ with $i \neq r$ have no events,

$$\begin{aligned} & (\forall r' \in \text{dom}(D') \cap \text{ReplicaID}. D'(r').N = C|_{r'}.N) \\ & \wedge (\forall i \in \text{dom}(D') \cap \text{ReplicaID}. \forall j \in \text{dom}(D'). \\ & \quad i \neq j \implies D'(i).N \cap D'(j).E = \emptyset) \\ & \wedge (\forall i \in \text{dom}(D') \cap \text{MessageID}. \forall j \in \text{dom}(D'). \\ & \quad (i \neq j \wedge j \notin \text{rcv}(C, i)) \implies D'(i).E \cap D'(j).E = \emptyset). \end{aligned}$$

The remaining requirements of $C [\mathcal{G}^{\text{op}}] D'$ are:

$$\begin{aligned} & (\forall i \in \text{dom}(D') \cap \text{MessageID}. \forall f \in D'(i).E. \\ & \quad D'(i).\text{info}(f).2 = \text{old}) \\ & \wedge (\forall i \in \text{dom}(D'). \forall f \in D'(i).E. \\ & \quad D'(i).\text{info}(f).3 = \text{abs}(C, \mathcal{V}^{\text{op}}).\text{ro}^{-1}(f)) \\ & \wedge (\forall i \in \text{dom}(D'). \forall f \in D'(i).E. \\ & \quad D'(i).\text{info}(f).4 = \text{abs}(C, \mathcal{V}^{\text{op}}).\text{vis}^{-1}(f)). \end{aligned}$$

The first holds because $\text{dom}(D') \cap \text{MessageID} = \emptyset$, and the other two holds because both $\text{abs}(C, \mathcal{V}^{\text{op}})$ and $\text{abs}(C, \mathcal{V}^{\text{op}})$ are the empty relations and

$$J.\text{info}(e).3 = J.\text{info}(e).4 = \emptyset.$$

Next, assume that $C.\text{act}(e) = \text{send}$. Let

$$(r, \text{mid}) = (C.\text{repl}(e), C.\text{msg}(e)).$$

Since $\text{dom}(D_0) = \text{ReplicaID}$,

$$\text{mid} \notin \text{dom}(D_0) \wedge r \in \text{dom}(D_0).$$

Let $I = D_0(r)$. Then, $I.E = \emptyset$ by the definition of D_0 . Thus, $\text{inv}(I)$ holds, from which and the fact that $I.E = \emptyset$ follows that

$$\text{send}^\sharp(I) \downarrow \wedge \text{send}^\sharp(I) = (I, I).$$

Since $\text{mid} \notin \text{dom}(D_0)$, this in turn implies that

$$\text{step}(C, e, D_0) \downarrow \wedge \text{step}(C, e, D_0) = D_0[r \mapsto I, \text{mid} \mapsto I].$$

Let $D' = D_0[r \mapsto I, \text{mid} \mapsto I]$. We should show that

$$C [\mathcal{G}^{\text{op}}] D'.$$

Since $\text{dom}(D') = \text{ReplicaID} \cup \{\text{mid}\}$, the first two conditions of $C \llbracket \mathcal{G}^{\text{op}} \rrbracket D'$ hold:

$$\begin{aligned} & \text{ReplicaID} \subseteq \text{dom}(D') \\ & \wedge \text{dom}(D') \cap \text{MessageID} = \{C.\text{msg}(f) \mid C.\text{act}(f) = \text{send}\}. \end{aligned}$$

For every $i \in \text{dom}(D')$, $D'(i).E = \emptyset$. Also, $\{f \in C.E \mid C.\text{act}(f) = \text{do}\} = \emptyset$. The following remaining requirements of $C \llbracket \mathcal{G}^{\text{op}} \rrbracket D'$ follow from these facts:

$$\begin{aligned} & (\forall i \in \text{dom}(D'). \text{inv}(D'(i))) \\ & \wedge (\forall i \in \text{dom}(D'). D'(i).A = \text{abs}(C, \mathcal{V}^{\text{op}})|_{D'(i).E}) \\ & \wedge (\forall i \in \text{dom}(D'). \forall f \in D'(i).E. \\ & \quad f \in C.E \wedge D'(i).\text{info}(f).1 = C.\text{time}(f)) \\ & \wedge (\forall r' \in \text{dom}(D') \cap \text{ReplicaID}. \\ & \quad D'(r').E \supseteq (\text{abs}(C, \mathcal{V}^{\text{op}}).\text{vis})^{-1}(D'(r')|_{r'}.E)) \\ & \wedge (\forall r' \in \text{dom}(D') \cap \text{ReplicaID}. \\ & \quad D'(r').E = (C|_{r'}.E \cup ((\text{roi}(C); \text{del}(C))^{-1}C|_{r'}.E))|_{\text{do}}) \\ & \wedge (\forall \text{mid}' \in \text{dom}(D') \cap \text{MessageID}. \forall E_0. \\ & \quad E_0 = \{f \in C.E \mid C.\text{msg}(f) = \text{mid}' \wedge C.\text{act}(f) = \text{send}\} \\ & \quad \implies D'(\text{mid}').E = (\text{roi}(C)^{-1}E_0)|_{\text{do}}) \\ & \wedge (\forall r' \in \text{dom}(D') \cap \text{ReplicaID}. D'(r').N = C|_{r'}.N) \\ & \wedge (\forall i \in \text{dom}(D') \cap \text{ReplicaID}. \forall j \in \text{dom}(D'). \\ & \quad i \neq j \implies D'(i).N \cap D'(j).E = \emptyset) \\ & \wedge (\forall i \in \text{dom}(D') \cap \text{MessageID}. \forall j \in \text{dom}(D'). \\ & \quad (i \neq j \wedge j \notin \text{rcv}(C, i)) \implies D'(i).E \cap D'(j).E = \emptyset). \\ & \wedge (\forall i \in \text{dom}(D') \cap \text{MessageID}. \forall f \in D'(i).E. \\ & \quad D'(i).\text{info}(f).2 = \text{old}) \\ & \wedge (\forall i \in \text{dom}(D'). \forall f \in D'(i).E. \\ & \quad D'(i).\text{info}(f).3 = \text{abs}(C, \mathcal{V}^{\text{op}}).\text{ro}^{-1}(f)) \\ & \wedge (\forall i \in \text{dom}(D'). \forall f \in D'(i).E. \\ & \quad D'(i).\text{info}(f).4 = \text{abs}(C, \mathcal{V}^{\text{op}}).\text{vis}^{-1}(f)). \end{aligned}$$

Condition G-VIS. Consider

$$\begin{aligned} C_1, C_2 & \in \text{CEX}[x], & e & \in \text{abs}(C_1, \mathcal{V}^{\text{op}}).E, \\ (R, T) & \in \text{Config}, & D & \in \text{DEX} \end{aligned}$$

such that C_1 is a prefix of C_2 under $C_2.\text{eo}$. We should show that

$$\text{ctxt}(\text{abs}(C_1, \mathcal{V}^{\text{op}}), e) = \text{ctxt}(\text{abs}(C_2, \mathcal{V}^{\text{op}}), e).$$

Let

$$\begin{aligned} (o_1, G_1, \text{oper}_1, \text{vis}_1, \text{ar}_1) & = \text{ctxt}(\text{abs}(C_1, \mathcal{V}^{\text{op}}), e), \\ (o_2, G_2, \text{oper}_2, \text{vis}_2, \text{ar}_2) & = \text{ctxt}(\text{abs}(C_2, \mathcal{V}^{\text{op}}), e). \end{aligned}$$

Then, $o_1 = o_2$. Also, if $G_1 = G_2$, then $\text{oper}_1 = \text{oper}_2$ and $\text{ar}_1 = \text{ar}_2$. Thus, it suffices to prove that

$$G_1 = G_2 \wedge \text{vis}_1 = \text{vis}_2. \quad (27)$$

We recall the definitions of G_1, G_2, vis_1 and vis_2 :

$$\begin{aligned} G_1 & = (\mathcal{V}^{\text{op}}(C_1))^{-1}(e), & G_2 & = (\mathcal{V}^{\text{op}}(C_2))^{-1}(e), \\ \text{vis}_1 & = \mathcal{V}^{\text{op}}(C_1)|_{G_1}, & \text{vis}_2 & = \mathcal{V}^{\text{op}}(C_2)|_{G_2}. \end{aligned}$$

Let $E = \{e' \mid (e', e) \in (C_2.\text{eo})^*\}$. Notice that since $e \in C_1.E$ and C_1 is a prefix of C_2 according to the $C_2.\text{eo}$ order,

$$E = \{e' \mid (e', e) \in (C_1.\text{eo})^*\}.$$

By the definition of \mathcal{V}^{op} ,

$$(\mathcal{V}^{\text{op}}(C_1) \subseteq C_1.\text{eo} \subseteq C_2.\text{eo}) \wedge (\mathcal{V}^{\text{op}}(C_2) \subseteq C_2.\text{eo}).$$

Hence,

$$\begin{aligned} \mathcal{V}^{\text{op}}(C_1) \cap (\text{Event} \times E) & = \mathcal{V}^{\text{op}}(C_1) \cap (E \times E), \\ \mathcal{V}^{\text{op}}(C_2) \cap (\text{Event} \times E) & = \mathcal{V}^{\text{op}}(C_2) \cap (E \times E). \end{aligned}$$

Then, we can discharge our proof obligation in (27) by showing that

$$\mathcal{V}^{\text{op}}(C_1) \cap (E \times E) = \mathcal{V}^{\text{op}}(C_2) \cap (E \times E). \quad (28)$$

For each C_i , define a relation $\text{rc}(C_i)$ by

$$\begin{aligned} \text{rc}(C_i) & = \{(e, f) \mid e, f \in C_i.E \wedge C_i.\text{act}(e) = \text{do} \\ & \quad \wedge C_i.\text{act}(f) = \text{do} \wedge \exists e', f'. e \xrightarrow{\text{ro}(C_i)} e' \xrightarrow{\text{del}(C_i)} f' \xrightarrow{\text{ro}(C_i)} f \\ & \quad \wedge \neg \exists e''. e \xrightarrow{\text{ro}(C_i)} e'' \xrightarrow{\text{ro}(C_i)} e' \wedge C_i.\text{act}(e'') = \text{send}\}. \end{aligned}$$

Then,

$$\begin{aligned} \mathcal{V}^{\text{op}}(C_1) \cap (E \times E) & = (\text{ro}(C_1) \cup \text{rc}(C_1))|_{\text{do}} \cap (E \times E) \\ & = ((\text{ro}(C_1) \cup \text{rc}(C_1)) \cap (E \times E))|_{\text{do}}. \end{aligned}$$

Similarly,

$$\begin{aligned} \mathcal{V}^{\text{op}}(C_2) \cap (E \times E) & = (\text{ro}(C_2) \cup \text{rc}(C_2))|_{\text{do}} \cap (E \times E) \\ & = ((\text{ro}(C_2) \cup \text{rc}(C_2)) \cap (E \times E))|_{\text{do}}. \end{aligned}$$

This means that we can prove (28) by showing that

$$\begin{aligned} \text{ro}(C_1) \cap (E \times E) & = \text{ro}(C_2) \cap (E \times E) \\ \wedge \text{rc}(C_1) \cap (E \times E) & = \text{rc}(C_2) \cap (E \times E). \end{aligned}$$

The first equality follows from the definition of ro , and the fact that

$$C_1.\text{eo} \cap (E \times E) = C_2.\text{eo} \cap (E \times E),$$

which holds because C_1 is a prefix of C_2 under $C_2.\text{eo}$. The second equality follows from the fact that for all events e, f , if $e \in E$,

$$\begin{aligned} (f \xrightarrow{\text{ro}(C_1)} e & \iff f \xrightarrow{\text{ro}(C_2)} e) \\ \wedge (f \xrightarrow{\text{del}(C_1)} e & \iff f \xrightarrow{\text{del}(C_2)} e) \\ \wedge (f \xrightarrow{\text{ro}(C_1)} e & \implies f \in E) \\ \wedge (f \xrightarrow{\text{del}(C_1)} e & \implies f \in E). \end{aligned}$$

□

C. Additional space bounds proofs

C.1 Encodings

Table 1 presents standard recursive encoding schemes for values from different domains and their asymptotic cost. These schemes are employed to encode replica state and return values throughout all our proofs.

C.2 Lower bound proofs for any implementation

We apply the proof technique presented in §6 to other data types discussed in the paper: registers `intreg` and `mvr`.

C.2.1 Lower bound for state-based last-writer-wins register (`intreg`)

THEOREM 24. *If $\mathcal{D}_{\text{intreg}} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{intreg}}$, then $\mathcal{D}_{\text{intreg}}$ is $\widehat{\Omega}(\lg m)$.*

We start by formulating a suitable experiment family.

LEMMA 25. *If $\mathcal{D}_{\text{intreg}} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{intreg}}$, $n \geq 2$ and $m \geq n$, then tuple $(Q, n, m, \mathbb{C}, \mathbf{e}, \text{rb})$ as defined in the left column of Fig. 20 is an experiment family.*

The idea of the experiments is to force replica 1 to remember number of writes dominated by write of the current value, which then introduces an overhead proportional to the $\lg m$,

Encoded object $s \in \mathcal{S}$	Encoding scheme $\text{enc}_{\mathcal{S}}(s)$	Length $\beta(s, \mathcal{S})$ s.t. $\text{len}_{\mathcal{S}}(s) \in \Theta(\beta)$
$u \in U$ with t.o. \leq_U (e.g. \mathbb{N}) s.t. $U^{-1}(u) = \{u' \mid u' \leq_U u\}$	variable length encoding	$1 + \lg U^{-1}(u) $
$(u_1, \dots, u_k) \in U_1 \times \dots \times U_k$	$\text{enc}(u_1, \dots, u_k) = \text{enc}_{\mathbb{N}}(k) \cdot \text{enc}_{U_1}(u_1) \cdot \dots \cdot \text{enc}_{U_k}(u_k)$	$1 + \sum_{i=1}^k \text{len}_{U_i}(u_i)$
$U \subseteq U'$ s.t. $U = \{u_1, \dots, u_k\}$	$\text{enc}(U) = \text{enc}(u_1, \dots, u_k)$	$1 + \sum_{u \in U} \text{len}_{U'}(u)$
$f : U \rightarrow V$ s.t. $U = \{u_1, \dots, u_k\}$	$\text{enc}(f) = \text{enc}(f(u_1), \dots, f(u_k))$	$1 + \sum_{u \in U} \text{len}_V(f(u))$
$f : U \rightarrow V$ s.t. $\text{dom}(f) = \{u_1, \dots, u_k\}$	$\text{enc}(f) = \text{enc}(u_1, f(u_1), \dots, (u_k, f(u_k)))$	$1 + \sum_{u \in \text{dom}(f)} (\text{len}_U(u) + \text{len}_V(f(u)))$

Table 1. Standard encoding schemes for components of replica state and return values and their cost. We note \cdot as a concatenation of strings.

Figure 20. Experiment families $(Q, n, m, \mathbb{C}, \mathbf{e}, \text{rb})$ used in the proofs of Theorem 24 (**intreg**) and Theorem 26 (**mvr**)

intreg	mvr
Conditions on n, m (number of replicas/updates) $m \geq n \geq 2$	
$(m-1) \bmod (n-1) = 0$	
Index set Q	
$Q = [1..m]$	$Q = ([2..n] \rightarrow [1.. \frac{m-1}{n-1}])$
Family size $ Q $	
$ Q = m$	$ Q = (\frac{m-1}{n-1})^{n-1}$
Driver programs	
<pre> procedure <i>init</i> for all $i \in [1..m]$ $u \leftarrow i \bmod 2$ do₂ $\text{wr}(u)^i$ $\text{send}_2(\text{mid}_i)$ for all $r \in [2+1..n]$ do_r rd^{m+r} // dummy procedure <i>exp</i>(α) $\text{receive}_1(\text{mid}_\alpha)$ do₁ rd^{m+n-1} // read \mathbf{e}_α procedure <i>test</i> $u \leftarrow \text{do}_1 \text{rd}^{2m}$ for all $i \in [1..m]$ $\text{receive}_1(\text{mid}_i)$ $u' \leftarrow \text{do}_1 \text{rd}^{2m+i}$ if $u' \neq u$ return $i-1$ return m </pre>	<pre> procedure <i>init</i> for all $r \in [2..n]$ for all $i \in [1.. \frac{m-1}{n-1}]$ do_r $\text{wr}(0)^{rm+i}$ $\text{send}_r(\text{mid}_{r,i})$ procedure <i>exp</i>(α) for all $r \in [2..n]$ $\text{receive}_1(\text{mid}_{r,\alpha(r)})$ do₁ $\text{wr}(1)^{(n+2)m}$ do₁ $\text{rd}^{(n+3)m}$ // read \mathbf{e}_α procedure <i>test</i>(r) for all $i \in [1.. \frac{m-1}{n-1}]$ $\text{receive}_1(\text{mid}_{r,i})$ $u \leftarrow \text{do}_1 \text{rd}^{(n+4)m+i}$ if $u = \{0, 1\}$ return $i-1$ return $\frac{m-1}{n-1}$ </pre>
Definition of executions \mathbb{C}_α	
$\mathbb{C}_\alpha = \text{exec}(\mathcal{D}_\tau, (R_0, T_0), \text{init}; \text{exp}(\alpha))$ where $(R_0, T_0) = ([x \mapsto \mathcal{D}_\tau.\bar{\sigma}_0], \emptyset)$	
Definition of read-back function $\text{rb} : \mathcal{D}_\tau.\Sigma \rightarrow Q$	
$\text{rb}(\sigma) = \text{result}(\mathcal{D}_\tau, (R_{\text{init}}[(x, 1) \mapsto \sigma], T_{\text{init}}), \text{test})$	$\text{rb}(\sigma) = \lambda r : [2..n]. \text{result}(\mathcal{D}_\tau, (R_{\text{init}}[(x, 1) \mapsto \sigma], T_{\text{init}}), \text{test}(r))$
where $(R_{\text{init}}, T_{\text{init}}) = \text{post}(\text{exec}(\mathcal{D}_\tau, (R_0, T_0), \text{init}))$	

the minimum cost of storing a timestamp; cf. the implementation in Fig. 2(c).

All experiments start with a common initialization phase, defined by *init*, where replica 2 performs m of wr operations, writing interchangeably values 0 and 1, starting with 1, and sending a message after each write. For each write operation we define a timestamp such that the execution of driver program consists of the arbitration order consistent with order of operations on replica 2.¹ These writes are followed by dummy reads on replicas 3.. n that only ensure execution uses at least n replicas. All messages send in the *init* phase of the program are undelivered until the second phase, defined by *exp*(α). There replica 1 receives exactly one

¹We ignored timestamp values in driver programs for other data types, because their semantics does not rely on arbitration order.

message from replica 2, selected using α . An experiment concludes with the read \mathbf{e}_α on the first replica.

The read-back works by performing a test defined by *test* program. To determine which message was delivered by replica 1 during the experiment, it first performs a reference read on the first replica, and then forces (re)delivery of all messages sent by replica 2 to replica 1, in the order they were sent, performing a read after each message delivered. By comparing return value of each such read with the reference read, the program identifies the first message that makes the read include an outcome of wr operation dominating in arbitration order all wr operations observed in the reference read; the index of this message corresponds to $\alpha + 1$.

PROOF OF LEMMA 25. Some parts of the proof are straightforward, such as checking that n and m match the number of replicas/updates, and that the driver programs use message identifiers and timestamps correctly. The only nontrivial obligation is to prove $\text{rb}(\text{state}(\mathbf{e}_\alpha)) = \alpha$. Let $(R_\alpha, T_\alpha) = \text{final}(\mathbb{C}_\alpha)$. Then

$$\begin{aligned}
\text{rb}(\text{state}(\mathbf{e}_\alpha)) &= \text{rb}(R_\alpha(x, 1)) \\
&= \text{result}(\mathcal{D}_{\text{intreg}}, (R_{\text{init}}[(x, 1) \mapsto R_\alpha(x, 1)], T_{\text{init}}), \text{test}) \\
&\stackrel{(ii)}{=} \text{result}(\mathcal{D}_{\text{intreg}}, (R_\alpha, T_\alpha), \text{test}) \\
&= \text{result}(\mathcal{D}_{\text{intreg}}, (R_0, T_0), (\text{init}; \text{exp}(\alpha); \text{test})) \stackrel{(i)}{=} \alpha,
\end{aligned}$$

where:

- (i) This is due to $\mathcal{D}_{\text{intreg}} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{intreg}}$, as we explained informally above. Let

$$C'_\alpha = \text{exec}(\mathcal{D}_{\text{intreg}}, (R_0, T_0), (\text{init}; \text{exp}(\alpha); \text{test})).$$

Then the operation context in $\text{abs}(C'_\alpha, \mathcal{V}^{\text{state}})$ of the reference read into u in *test* contains α first wr operations. By $\mathcal{F}_{\text{intreg}}$, u must reflect value of the visible wr that is the latest in the arbitration order $\text{ar}(C'_\alpha)$. By assigned timestamps the arbitration order corresponds to the order of wr operations on replica 2. Hence, $u = \alpha \bmod 2$.

Let u'_i denote value of a read into u' in the i -th iteration of the loop in the program *test*. The operation context of read into u'_i in $\text{abs}(C'_\alpha, \mathcal{V}^{\text{state}})$ contains $\max(i, \alpha)$ first wr operations performed at replica 2. Again, by specification $\mathcal{F}_{\text{intreg}}$ the value must be $u'_i = \max(i, \alpha) \bmod 2$.

By comparing values of u and all u'_i , we obtain:

$$\alpha = \min\{i \mid u \neq u'_{i+1} \vee i = m\}.$$

- (ii) We have $T_{\text{init}} = T_\alpha$ because *exp*(α) does not send any messages. Besides, $R_{\text{init}}[(x, 1) \mapsto R_\alpha(x, 1)]$ and R_α can differ only in the states of the replicas 2.. n . These cannot influence the run of *test*, since it performs events on replica 1 only. \square

PROOF OF THEOREM 24. Given n_0, m_0 , we pick $n = n_0$ and $m = m_0$. Take the experiment family $(Q, n, m, \mathbb{C}, \mathbf{e}, \text{rb})$ given by Lemma 25. Then for any α , return value $\mathbb{C}_\alpha.\text{rval}(\mathbf{e}_\alpha) \in \{0, 1\}$, which can be encoded with a constant length string. Using Lemma 14 for some constants K_1, K independent from n_0, m_0

we get:

$$\begin{aligned} \text{wcmo}(\mathcal{D}_{\text{intreg}}, n, m) &\geq \\ & \lfloor \lg_{|\Lambda|} |Q| \rfloor / (\max_{\alpha \in Q} \text{len}(\mathbb{C}_\alpha.\text{rval}(\mathbf{e}_\alpha))) \geq \\ & K_1 \frac{\lg_{|\Lambda|}(m)}{1} \geq K \lg m. \quad \square \end{aligned}$$

C.2.2 Lower bound for state-based multi-value register (mvr)

THEOREM 26. *If $\mathcal{D}_{\text{mvr}} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{mvr}}$, then \mathcal{D}_{mvr} is $\hat{\Omega}(n \lg m)$.*

LEMMA 27. *If $\mathcal{D}_{\text{mvr}} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{mvr}}$, $n \geq 2$ and $m \geq n$ is such that $(m-1)$ is a multiple of $(n-1)$, then tuple $(Q, n, m, \mathbb{C}, \mathbf{e}, \text{rb})$ as defined in the right column of Fig. 20 is an experiment family.*

The idea of the experiments is to force replica 1 to remember a number of wr operations from each replica dominated by the current wr , which then introduces an overhead proportional to $n \lg m$; cf. the implementation in Fig. 17. All experiments start with a common initialization phase, defined by *init*, where each of the replicas $2..n$ performs $(m-1)/(n-1)$ writes $\text{wr}(0)$ and sends a message after each write. All messages remain undelivered until the second phase, defined by *exp*(α). There replica 1 receives exactly one message from each replica $r = 2..n$, selected using $\alpha(r)$, and subsequently overwrites the register with $\text{wr}(1)$. An experiment concludes with the read \mathbf{e}_α on the first replica.

The read-back works by performing separate tests for each of the replicas $r = 2..n$, defined by *test*(r). To determine which message was sent by replica r during the experiment, the program *test*(r) enforces (re)delivery of all messages sent by replica r to replica 1, in the order they were sent, performing a read after each message delivered. By inspecting the return value of each such read, the program identifies the first message that causes read to return both values 0 and 1, which corresponds to observing an outcome of a $\text{wr}(0)$ performed at r that was not visible at the time of $\text{wr}(1)$, i.e., that is concurrent to $\text{wr}(1)$; the index of this message corresponds to $\alpha(r) + 1$.

PROOF OF LEMMA 27. Some parts of the proof are straightforward, such as checking that n and m match the number of replicas/updates, and that the driver programs use message identifiers and timestamps correctly. The only nontrivial obligation is to prove $\text{rb}(\text{state}(\mathbf{e}_\alpha)) = \alpha$. Let $(R_\alpha, T_\alpha) = \text{final}(\mathbb{C}_\alpha)$. Then

$$\begin{aligned} \text{rb}(\text{state}(\mathbf{e}_\alpha))(r) &= \text{rb}(R_\alpha(x, 1))(r) \\ &= \text{result}(\mathcal{D}_{\text{mvr}}, (R_{\text{init}}[(x, 1) \mapsto R_\alpha(x, 1)], T_{\text{init}}), \text{test}(r)) \\ &\stackrel{\text{(ii)}}{=} \text{result}(\mathcal{D}_{\text{mvr}}, (R_\alpha, T_\alpha), \text{test}(r)) \\ &= \text{result}(\mathcal{D}_{\text{mvr}}, (R_0, T_0), (\text{init}; \text{exp}(\alpha); \text{test}(r))) \stackrel{\text{(i)}}{=} \alpha(r), \end{aligned}$$

where:

- (i) This is due to $\mathcal{D}_{\text{mvr}} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{mvr}}$, as we explained informally above. Let u_i denote value of a read into u in the i -th iteration of the loop in the program *test*(r). By \mathcal{F}_{mvr} the value of u_i is determined by the set and relation of visible writes in the operation context. Let

$$C'_\alpha = \text{exec}(\mathcal{D}_{\text{mvr}}, (R_0, T_0), (\text{init}; \text{exp}(\alpha); \text{test}(r))).$$

Then the operation context in $\text{abs}(C'_\alpha, \mathcal{V}^{\text{state}})$ of the read into u_i in *test*(r) contains:

- (a) first $\max(\alpha(r), i)$ operations $\text{wr}(0)$ from replica r , and
- (b) first $\alpha(r')$ operations $\text{wr}(0)$ from every replica $r' \neq r$, and
- (c) and a single $\text{wr}(1)$ made on replica 1.

We will now analyze the visibility relation between writes $\text{wr}(0)$ and $\text{wr}(1)$. None of writes $\text{wr}(0)$ observed $\text{wr}(1)$ in its operation context in $\text{abs}(C'_\alpha, \mathcal{V}^{\text{state}})$. The operation context of $\text{wr}(1)$ includes first $\alpha(r')$ operations $\text{wr}(0)$ from every replica

r' , including replica r . That means, read into u_i observes $\max(0, i - \alpha(r))$ operations $\text{wr}(0)$ from r that were not visible to $\text{wr}(1)$. By specification \mathcal{F}_{mvr} the return value of the read is in this case:

$$u_i = \begin{cases} \{0, 1\} & \text{if } i - \alpha(r) > 0 \\ \{1\} & \text{otherwise} \end{cases}$$

i.e., 0 appears in the return value whenever there is $\text{wr}(0)$ concurrent to $\text{wr}(1)$ visible. Therefore:

$$\alpha(r) = \min\{i \mid u_{i+1} = \{0, 1\} \vee i = \frac{m-1}{n-1}\}.$$

- (ii) We have $T_{\text{init}} = T_\alpha$ because *exp*(α) does not send any messages. Besides, $R_{\text{init}}[(x, 1) \mapsto R_\alpha(x, 1)]$ and R_α can differ only in the states of the replicas $2..n$. These cannot influence the run of *test*(r), since it performs events on replica 1 only. \square

PROOF OF THEOREM 26. Given n_0, m_0 , we pick $n = n_0$ and some $m \geq n_0$ such that $(m-1)$ is a multiple of $(n-1)$ and $m \geq n^2$. Take the experiment family $(Q, n, m, \mathbb{C}, \mathbf{e}, \text{rb})$ given by Lemma 27. Then for any α , $\mathbb{C}_\alpha.\text{rval}(\mathbf{e}_\alpha) = \{1\}$, which can be encoded with a constant length string. Using Lemma 14 and $m \geq n^2$, for some constants K_1, K_2, K_3, K independent from n_0, m_0 we get:

$$\begin{aligned} \text{wcmo}(\mathcal{D}_{\text{mvr}}, n, m) &\geq \\ & \lfloor \lg_{|\Lambda|} |Q| \rfloor / (\max_{\alpha \in Q} \text{len}(\mathbb{C}_\alpha.\text{rval}(\mathbf{e}_\alpha))) \geq \\ & K_1 \frac{\lg_{|\Lambda|}(\frac{m-1}{n-1})^{n-1}}{\text{len}_{\mathcal{P}(\mathbb{Z})}(\{0\})} \geq K_2 \frac{n \lg(m/n)}{1} \geq \\ & K_3 n \lg \sqrt{m} \geq K n \lg m. \quad \square \end{aligned}$$

C.3 Upper bounds for known implementations

The following upper bounds aim to prove asymptotic worst-case optimality of known object implementations if they match the lower bounds established for any 1-push convergent implementation.

THEOREM 28. *Let $\mathcal{D}_{\text{intreg}}^*$ be the state-based last-writer-wins register implementation defined in Fig. 2(c), such that $\mathcal{D}_{\text{intreg}}^* \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{intreg}}(\S 5.3)$. $\mathcal{D}_{\text{intreg}}^*$ is $\hat{O}(\lg m)$.*

PROOF. Consider any execution $C \in \llbracket \mathcal{D}_{\text{intreg}}^* \rrbracket$ with n replicas, m updates, and any rd event e in this execution (assuming there is at least one).

Recall that the state of a register replica used in the read is a tuple $\langle a, t \rangle = \text{state}(e)$, where a is a register value and t is a timestamp of the value. Practical implementations of LWW-register and replicated stores, often do not rely on integers as timestamps as we simplified so far, but use Lamport clocks [21] instead. Lamport clock is a tuple $t = (k, r) \in \mathbb{N} \times \text{ReplicaID}$ where k is an integer clock and r is an id of replica that generated the tuple. Assuming the optimal Lamport clock implementation is used by the register, such that $k \leq m$, for some constant K_1 , state $\langle a, (k, r) \rangle$ encoding is bounded by: $\text{len}(a, (k, r)) \leq K_1(\text{len}_{\mathbb{Z}}(a) + \lg m + \lg n)$

Note that return value $u = C.\text{rval}(e)$ of rd matches the register value stored in the state (c.f. the definition of rd for $\mathcal{D}_{\text{intreg}}^*$), $u = a$. Therefore, for some constants K_2, K (picked independently of a, n and m) we obtain the following bound on metadata overhead:

$$\frac{\text{len}(a, (k, r))}{\text{len}(u)} \leq K_2 \frac{\text{len}(a) + \lg m + \lg n}{\text{len}(a)} \leq K \lg m$$

which satisfies the \hat{O} definition. \square

THEOREM 29. *Let $\mathcal{D}_{\text{ctr}}^*$ be the state-based counter implementation defined in Fig. 2(b), such that $\mathcal{D}_{\text{ctr}}^* \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{ctr}}(\S 5.3)$. $\mathcal{D}_{\text{ctr}}^*$ is $\hat{O}(n)$.*

PROOF. Consider any execution $C \in \llbracket \mathcal{D}_{\text{ctr}}^* \rrbracket$ with n replicas, m updates, and any rd event e in this execution (assuming there is at least one). Let m' be the number of inc operations visible to the read in the witness abstract execution of this execution, i.e., $m' = |\{e' \mid e' \in \mathcal{V}^{\text{state}}(C)^{-1}(e) \wedge C.\text{oper}(e) = \text{inc}\}|$; clearly, $m' \leq m$.

Recall that the state of a counter replica used in the read is a tuple $\langle r, v \rangle = \text{state}(e)$, where r is a replica id and v is a total function from replica id to the number of increments made at that replica. By definitions of inc and receive, the two operations that mutate component v , $\forall s \in \text{ReplicaID}. v(s) \leq m'$. Therefore, for some constants K_1 and K_2 , state $\langle r, v \rangle$ can be encoded as a string bounded by:

$$\text{len}(r, v) \leq K_1(\lg n + n \lg(m' + 1)) \leq K_2 n \lg(m' + 1)$$

By \mathcal{F}_{ctr} , the value of a counter must be m' , hence the return value $v = C.\text{rval}(e)$ encoding is bounded by: $\text{len}_{\mathbb{N}}(v) \geq K_3(\lg(m' + 1))$ for some constant K_3 . Finally, for some constant K , metadata overhead is bounded by:

$$\frac{\text{len}(r, v)}{\text{len}(u)} \leq K \frac{n \lg(m' + 1)}{\lg(m' + 1)} = Kn$$

□

THEOREM 30. Let $\mathcal{D}_{\text{mvr}}^*$ be the state-based multi-value register implementation defined in Fig. 17, such that $\mathcal{D}_{\text{mvr}}^* \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{mvr}}(\S A)$. $\mathcal{D}_{\text{mvr}}^*$ is $\widehat{O}(n \lg m)$.

PROOF. Consider any execution $C \in \llbracket \mathcal{D}_{\text{mvr}}^* \rrbracket$ with n replicas, m updates, and any rd event e in this execution (assuming there is at least one).

Recall that the state of a register replica used in the read is a tuple $\langle r, V \rangle = \text{state}(e)$, where r is a replica id and V is a set of active entries; each entry $(a, v) \in V$ is made of value a and a vector of integers v (a mapping from a replica id to an integer). By definitions of wr and receive, the two operations that mutate entries in V , $\forall v \in V, s \in \text{ReplicaID}. v(s) \leq m$. Let $U(V) = \{a \mid \exists (a, v) \in V\}$. The definition of receive ensures that every a is unique in V , i.e., $|U(V)| = |V|$. Hence, for some constants K_1, K_2 and K_3 , state $\langle r, V \rangle$ can be encoded with a string bounded by:

$$\begin{aligned} \text{len}(r, V) &\leq K_1(\lg n + \sum_{(a,v) \in V} (\text{len}_{\mathbb{Z}}(a) + n \lg m)) \leq \\ &K_2(\lg n + \sum_{a \in U(V)} (\text{len}_{\mathbb{Z}}(a) + n \lg m)) \leq \\ &K_3(n \lg m(1 + \sum_{a \in U(V)} \text{len}_{\mathbb{Z}}(a))) \end{aligned}$$

For some constant K_4 , return value $u = C.\text{rval}(e)$ of rd consumes no less than:

$$\text{len}_{\mathcal{P}(\mathbb{Z})}(u) \geq K_4(1 + \sum_{a \in u} \text{len}_{\mathbb{Z}}(a))$$

From the definition of rd, the set of returned values u matches the set of entries V , i.e., $u = U(V)$. This leads us to the following overhead limit, for some constant K :

$$\frac{\text{len}(r, V)}{\text{len}(u)} \leq K \frac{n \lg m(1 + \sum_{a \in U(V)} \text{len}_{\mathbb{Z}}(a))}{1 + \sum_{a \in U(V)} \text{len}_{\mathbb{Z}}(a)} = Kn \lg m$$

□

The original implementation of multi-value register has a higher upper bound, as shown next.

THEOREM 31. Let \mathcal{D}_{mvr} be the state-based multi-value register implementation defined in Fig. 16, such that $\mathcal{D}_{\text{mvr}} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{mvr}}(\S A)$. \mathcal{D}_{mvr} is $\widehat{O}(n^2 \lg m)$.

PROOF. Consider any execution $C \in \llbracket \mathcal{D}_{\text{mvr}} \rrbracket$ with n replicas, m updates, and any rd event e in this execution (assuming there is at least one).

Recall that the state of a register replica used in the read is a tuple $\langle r, V \rangle = \text{state}(e)$, where r is a replica id and V is a set of active entries; each entry $(a, v) \in V$ is made of value a and a vector of integers v (a mapping from a replica id to an integer). By definitions of wr and receive, the two operations that mutate entries in V , $\forall v \in V, s \in \text{ReplicaID}. v(s) \leq m$. Also, by definitions of wr and receive (alternatively, by specification \mathcal{F}_{mvr}), the number of concurrent entries in the state cannot exceed the number of sources of concurrency: $|V| \leq n$. Let $U(V) = \{a \mid \exists (a, v) \in V\}$. Then, for some constants K_1, K_2 and K_3 , state $\langle r, V \rangle$ can be encoded with a string no bigger than:

$$\begin{aligned} \text{len}(r, V) &\leq K_1(\lg n + \sum_{(a,v) \in V} (\text{len}_{\mathbb{Z}}(a) + n \lg m)) \leq \\ &K_2(\lg n + n \sum_{a \in U(V)} (\text{len}_{\mathbb{Z}}(a) + n \lg m)) \leq \\ &K_3 n^2 \lg m(1 + \sum_{a \in U(V)} \text{len}_{\mathbb{Z}}(a)) \end{aligned}$$

For some constant K_4 , return value $u = C.\text{rval}(e)$ of rd consumes no less than:

$$\text{len}_{\mathcal{P}(\mathbb{Z})}(u) \geq K_4(1 + \sum_{a \in u} \text{len}_{\mathbb{Z}}(a))$$

From the definition of rd, we observe $u = U(V)$. This leads to the overhead limit that holds for some constant K :

$$\frac{\text{len}(r, V)}{\text{len}(u)} \leq K \frac{n^2 \lg m(1 + \sum_{a \in U(V)} \text{len}_{\mathbb{Z}}(a))}{1 + \sum_{a \in U(V)} \text{len}_{\mathbb{Z}}(a)} = Kn^2 \lg m$$

□

THEOREM 32. Let $\mathcal{D}_{\text{orset}}^*$ be the optimized state-based OR-set implementation defined in Fig. 7, such that $\mathcal{D}_{\text{orset}}^* \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{orset}}(\S 5.3)$. $\mathcal{D}_{\text{orset}}^*$ is $\widehat{O}(n \lg m)$.

PROOF. Consider any execution $C \in \llbracket \mathcal{D}_{\text{orset}}^* \rrbracket$ with n replicas, m updates, and any rd event e in this execution (assuming there is at least one).

Recall that the state of a set replica at the time of read is a tuple $\langle r, V, w \rangle = \text{state}(e)$, where r is a replica id, V is a map of active set entries and w is a version vector, a map from replica id to integers. Each entry $(a, r) \in \text{dom}(V)$ is made of a value a and replica id r , mapped to epoch number, an integer $V(a, r)$. For sake of efficient encoding, we apply the optimization of Bieniusa et al. [6], and assume that V is undefined for (a, r) if $V(a, r) = 0$. Therefore, we treat V as a partial function during encoding. By definitions of add, remove and receive, the three operations that mutate the map of entries V , $\forall (a, r) \in \text{dom}(V). V(a, r) \leq m$. Similarly, by definition of add and receive, the two operations that mutate vector w , $\forall r. w(r) \leq m$. Let $U(V) = \{a \mid \exists (a, r) \in \text{dom}(V)\}$. Then, for some constants K_1 and K_2 , component V of the state can be encoded with a string no bigger than:

$$\begin{aligned} \text{len}(V) &\leq K_1(1 + \sum_{(a,r) \in \text{dom}(V)} (\text{len}_{\mathbb{Z}}(a) + \lg n + \lg m)) \leq \\ &K_2(1 + n \sum_{a \in U(V)} (\text{len}_{\mathbb{Z}}(a) + \lg m)) \end{aligned}$$

This means that, for some constants K_3 and K_4 , complete state occupies at most:

$$\begin{aligned} \text{len}(r, V, w) &\leq K_3(\lg n + n \lg m + n \sum_{a \in U} (\text{len}_Z(a) + \lg m)) \leq \\ &K_4 n \lg m (1 + \sum_{a \in U} \text{len}_Z(a)) \end{aligned}$$

For some constant K_5 , return value $u = C.\text{rval}(e)$ of `rd` consumes no less than:

$$\text{len}_{\mathcal{P}(Z)}(u) \geq K_5 (1 + \sum_{a \in u} \text{len}_Z(a))$$

From the definition of `rd`, we observe: $u = U(V)$. Hence, for some constant K , the overhead is bounded by:

$$\frac{\text{len}(r, V, w)}{\text{len}(u)} \leq K \frac{n \lg m (1 + \sum_{a \in U(V)} \text{len}_Z(a))}{1 + \sum_{a \in U(V)} \text{len}_Z(a)} = K n \lg m$$

□

The original implementation of OR-set has a higher upper bound, as shown next.

THEOREM 33. *Let $\mathcal{D}_{\text{orset}}$ be the original state-based OR-set implementation defined in Fig. 15, such that $\mathcal{D}_{\text{orset}} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{orset}}(\S A)$. $\mathcal{D}_{\text{orset}}$ is $\hat{\Omega}(m \lg m)$.*

PROOF. Consider any execution $C \in \llbracket \mathcal{D}_{\text{orset}} \rrbracket$ with n replicas, m updates, and any `rd` event e in this execution (assuming there is at least one).

Recall that the state of a set replica at the time of read is a tuple $\langle S, D \rangle = \text{state}(e)$, where S is a set of active entries $(a, t) \in S$ with value a and timestamp t , and D is a set of timestamps of removed elements. We first establish some properties of state. By inspection of `add`, `remove`, and `receive`, the only operations mutating sets S and D , the invariant $(a, t) \in S \implies t \notin D$ holds. By definitions of `add` and `receive`, $|S| \leq m$. Similarly, by definitions of `remove` and `receive`, $|D| \leq m$. Moreover, since only `add` generate new timestamps and remaining functions only transfer them between components S and D , $|S| + |D| \leq m$. Let $U(S) = \{a \mid \exists (a, t) \in S\}$. Assuming Lamport clocks as timestamps, for some constants K_1, K_2, K_3 and K_4 , the state can be encoded with at most: $\text{len}(S, D) \leq$

$$\begin{aligned} K_1 (1 + \sum_{(a,t) \in S} (\text{len}_Z(a) + \lg n + \lg m) + |D|(\lg n + \lg m)) &\leq \\ K_2 (1 + m \lg m + \sum_{(a,.) \in S} \text{len}_Z(a)) &\leq \\ K_3 (1 + m \lg m + m \sum_{a \in U(S)} \text{len}_Z(a)) &\leq \\ K_4 m \lg m (1 + \sum_{a \in U(S)} \text{len}_Z(a)) &\leq \end{aligned}$$

For some constant K_5 , return value $u = C.\text{rval}(e)$ of `rd` consumes no less than:

$$\text{len}_{\mathcal{P}(Z)}(u) \geq K_5 (1 + \sum_{a \in u} \text{len}_Z(a))$$

By the definition of `rd`, the set of returned values u matches the active elements, i.e., $u = U(S)$. Hence, for some constant K the overhead is:

$$\frac{\text{len}(S, D)}{\text{len}(u)} \leq K \frac{m \lg m (1 + \sum_{a \in U(S)} \text{len}_Z(a))}{1 + \sum_{a \in U(S)} \text{len}_Z(a)} = K m \lg m$$

□

C.4 Lower bounds for suboptimal implementations

In order to demonstrate an implementation of a data type is suboptimal, we need to show a lower bound higher than the upper bound of the optimal implementation, using counter-example executions.

THEOREM 34. *Let $\mathcal{D}_{\text{orset}}$ be the original state-based OR-set implementation defined in Fig. 15, such that $\mathcal{D}_{\text{orset}} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{orset}}(\S A)$. $\mathcal{D}_{\text{orset}}$ is $\hat{\Omega}(m \lg m)$.*

PROOF. Consider the following driver program defined for any positive n and m such that $m \bmod 2 = 0$:

```

procedure inflate( $n, m$ )
  for all  $i \in [1..(m/2)]$ 
     $\text{do}_1 \text{ add}(i)^{2i}$ 
     $\text{do}_1 \text{ remove}(i)^{2i+1}$ 
   $\text{do}_1 \text{ rd}^{m+2}$  // read  $e^*$ 
  for all  $i \in [2..n]$ 
     $\text{do}_2 \text{ rd}^{m+3}$  // meaningless read on other replicas

```

Given any number of replicas n_0 and any number of updates $m_0 \geq n_0$, we pick $n = n_0$ and $m = 2m_0$ to demonstrate overhead on execution:

$$C = \text{exec}(\mathcal{D}_{\text{orset}}, ([x \mapsto \mathcal{D}_{\text{orset}}.\bar{\sigma}_0], \emptyset), \text{inflate}(n, m))$$

Clearly, the execution has n replicas and m updates. The execution is well-defined since the driver program uses message identifiers and timestamps correctly. Let $\langle S, D \rangle = \text{state}(e^*)$ be the state of replica 1 at the time of read e^* , where S is a set of active entries and D is a set of timestamps of deleted elements. By definitions of `add` and `remove` in $\mathcal{D}_{\text{orset}}$, the set of active elements is empty for this execution, $S = \emptyset$, whereas the set of deleted elements contains tombstones for all removed elements, $|D| = m/2$. Assuming Lamport clock timestamps, for some constants K_1 and K_2 , state $\langle S, D \rangle$ needs to occupy at least:

$$\text{len}(S, D) \geq K_1 \sum_{i=1}^{m/2} (\lg n + \lg m + \text{len}_Z(i)) \geq K_2 m \lg m$$

By definition of `rd`, the return value u of read e^* is an empty set, $C.\text{rval}(e^*) = \emptyset$, which can be encoded with a constant length string. Therefore, for some constant K the overhead at the time of read e^* fulfills the definition of $\hat{\Omega}$:

$$\text{wcmo}(\mathcal{D}_{\text{orset}}, m, n) \geq \text{mmo}(\mathcal{D}_{\text{orset}}, C) \geq \frac{\text{len}(S, D)}{\text{len}(C.\text{rval}(e^*))} \geq K m \lg m$$

□

THEOREM 35. *Let \mathcal{D}_{mvr} be the original state-based multi-value register implementation defined in Fig. 16, such that $\mathcal{D}_{\text{mvr}} \text{ sat}[\mathcal{V}^{\text{state}}, \text{T-Any}] \mathcal{F}_{\text{mvr}}(\S A)$. \mathcal{D}_{mvr} is $\hat{\Omega}(n^2 \lg m)$.*

PROOF. Consider the following driver program defined for any positive n and m such that $\frac{m-n}{n} \geq 1$ and $(m-n) \bmod n = 0$:

```

procedure inflate( $n, m$ )
  for all  $r \in [1..n]$ 
    for all  $i \in [1..(\frac{m-n}{n})]$ 
       $\text{do}_r \text{ wr}(0)^{r \frac{m-n}{n} + i}$ 
       $\text{send}_r(\text{mid}_{r,0})$ 
    for all  $r \in [1..n]$ 
      for all  $r' \in [1..n] \setminus \{r\}$ 
         $\text{receive}_{r'}(\text{mid}_{r',0})$ 
       $\text{do}_r \text{ wr}(1)^{2m+r}$ 
       $\text{send}_r(\text{mid}_{r,1})$ 
      if  $r \neq 1$ 
         $\text{receive}_1(\text{mid}_{r,1})$ 
   $\text{do}_1 \text{ rd}^{2m+n+1}$  // read  $e^*$ 

```

Figure 21. Transport layer specifications for validating various consistency axioms. The relation $\text{roo}(C)$ is defined in §7.

$$\begin{aligned}
& (\forall e \in C.E. \forall r, r'. C.\text{act}(e) = \text{send} \wedge C.\text{repl}(e) = r \wedge r \neq r' \\
& \quad \implies \exists f. C.\text{repl}(f) = r' \wedge e \xrightarrow{\text{del}(C)} f) \wedge \\
& (\forall e \in C.E. C.\text{act}(e) = \text{do} \implies \exists f. \text{act}(f) = \text{send} \wedge e \xrightarrow{\text{roo}(C)} f) \\
& \quad \text{(T-Reliable)} \\
& \forall e, f \in C.E. (C.\text{act}(e) = \text{send}) \wedge e \xrightarrow{(\text{roo}(C) \cup \text{del}(C))^+} f \implies \\
& e \xrightarrow{\text{roo}(C)} f \vee \exists g \in C.E. e \xrightarrow{\text{del}(C)} g \xrightarrow{(\text{roo}(C))^*} f \quad \text{(T-PO-Causal)} \\
& \quad (\text{roo}(C) \cup \text{del}(C))^+|_{\text{do}} \subseteq \text{ar}(C) \quad \text{(T-PO-Lamport)} \\
& \forall e, f \in C.E. (C.\text{act}(e) = \text{do}) \wedge e \xrightarrow{(\text{roo}(C) \cup \text{del}(C))^+} f \implies \\
& e \xrightarrow{\text{roo}(C)} f \vee (\exists e', f'. e \xrightarrow{\text{roo}(C)} e' \xrightarrow{\text{del}(C)} f' \xrightarrow{(\text{roo}(C))^*} f \wedge \\
& \quad \neg \exists e''. e \xrightarrow{\text{roo}(C)} e'' \xrightarrow{\text{roo}(C)} e' \wedge C.\text{act}(e'') = \text{send}) \\
& \quad \text{(T-CO-Causal)} \\
& \forall e, f \in C.E. (e, f) \in (\text{ro}(C) \cup \text{del}(C))^+|_{\text{do}} \implies \\
& \quad C.\text{time}(e) < C.\text{time}(f) \quad \text{(T-CO-Lamport)}
\end{aligned}$$

Given any number of replicas n_0 and any number of updates $m_0 \geq n_0$, we pick $n = n_0$ and $m \geq n_0$ such that $m - n$ is divisible by n , $(m - n) \geq n^2$ and $(m - n)^2 \geq m$. We demonstrate overhead on execution:

$$C = \text{exec}(\mathcal{D}_{\text{mvr}}, ([x \mapsto \mathcal{D}_{\text{mvr}}.\bar{\sigma}_0], \emptyset), \text{inflate}(n, m))$$

Clearly, the execution has n replicas and m updates. The execution is well-defined since the driver program uses message identifiers and timestamps correctly. Let $\langle r, V \rangle = \text{state}(e^*)$ be the state of replica 1 at the time of read e^* , where r is a replica id and V is a set of entries (a, v) with a value a and integers vector v (map from replica id to integer). By definition of implementation and program execution, V contains n entries $(1, v)$ with concurrent vectors, such that each vector v has values: $\forall s. v(s) \geq \frac{m-n}{n}$. Therefore, for some constants $K_1, K_2, K_3, K_4, K_5, K_6$ encoded state $\langle r, v \rangle$ string has a length bounded by:

$$\begin{aligned}
\text{len}(r, V) & \geq K_1 \lg n + \sum_{i=1}^n (\text{len}_{\mathbb{Z}}(1) + \sum_{j=1}^n \lg \frac{m-n}{n}) \geq \\
& K_2 n^2 \lg \frac{m-n}{n} \geq K_3 n^2 \lg \sqrt{m-n} \geq K_4 n^2 \lg(m-n) \geq \\
& K_5 n^2 \lg \sqrt{m} \geq K_6 n^2 \lg m
\end{aligned}$$

By definition of rd , the return value u of read e^* is a singleton set, $C.\text{rval}(e^*) = \{1\}$, which can be encoded with a constant length string. Therefore, for some constant K the overhead at the time of read e^* fulfills the definition of $\hat{\Omega}$:

$$\begin{aligned}
\text{wcmo}(\mathcal{D}_{\text{mvr}}, m, n) & \geq \text{mmo}(\mathcal{D}_{\text{mvr}}, C) \geq \frac{\text{len}(r, V)}{\text{len}(C.\text{rval}(e^*))} \geq \\
& K \frac{n^2 \lg m}{1} = K n^2 \lg m
\end{aligned}$$

□

D. Additional material on consistency axioms

D.1 Correspondence between transport layer specifications and consistency axioms

Fig. 21 summarizes the transport layer specifications validating various consistency axioms. In the context of an execution C , we denote by $\cdot|_{\text{do}}$ the restriction to events e with $C.\text{act}(e) = \text{do}$. T-CO-Lamport essentially requires all timestamps to behave like a single system-wide Lamport clock. The T-CO-Causal condition is somewhat more complicated than T-PO-Causal due to the subtlety of stating that a message in an op-based implementation captures an operation, explained in §4.

The following theorem formalizes the correspondence between various transport layer specifications and axioms stated in §7.

THEOREM 36. *Sets W of visibility witnesses and transport layer specifications \mathcal{T} validate axioms X as per the following table:*

	W	\mathcal{T}	X
1	$\{\mathcal{V}^{\text{state}}, \mathcal{V}^{\text{op}}\}$	T-Any	THINAIR
2	$\{\mathcal{V}^{\text{state}}, \mathcal{V}^{\text{op}}\}$	T-Reliable	EVENTUAL
3	$\{\mathcal{V}^{\text{state}}\}$	T-Any	POCV
4	$\{\mathcal{V}^{\text{state}}, \mathcal{V}^{\text{op}}\}$	T-PO-Causal	POCV
5	$\{\mathcal{V}^{\text{state}}, \mathcal{V}^{\text{op}}\}$	T-PO-Lamport	POCA
6	$\{\mathcal{V}^{\text{state}}, \mathcal{V}^{\text{op}}\}$	T-CO-Causal	COCV
7	$\{\mathcal{V}^{\text{state}}, \mathcal{V}^{\text{op}}\}$	T-CO-Lamport	COCA

PROOF. Consider an execution $C \in \llbracket \mathcal{T} \rrbracket$ and \mathbb{V} such that $\{\mathbb{V}(\tau) \mid \tau \in \text{dom}(\mathbb{V})\} \subseteq W$ and let

$$\text{abs}(C, \text{witness}(\mathbb{V})) = (E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar}).$$

We use sameobj , hbo and hb as defined in Fig. 13.

1. From the definitions of $\mathcal{V}^{\text{state}}$ and \mathcal{V}^{op} , we get

$$\text{vis} \subseteq (\text{ro}(C) \cup \text{del}(C))^+|_{\text{do}}. \quad (29)$$

Hence,

$$\text{ro} \cup \text{vis} \subseteq (\text{ro}(C) \cup \text{del}(C))^+|_{\text{do}},$$

which is acyclic.

2. Consider $e \in E$ and let $\text{repl}(e) = r$. By T-Reliable, there exists $g \in E$ such that

$$\begin{aligned}
\text{act}(g) & = \text{send} \wedge e \xrightarrow{\text{roo}(C)} g \wedge \\
& \neg \exists g'. \text{act}(g') = \text{send} \wedge e \xrightarrow{\text{roo}(C)} g' \xrightarrow{\text{roo}(C)} g.
\end{aligned}$$

Then by T-Reliable for every replica $r' \neq r$ there exists $e_{r'}$ such that

$$\text{act}(e_{r'}) = \text{receive} \wedge g \xrightarrow{\text{del}(C)} e_{r'}.$$

Hence, by the definition of $\mathcal{V}^{\text{state}}$ and \mathcal{V}^{op} , for any event f following $e_{r'}$ in $\text{roo}(C)$, we have $e \xrightarrow{\text{vis}} f$. Furthermore, for any event f following e in $\text{roo}(C)$, we also have $e \xrightarrow{\text{vis}} f$. Thus, the set of events $f \in E$ such that $\text{sameobj}(e, f)$ and $\neg(e \xrightarrow{\text{vis}} f)$ is finite.

3. By the definition of $\mathcal{V}^{\text{state}}$, we get:

$$\begin{aligned}
\text{hbo} & = ((\text{ro} \cap \text{sameobj}) \cup \text{vis})^+ \\
& = (\text{ro}(C) \cup \text{del}(C))^+|_{\text{do}} \\
& = \text{vis}.
\end{aligned}$$

4. By (29), we get:

$$\text{hbo} = ((\text{ro} \cap \text{sameobj}) \cup \text{vis})^+ \subseteq (\text{ro}(C) \cup \text{del}(C))^+|_{\text{do}}.$$

Consider $(e, f) \in \text{hbo}$. If $\mathbb{V}(\text{type}(\text{obj}(e))) = \mathcal{V}^{\text{state}}$, then $(e, f) \in \text{vis}$ follows from the above and the definition of $\mathcal{V}^{\text{state}}$. Now consider the case when $\mathbb{V}(\text{type}(\text{obj}(e))) = \mathcal{V}^{\text{op}}$. If $\text{repl}(e) = \text{repl}(f)$, then $(e, f) \subseteq \text{roo}(C) \subseteq \text{vis}$. Otherwise, for some $e' \in C.E$ we

Figure 22. Axioms defining session guarantees over an execution $(E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar})$. We write r_1, r_2 for the composition of binary relations r_1 and r_2 .

Auxiliary relations

sameobj(e, f) \iff obj(e) = obj(f)

Per-object replica order: $\text{roo} = (\text{ro} \cap \text{sameobj})$

Axioms

RYW (Read Your Writes). An operation sees all previous operations at the same replica: $\text{roo} \subseteq \text{vis}$

MR (Monotonic Reads). An operation sees all operations previously seen by the same replica: $(\text{vis}; \text{roo}) \subseteq \text{vis}$

WFRV (Writes Follow Reads in Visibility). Operations are made visible at other replicas after operations on the same object that were previously seen at the same replica: $(\text{vis}; \text{roo}^*; \text{vis}) \subseteq \text{vis}$

WFRA (Writes Follow Reads in Arbitration). Arbitration orders an operation after other operations previously seen at the same replica: $(\text{vis}; \text{roo}^*) \subseteq \text{ar}$

MWV (Monotonic Writes in Visibility). Operations are made visible at other replicas after all previous operations on the same object at the same replica: $(\text{roo}; \text{vis}) \subseteq \text{vis}$

MWA (Monotonic Writes in Arbitration). Arbitration orders an operation after all previous operations at the same replica: $\text{roo} \subseteq \text{ar}$

have

$$(C.\text{act}(e') = \text{send}) \wedge e \xrightarrow{\text{roo}(C)} e' \xrightarrow{(\text{roo}(C) \cup \text{del}(C))^+} f \wedge \\ \neg \exists e''. e \xrightarrow{\text{roo}(C)} e'' \xrightarrow{\text{roo}(C)} e' \wedge (C.\text{act}(e'') = \text{send}).$$

Then by T-PO-Causal, for some $f' \in C.E$ we have

$$e' \xrightarrow{\text{del}(C)} f' \xrightarrow{\text{roo}(C)} f.$$

Then $e \xrightarrow{\text{vis}} f$ by the definition of \mathcal{V}^{op} .

5. Using (29) and T-PO-Lamport, we get:

$$\text{hbo} = ((\text{ro} \cap \text{sameobj}) \cup \text{vis})^+ \\ \subseteq (\text{roo}(C) \cup \text{del}(C))^+ |_{\text{do}} \\ \subseteq \text{ar}.$$

6. We have:

$$\text{hb} = (\text{vis} \cup \text{roo})^+ \subseteq (\text{ro}(C) \cup \text{del}(C))^+ |_{\text{do}}.$$

Consider $(e, f) \in \text{hb} \cap \text{sameobj}$. If $\text{repl}(e) = \text{repl}(f)$, then $(e, f) \subseteq \text{roo}(C) \subseteq \text{vis}$. Otherwise, by T-CO-Causal for some $e', f' \in C.E$ we have

$$e \xrightarrow{\text{roo}(C)} e' \xrightarrow{\text{del}(C)} f' \xrightarrow{\text{roo}(C)} f \wedge \\ \neg \exists e''. e \xrightarrow{\text{roo}(C)} e'' \xrightarrow{\text{roo}(C)} e' \wedge C.\text{act}(e'') = \text{send}.$$

Then $e \xrightarrow{\text{vis}} f$ by the definition of $\mathcal{V}^{\text{state}}$ and \mathcal{V}^{op} .

7. COCA. Using (29) and T-CO-Lamport, we get:

$$\text{hb} \cup \text{ar} = (\text{ro} \cup \text{vis})^+ \cup \text{ar} \\ \subseteq (\text{ro}(C) \cup \text{del}(C))^+ |_{\text{do}} \cup \text{ar}(C) \\ \subseteq \{(e, f) \mid e, f \in C.E |_{\text{do}} \wedge \text{time}(e) < \text{time}(f)\},$$

which is acyclic. \square

D.2 Session guarantees

The axioms RYW–MWA in Fig. 13 formalize *session guarantees*, which are used to define levels of consistency in between basic eventual consistency and per-object causal consistency. The guarantees are due to Terry et al. [34], who defined them in a low-level operational framework. Here we recast them into axioms appropriate for arbitrary replicated data types. However, we have preserved the original terminology, and thus refer to reads and writes in the

names of the axioms. As a test-case for our formalization, we prove that the conjunction of all session guarantees is equivalent to per-object causal consistency².

PROPOSITION 37. POCV is equivalent to the conjunction of RYW, MR, WFRV and MWV. POCA is equivalent to the conjunction of WFRA and MWA.

PROOF. POCV is equivalent to the conjunction of RYW, MR, WFRV and MWV. Pick an abstract execution

$$A = (E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar}).$$

Let roo and hbo be, respectively, the per-object replica order and the per-object causality order, both induced by A . Let

$$r_1 = \text{roo} \cup (\text{vis}; \text{roo}) \cup (\text{vis}; \text{roo}^*; \text{vis}) \cup (\text{roo}; \text{vis}); \\ r_2 = (\text{vis} \cup \text{roo})^+.$$

Then POCV is equivalent to $r_2 \subseteq \text{vis}$, and the conjunction of RYW, MR, WFRV and MWV to $r_1 \subseteq \text{vis}$.

If $r_2 \subseteq \text{vis}$, then $r_1 \subseteq r_2 \subseteq \text{vis}$. We now assume that $r_1 \subseteq \text{vis}$ and show $r_2 \subseteq \text{vis}$. Pick $e, e' \in E$ such that $(e, e') \in r_2$. Since $\text{roo} \subseteq r_1$, if $(e, e') \in \text{roo}$, then $(e, e') \in r_1 \subseteq \text{vis}$, as required.

If $(e, e') \notin \text{roo}$, then there exist $e_1, f_1, \dots, e_n, f_n$ with $n \geq 1$ such that

$$(e, e_1) \in \text{roo}^* \wedge (f_n, e') \in \text{roo}^* \\ \wedge (\forall i \in \{1, \dots, n\}. (e_i, f_i) \in \text{vis}) \\ \wedge (\forall i \in \{1, \dots, n-1\}. (f_i, e_{i+1}) \in \text{roo}).$$

Using our assumption that $r_1 \subseteq \text{vis}$ can prove the desired $(e, e') \in \text{vis}$ as follows:

$$(e, e_1) \in \text{roo}^* \wedge (e_1, f_1) \in \text{vis} \wedge (f_1, e_2) \in \text{roo} \wedge \dots \\ (e_n, f_n) \in \text{vis} \wedge (f_n, e') \in \text{roo}^* \\ \implies (e, e_1) \in \text{roo}^* \wedge (e_1, f_n) \in \text{vis} \wedge (f_n, e') \in \text{roo}^* \\ \implies (e, e_1) \in \text{roo}^* \wedge (e_1, e') \in \text{vis} \\ \implies (e, e') \in \text{vis}.$$

The first implication comes from the fact that $(\text{vis}; \text{roo}^*; \text{vis}) \subseteq r_1$, the second from $(\text{vis}; \text{roo}) \subseteq r_1$, and the third from $(\text{roo}; \text{vis}) \subseteq r_1$.

POCA is equivalent to the conjunction of WFRA and MWA. Consider the abstract execution A . POCA obviously implies WFRA and MWA, so it remains to prove the converse. Consider $(e, f) \in \text{hbo}$. By the definition of hbo , there exist e_1, \dots, e_n with $n \geq 2$ such that

$$e_1 = e \wedge e_n = f \wedge \forall i \in \{1, \dots, n-1\}. (e_i, e_{i+1}) \in (\text{roo} \cup \text{vis}).$$

Since WFRA and MWA hold, the last conjunct above implies that

$$\forall i \in \{1, \dots, n-1\}. (e_i, e_{i+1}) \in \text{ar}.$$

Furthermore, ar is transitive. Hence, we have $(e, f) \in \text{ar}$, as desired. \square

D.3 Comparison with quiescent consistency

We show that our specifications of eventual consistency describe the semantics of a replicated store more precisely than quiescent consistency, as stated by (1). To formalize the latter, assume that all operations are divided into queries (Query) and updates (Upd), and that return values computed by \mathcal{F}_τ are insensitive to queries:

$$\mathcal{F}_\tau(o, E, \text{oper}, \text{vis}, \text{ar}) = \mathcal{F}_\tau(o, E|_u, \text{oper}|_u, \text{vis}|_u, \text{ar}|_u)$$

where $\cdot|_u$ denotes a restriction to update events. The following straightforward proposition shows that even the basic notion of eventual consistency from §7 implies quiescent consistency.

²J. Brzeziński, C. Sobaniec, and D. Wawrzyniak. From session causality to causal consistency. In *PDP*, 2004.

PROPOSITION 38. Consider an abstract execution A containing only finitely many update events and such that $A \models \mathbb{F}$ and $A \models \text{EVENTUAL}$. Then there are only finitely many query events $e \in A.E$ that do not see all updates on the object they are accessing:

$$\neg(\forall f \in A.E. A.\text{oper}(f) \in \text{Update} \wedge \text{sameobj}(e, f) \implies f \xrightarrow{A.\text{vis}} e).$$

Furthermore, all other query events return the same value.

The converse is not true. Consider a program with a counter object x (initially zero), where a client adds 1 to x at one replica, and another client continuously adds 2 to x at another replica until it is odd:

```
x.add(1) || do { x.add(2) } while (x.rd()%2 == 0)
```

Under basic eventual consistency, termination is guaranteed: the axiom **EVENTUAL** guarantees that all but finitely many reads of x must see the $x.\text{add}(1)$ operation, and thus the loop cannot repeat forever. However, under quiescent consistency termination is not guaranteed: since x is updated continuously, the store is not quiescent, and all bets are off.

D.4 Comparison with the C/C++ memory model

Our formalization allows us to compare consistency models implemented by replicated stores and common hardware and programming languages. To this end, in Fig. 23 we specialize the axioms in Fig. 13 and 22 to the case when all objects are of the type `intreg`, which corresponds to the setting of the latter. We also specialize the notion of an abstract execution as follows.

DEFINITION 39. A *shared-memory execution* is a tuple $R = (H, \text{rf}, \text{ar})$, where

- $H = (E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro})$ is a history such that

$$\forall e \in E. \text{type}(\text{obj}(e)) = \text{intreg};$$

- $\text{rf} \subseteq E \times E$ is an acyclic **reads-from relation** such that

$$(\forall e, f \in E. e \xrightarrow{\text{rf}} f \implies \text{obj}(e) = \text{obj}(f) \wedge \text{oper}(e) = \text{wr}(_) \wedge \text{oper}(f) = \text{rd}) \wedge$$

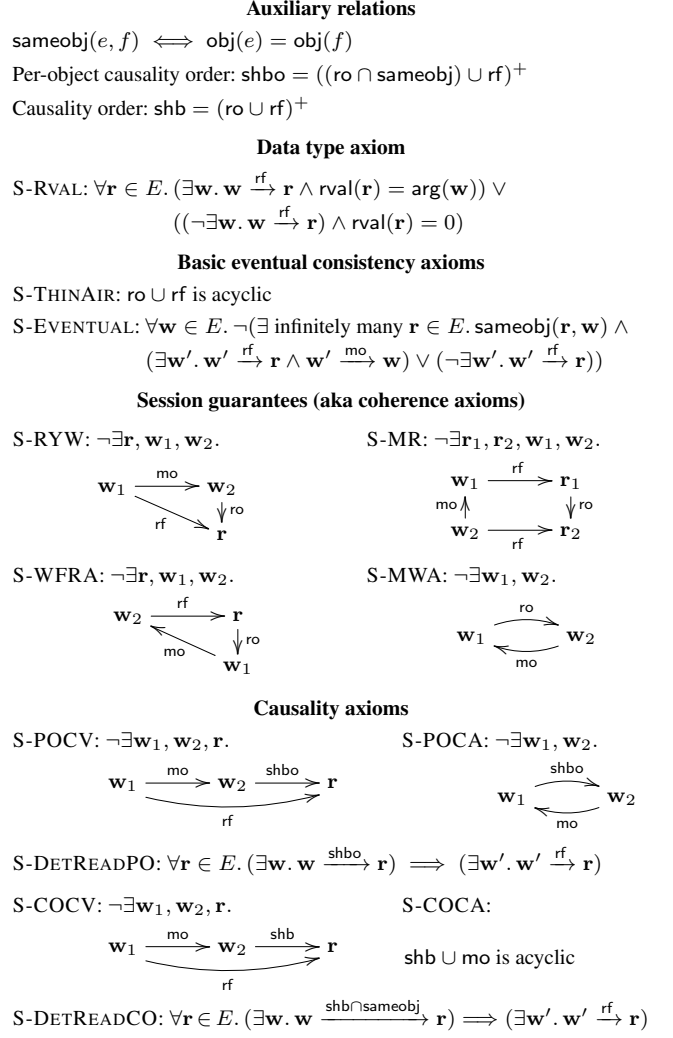
$$(\forall e_1, e_2, f \in E. e_1 \xrightarrow{\text{rf}} f \wedge e_2 \xrightarrow{\text{rf}} f \implies e_1 = e_2);$$

- $\text{mo} \subseteq E \times E$ is a **modification order**, which is a union of transitive, irreflexive and total orders on writes to each object (as defined by `oper` and `obj`).

The reads-from relation and the modification order replace visibility and arbitration, respectively. Unlike `vis`, the `rf` relation does not contain edges with the target being a write event, and contains at most one edge with the target being a given read event; its source corresponds to the write event the read get its return value from. Unlike `ar`, the `mo` relation totally orders only writes to the same object. According to the S-RVAL axiom in Fig. 23, a read returns the value written by the `rf`-related write, or the default value 0 in its absence. The **EVENTUAL** axiom is adjusted to account for using `rf` instead of `vis`: it prohibits infinitely many reads to read from writes preceding any given one in `mo`.

The other axioms in Fig. 23 are formulated in the ‘negative’ form, prohibiting certain configurations in an execution, rather than in the ‘positive’ form, like in Fig. 13, since here this is more concise. For example, S-RYW says that a read r that happened after a write w_2 at the same replica cannot read from a write w_1 earlier in `mo`. This follows from the RYW axiom in Fig. 13: in an execution in the sense of Def. 2 satisfying that axiom, the read r has to see the write w_2 ; then according to $\mathcal{F}_{\text{intreg}}$, r cannot take its value from w_1 . Similarly, the S-COCV axiom says that, if a read r causally happens after a write w_2 , then it cannot read from a write w_1 earlier in `mo`. As before, this follows from S-COCV in Fig. 13, which in this case guarantees that, r sees both writes. There are no

Figure 23. Eventual consistency axioms over a shared-memory execution $(E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{rf}, \text{mo})$. Variables r, r_1, r_2 and w, w_1, w_2 range over read and write events, respectively; `arg` selects the argument of a write.



axioms corresponding to the WFRV and MWV session guarantees, because for integer registers, the session guarantee axioms in Fig. 23 are enough to imply per-object causality.

The resulting consistency model corresponds to that of C/C++ [4, 5] with the caveats explained in §7. In particular, C/C++ has the following weaker analog of S-COCA:

$$\neg \exists w_1, w_2. \begin{array}{c} \xrightarrow{\text{shb}} \\ w_1 \quad w_2 \\ \xleftarrow{\text{mo}} \end{array}$$

This does not allow `mo` to contradict `shb`, but allows `mo` edges for different objects to form cycles with it. The stronger version of COCA is validated by common implementations of causal consistency [26]. Additionally, the release-acquire fragment of C/C++ has an analog of S-DETRREADCO, but the relaxed fragment does not have an analog of S-DETRREADPO.

The correspondence between the axioms for replicated stores in Fig. 13 and 22 and shared-memory axioms in Fig. 23 is formalized as follows.

DEFINITION 40. For an abstract execution

$$A = (E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar})$$

over intreg objects we let the corresponding shared-memory execution be

$$\text{shared}(A) = (E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{rf}, \text{mo}),$$

where

$$\begin{aligned} e \xrightarrow{\text{rf}} f &\iff \text{oper}(e) = \text{wr}(\cdot) \wedge \text{oper}(f) = \text{rd} \wedge e \xrightarrow{\text{vis}} f \wedge \\ &\quad \neg \exists e'. \text{oper}(e') = \text{wr}(\cdot) \wedge e \xrightarrow{\text{ar}} e' \xrightarrow{\text{vis}} f; \\ e \xrightarrow{\text{mo}} f &\iff \text{oper}(e) = \text{wr}(\cdot) \wedge \text{oper}(f) = \text{wr}(\cdot) \wedge e \xrightarrow{\text{ar}} f. \end{aligned}$$

PROPOSITION 41. 1. If A over intreg objects satisfies one of the axioms from Fig. 13, then $\text{shared}(A)$ satisfies the corresponding axioms from Fig. 23, with POCV corresponding to both S-POCV and S-DETRREADPO and COCV corresponding to both S-COCV and S-DETRREADCO. Furthermore, if A satisfies $\mathcal{F}_{\text{intreg}}$, then $\text{shared}(A)$ satisfies S-RVAL.

2. If R satisfies one of the following sets of axioms:

- (a) S-RVAL, S-THINAIR;
- (b) S-RVAL, S-THINAIR, S-POCV, S-POCA, S-DETRREADPO;
- (c) S-RVAL, S-THINAIR, S-COCV, S-COCA, S-DETRREADCO,

then there exists A over intreg objects such that $R = \text{shared}(A)$ and A satisfies

- (a) $\mathcal{F}_{\text{intreg}}$, THINAIR;
- (b) $\mathcal{F}_{\text{intreg}}$, THINAIR, POCV, POCA;
- (c) $\mathcal{F}_{\text{intreg}}$, THINAIR, COCV, COCA.

3. The conjunction of S-RYW, S-MR, S-WFRA, S-MWA is equivalent to that of S-POCV and S-POCA.

PROOF. The proof of item 1 is trivial.

To prove item 2, consider a shared-memory execution

$$R = (E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{rf}, \text{mo}).$$

and let

$$A = (E, \text{repl}, \text{obj}, \text{oper}, \text{rval}, \text{ro}, \text{vis}, \text{ar}),$$

where vis and ar are defined as follows:

- (a) $\text{vis} = \text{rf}$ and ar is any extension of mo to a union of transitive, irreflexive and total orders on events on each object.
- (b) $\text{vis} = \text{shbo}$ and ar is any extension of $\text{mo} \cup \text{shbo}$ to a union of transitive, irreflexive and total orders on events on each object. S-THINAIR implies that vis is well-formed. S-POCA and the totality of mo on writes to the same object imply that $\text{mo} \cup \text{shbo}$ is acyclic, and thus ar can indeed be constructed.
- (c) $\text{vis} = \text{shb} \cap \text{sameobj}$ and $\text{ar} = \overline{\text{ar}} \cap \text{sameobj}$, where $\overline{\text{ar}}$ is any transitive, irreflexive and total order on events in E containing $\text{shb} \cup \text{mo}$. S-COCA implies that both vis is well-formed and $\overline{\text{ar}}$ can indeed be constructed.

In case (b), $\text{hbo} = \text{shbo}$, so that A satisfies POCV and POCA. In case (c), $\text{hb} = \text{shb}$. Then

$$\text{hb} \cap \text{sameobj} = \text{shb} \cap \text{sameobj} \subseteq \text{vis},$$

so that A satisfies COCV. Besides,

$$\text{hb} \cup \text{ar} \subseteq \text{shb} \cup \overline{\text{ar}} = \overline{\text{ar}},$$

which is acyclic. Thus, A satisfies COCA.

In all cases, the definitions of vis imply

$$\text{vis} \subseteq (\text{ro} \cup \text{rf})^+,$$

so that

$$(\text{ro} \cup \text{vis})^+ \subseteq (\text{ro} \cup \text{rf})^+.$$

Hence if R satisfies S-THINAIR, then A satisfies THINAIR.

We now prove that $\text{shared}(A) = R$; this straightforwardly implies that A satisfies $\mathcal{F}_{\text{intreg}}$. Assume that $\text{shared}(A) \neq R$. Since in all cases we have $\text{mo} \subseteq \text{ar}$, the projection of ar onto writes to the same object is equal to mo . Furthermore, we have $\text{rf} \subseteq \text{vis}$ in all cases. Thus, we are left with the following two possibilities.

- $e \xrightarrow{\text{rf}} f$, but for some $e, f, e' \in E$ we have

$$\begin{aligned} \text{oper}(e) = \text{wr}(\cdot) \wedge \text{oper}(f) = \text{rd} \wedge \\ \text{oper}(e') = \text{wr}(\cdot) \wedge e \xrightarrow{\text{ar}} e' \xrightarrow{\text{vis}} f, \end{aligned}$$

which implies $e \xrightarrow{\text{mo}} e'$.

In case (a), $\text{vis} = \text{rf}$, so that the above is impossible by the properties of rf and mo . In case (b), $\text{vis} = \text{shbo}$, so that the above contradicts S-POCV. In case (c), $\text{vis} = \text{shb} \cap \text{sameobj}$, so that the above contradicts S-COCV.

- For some $e, f \in E$ we have

$$\text{oper}(e) = \text{wr}(\cdot) \wedge \text{oper}(f) = \text{rd} \wedge e \xrightarrow{\text{vis}} f,$$

but $\neg \exists e'. e' \xrightarrow{\text{rf}} f$. In case (a), $\text{vis} = \text{rf}$ and the above immediately yields a contradiction. In case (b), $\text{vis} = \text{shbo}$, and we get a contradiction with S-DETRREADPO. In case (c), $\text{vis} = \text{shb} \cap \text{sameobj}$, and we get a contradiction with S-DETRREADCO.

Finally, we prove item 3. It is easy to see that S-POCV implies S-RYW, S-MR and S-POCA implies S-MWA, S-WFRA. Assume now that S-RYW, S-MR, S-MWA, S-WFRA hold. Consider $e, e' \in E$ such that $(e, e') \in \text{shbo} \setminus \text{roo}$. Then there exist $\mathbf{w}_1, \mathbf{r}_1, \dots, \mathbf{w}_n, \mathbf{r}_n$ with $n \geq 1$ such that

$$\begin{aligned} (e, \mathbf{w}_1) \in \text{roo}^* \wedge (\mathbf{r}_n, e') \in \text{roo}^* \\ \wedge (\forall i \in \{1, \dots, n\}. (\mathbf{w}_i, \mathbf{r}_i) \in \text{vis}) \\ \wedge (\forall i \in \{1, \dots, n-1\}. (\mathbf{r}_i, \mathbf{w}_{i+1}) \in \text{roo}). \end{aligned}$$

Since mo is total on writes to the same object, by S-MR we get

$$\begin{aligned} (e, \mathbf{w}_1) \in \text{roo}^* \wedge (\mathbf{w}_1, \mathbf{r}_1) \in \text{rf} \wedge (\mathbf{r}_1, \mathbf{w}_2) \in \text{roo} \wedge \dots \\ (\mathbf{w}_n, \mathbf{r}_n) \in \text{rf} \wedge (\mathbf{r}_n, e') \in \text{roo}^* \\ \implies (e, \mathbf{w}_1) \in \text{roo}^* \wedge (\mathbf{w}_1, \mathbf{r}_n) \in \text{rf} \wedge (\mathbf{r}_n, e') \in \text{roo}^*. \end{aligned}$$

Hence, $\text{shbo} = \text{roo} \cup (\text{roo}^*; \text{rf}; \text{roo}^*)$.

To prove S-POCV, assume there exist $\mathbf{w}_1, \mathbf{w}_2, \mathbf{r} \in E$ such that

$$\begin{array}{ccccc} \mathbf{w}_1 & \xrightarrow{\text{mo}} & \mathbf{w}_2 & \xrightarrow{\text{shbo}} & \mathbf{r} \\ & & & \searrow & \\ & & & \text{rf} & \end{array}$$

If $\mathbf{w}_2 \xrightarrow{\text{roo}} \mathbf{r}$, then this contradicts S-RYW. Otherwise, for some $\mathbf{w}', \mathbf{r}' \in E$ we have

$$\begin{array}{ccccccc} \mathbf{w}_1 & \xrightarrow{\text{mo}} & \mathbf{w}_2 & \xrightarrow{\text{roo}^*} & \mathbf{w}' & \xrightarrow{\text{rf}} & \mathbf{r}' & \xrightarrow{\text{roo}^*} & \mathbf{r} \\ & & & & & \searrow & & & \\ & & & & & \text{rf} & & & \end{array}$$

Then by S-MWA and the irreflexivity of mo we get

$$\begin{array}{ccccccc} \mathbf{w}_1 & \xrightarrow{\text{mo}} & \mathbf{w}' & \xrightarrow{\text{rf}} & \mathbf{r}' & \xrightarrow{\text{roo}} & \mathbf{r} \\ & & & \searrow & & & \\ & & & \text{rf} & & & \end{array}$$

which contradicts S-MR.

To prove S-POCA, assume there exist $\mathbf{w}_1, \mathbf{w}_2 \in E$ such that

$$\begin{array}{ccc} & \xrightarrow{\text{shbo}} & \\ \mathbf{w}_1 & & \mathbf{w}_2 \\ & \xleftarrow{\text{mo}} & \end{array}$$

If $w_1 \xrightarrow{roo} w_2$, then this contradicts S-MWA. Otherwise, for some $w', r' \in E$ we have

$$w_1 \xrightarrow{roo^*} w' \xrightarrow{rf} r' \xrightarrow{roo} w_2$$

\xleftarrow{mo}

By S-MWA, this entails

$$w' \xrightarrow{rf} r' \xrightarrow{roo} w_2$$

\xleftarrow{mo}

which contradicts S-WFRA.

□