

Automatically Verifying Replication-aware Linearizability

VIMALA SOUNDARAPANDIAN, IIT Madras, India

KARTIK NAGAR, IIT Madras, India

ASEEM RASTOGI, Microsoft Research, India

KC SIVARAMAKRISHNAN, IIT Madras and Tarides, India

Data replication is crucial for enabling fault tolerance and uniform low latency in modern decentralized applications. Replicated Data Types (RDTs) have emerged as a principled approach for developing replicated implementations of basic data structures such as counter, flag, set, map, etc. While correctness of RDTs is generally specified using the notion of strong eventual consistency—which guarantees that replicas which have received the same set of updates would converge to the same state—a more expressive specification which relates the converged state to updates received at a replica would be more beneficial to RDT users. Replication-aware linearizability is one such specification, which requires all replicas to always be in a state which can be obtained by linearizing the updates received at the replica. In this work, we develop a novel fully automated technique for verifying replication-aware linearizability for Mergeable Replicated Data Types (MRDTs). We identify novel algebraic properties for MRDT operations and the merge function which are sufficient for proving an implementation to be linearizable and which go beyond the standard notions of commutativity, associativity and idempotence. We also develop a novel inductive technique called bottom-up linearization to automatically verify the required algebraic properties. Our technique can be used to verify both MRDTs and state-based CRDTs. We have successfully applied our approach on a number of complex MRDT and CRDT implementations including a novel JSON MRDT.

1 Introduction

Modern decentralized applications often employ data replication across geographically distributed locations to enhance fault tolerance, minimize data access latency and improve scalability. This practice is crucial for mitigating the impact of network failures and reducing data transmission delays to end users. However, these systems encounter the challenge of concurrent conflicting data updates across different replicas.

Recently, Mergeable Replicated Data Types (MRDTs) [11, 12, 23] have emerged as a systematic approach to the problem of ensuring that replicas remain eventually consistent despite concurrent conflicting updates. MRDTs draw inspiration from the Git version control system, where each update creates a new version and any two versions can be merged explicitly through a user-defined merge function. `merge` is a ternary function which takes as input the two versions to be merged and their Lowest Common Ancestor (LCA), i.e., the most recent version from which the two versions diverged. As opposed to Conflict-Free Replicated Data Types (CRDTs)[21] which may have to carry around causal context metadata to ensure consistency, MRDTs can rely on the underlying system model to provide the causal context through the LCA. This results in implementations that are comparatively simpler and also more efficient. For example, if we consider state-based CRDTs, which are the closest analogue to the MRDT model, then any counter CRDT implementation would require $O(n)$ space, where n is the number of replicas (a lower bound proved by [4]), whereas a counter MRDT implementation only requires $O(1)$ space. The states maintained by CRDT implementations need to form a join semi-lattice, with all CRDT operations restricted to being monotonic functions and merge restricted to the lattice join. While these restrictions simplify the task of reasoning about correctness [5, 13, 18], crafting correct and efficient CRDT implementations itself becomes much harder.

MRDTs do not require any of the above restrictions, which helps in developing implementations with better space and time complexity. However, reasoning about correctness now becomes harder. Indeed, the MRDT system model allows arbitrary replicas to merge their states at arbitrary points of time, and this can result in subtle bugs requiring a very specific orchestration of merge actions. As part of this work, we discovered such subtle bugs in MRDT implementations claimed to be verified by previous works [23] (more details can be found in §5.2). The MRDT state as well as the implementation of data type operations and the merge function have to be cleverly designed to ensure strong eventual consistency. That is, despite concurrent conflicting updates and arbitrary ordering of merges, all replicas will eventually converge to the same state. Further, we would also like to show that an MRDT satisfies the functional behavior of the data type, along with the user-defined conflict resolution policy for concurrent conflicting updates (e.g., for a set data type, an *add-wins* policy which favors the add operation over a concurrent remove of the same element at different replicas). There have been few works [11, 12, 23] which have looked at the problem of specifying and verifying MRDTs. However, they either restrict the system model by disallowing concurrent merges [12], focus only on convergence as the correctness specification [11, 12], or do not support automated verification [23].

In this work, we couch correctness of MRDTs using the notion of *Replication-Aware Linearizability* (RA-linearizability) [25], which says that the state at any replica must be obtained by linearizing (i.e., constructing a sequence of) update operations that have been applied at the replica. As a first contribution, we adapt RA-linearizability to the MRDT system model (§3), and develop a simple specification framework for MRDTs based on conflict resolution policy for concurrent update operations. We show that an MRDT implementation can be linearized only under certain technical constraints on the conflict resolution policy and if the merge operation satisfies a weaker notion of commutativity called *conditional commutativity*. By ensuring that the linearization order obeys the conflict resolution policy for concurrent update operations and it remains the same across all replicas, we guarantee both strong eventual consistency and adherence to the user-provided specification.

Next, we propose a sound but not complete technique for proving RA-linearizability for MRDT implementations. The main challenge lies in showing that the merge function generates a state which is a linearization of its inputs. We develop a technique called *bottom-up linearization*, which relies on certain simple algebraic properties of the merge function to prove that it generates the correct linearization. We then design an induction scheme to *automatically* verify the required algebraic properties of merge for an arbitrary MRDT implementation. Our main insight here is to leverage the fact that the merge inputs are themselves linearizations, and hence we can use induction over their operation sequences. We extract a set of verification conditions (VCs) which are amenable to automated reasoning, and prove that if an MRDT implementation satisfies the VCs, it is linearizable (§4). While our development is focussed on MRDTs, our technique can be directly applied on state-based CRDTs. State-based CRDTs also have a merge-based system model which is slightly simpler than MRDTs as the merge function does not require any LCA.

Finally, we develop a framework in the F* [24] programming language that allows implementing MRDTs and automatically mechanically proving the VCs required by our technique. The framework provides several advantages over previous works. First, we require the programmer to specify only the MRDT operations, the merge function, and the conflict resolution policy, in contrast to the earlier work that also requires proof constructs such as abstract simulation relations [23]. Second, the VCs are simple enough that in *all* the case studies we have done, including data types such as counter, set, map, boolean flag, and list, they are automatically discharged by F*. Finally, we extract the verified implementations to OCaml using the F* extraction pipeline and run them (§5).

We have also implemented and verified few state-based CRDTs using our framework. In the next section, we present the main ideas of our work informally through a series of examples.

2 Overview

2.1 System Model

The MRDT system model resembles a distributed version control system, such as Git [6], with replication centred around versioned states in branches and explicit merges. A replicated data store handles multiple objects independently [9, 19]; in our presentation, we focus on modelling a store with a single object. The state of the object is replicated across multiple replicas $r_1, r_2, \dots \in \mathcal{R}$ in the store. Clients interact with the store by performing query or update operations on one of the replicas, with update operations modifying its state. These replicas operate concurrently, allowing independent modifications without synchronization. They periodically (and non-deterministically) exchange updates with each other through a process called *merge*. Due to concurrent operations happening at multiple replicas, conflicts may arise, which must be resolved by the merge operation in an appropriate and consistent manner. An object has a type $\tau \in \text{Type}$, whose type signature $\langle O_\tau, Q_\tau, Val_\tau \rangle$ contains the set of supported update operations O_τ , query operations Q_τ and their return values Val_τ .

Definition 2.1. A MRDT implementation for a data type τ is a tuple $\mathcal{D}_\tau = \langle \Sigma, \sigma_0, \text{do}, \text{merge}, \text{query}, \text{rc} \rangle$, where:

- Σ is the set of states, $\sigma_0 \in \Sigma$ is the initial state.
- $\text{do} : \Sigma \times \mathcal{T} \times \mathcal{R} \times O_\tau \rightarrow \Sigma$ implements all update operations in O_τ , where \mathcal{T} is the set of timestamps.
- $\text{merge} : \Sigma \times \Sigma \times \Sigma \rightarrow \Sigma$ is a three-way merge function.
- $\text{query} : \Sigma \times Q_\tau \rightarrow Val_\tau$ implements all query operations in Q_τ , returning a value in Val_τ .
- $\text{rc} \subseteq O_\tau \times O_\tau$ is the conflict resolution policy to be followed for concurrent update operations.

An MRDT \mathcal{D}_τ provides implementations of *do*, *merge* and *query* which will be invoked by the data store appropriately. A client request to perform an update operation $o \in O_\tau$ at a replica r triggers the call $\text{do}(\sigma, t, r, o)$. This takes as input the current state $\sigma \in \Sigma$ of r , a unique timestamp $t \in \mathcal{T}$ and produces an updated state which is then installed at r . The data store ensures that timestamps are unique across all operations (which can be achieved through e.g. Lamport timestamps [14]).

Replicas can also receive states from other replicas, which are merged with the receiver's state using *merge*. The *merge* function is called with the current states of both the sender and receiver replicas and their lowest common ancestor (LCA), which represents the most recent common state from which the two replicas diverged. Clients can query the state of the MRDT using the *query* method. This takes a MRDT state $\sigma \in \Sigma$ and a query operation as input and produces a return value. Note that a query operation cannot change the state at a replica.

While merging, it may happen that conflicting update operations may have been performed on the two states, in which case, the implementation also provides a conflict resolution policy *rc*. The merge function should make sure that this policy is followed while computing the merged state. To illustrate, we now present a couple of MRDT implementations: an increment-only counter and an observed-remove set.

- 1: $\Sigma = \mathbb{N}$
- 2: $O = \{\text{inc}\}$
- 3: $Q = \{\text{rd}\}$
- 4: $\sigma_0 = 0$
- 5: $\text{do}(\sigma, _, _, \text{inc}) = \sigma + 1$
- 6: $\text{merge}(\sigma_\top, \sigma_1, \sigma_2) = \sigma_1 + \sigma_2 - \sigma_\top$
- 7: $\text{query}(\sigma, \text{rd}) = \sigma$
- 8: $\text{rc} = \emptyset$

Fig. 1. Counter MRDT implementation

99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147

The counter MRDT implementation is given in Fig. 1. The state space of the counter MRDT is simply the set of natural numbers, and it allows clients to perform only one update operation (*inc*) which increments the value of the counter. For merging two counter states σ_1 and σ_2 , whose lowest common ancestor is σ_\top , intuitively, we want to find the total number of increment operations across σ_1 and σ_2 . Since σ_\top already accounts for the effect of the common increments in σ_1 and σ_2 , we need to count the newer increments and then add them to σ_\top . This is achieved by adding $\sigma_1 - \sigma_\top$ and $\sigma_2 - \sigma_\top$ to σ_\top , which simplifies to the merge definition in Fig. 1. For example, suppose we have replicas r_1 and r_2 whose initial state was $\sigma_\top = 2$. Now, if there are 2 *inc* operations at r_1 and 3 *inc* operation at r_2 , their states will be $\sigma_1 = 4$ and $\sigma_2 = 5$. On merging r_2 at r_1 , $\text{merge}(\sigma_\top, \sigma_1, \sigma_2)$ will return 7, which reflects the total number of increments. The query method simply returns the current state of the counter. Finally, the increment operation commutes with itself, so there is no need to define a conflict resolution policy.

An observed-remove set (OR-set) [21] is an implementation of a set data type which employs an add-wins conflict-resolution strategy, prioritizing addition in cases of concurrent addition and removal of the same element. Fig. 2 shows the OR-set MRDT implementation. This implementation is quite similar to the operation-based (op-based) CRDT implementation of OR-set [22]. The state of the OR-set is a set of element-timestamp pairs, with the initial state being an empty set. Clients can perform two operations for every element $a \in \mathbb{E}$: add_a and rem_a . The add_a method adds the element a along with the (unique) timestamp at which the operation was performed. The rem_a method removes all entries in the set corresponding

- 1: $\Sigma = \mathcal{P}(\mathbb{E} \times \mathcal{T})$
- 2: $O = \{\text{add}_a, \text{rem}_a \mid a \in \mathbb{E}\}$
- 3: $Q = \{\text{rd}\}$
- 4: $\sigma_0 = \{\}$
- 5: $\text{do}(\sigma, t, _, \text{add}_a) = \sigma \cup \{(a, t)\}$
- 6: $\text{do}(\sigma, _, _, \text{rem}_a) = \sigma \setminus \{(a, i) \mid (a, i) \in \sigma\}$
- 7: $\text{merge}(\sigma_\top, \sigma_1, \sigma_2) =$
 $(\sigma_\top \cap \sigma_1 \cap \sigma_2) \cup (\sigma_1 \setminus \sigma_\top) \cup (\sigma_2 \setminus \sigma_\top)$
- 8: $\text{query}(\sigma, \text{rd}) = \{a \mid (a, _) \in \sigma\}$
- 9: $\text{rc} = \{(\text{rem}_a, \text{add}_a) \mid a \in \mathbb{E}\}$

Fig. 2. OR-set MRDT implementation

to the element a . An element a is considered to be present in the set if there is some (a, t) in the state.

The merge method takes as input the LCA set σ_\top and the two sets σ_1 and σ_2 to be merged, retains elements of σ_\top that were not removed in both sets, and includes the newly added elements from both sets. Since σ_\top is the most recent state from which the two sets diverged, the intersection of all three sets is the set of elements that were not removed from σ_\top in either branch, while the difference of either set with the σ_\top corresponds to the newly added elements. The query operation rd returns all the elements in the set. The conflict resolution relation rc orders rem_a before add_a of the same element in order to achieve the add-wins semantics. Note that all other pairs of operations ($\text{add}__$ and $\text{add}__$, $\text{rem}__$ and $\text{rem}__$, and add_x and rem_y with $x \neq y$) commute with each other, hence rc does not specify their ordering. We now consider whether the merge operation adheres to the conflict resolution policy.

2.2 RA-linearizability for MRDTs

We would like to verify that an MRDT implementation is correct, in the sense that in every execution, (a) replicas which have observed the same set of update operations converge to the same state, and (b) this state reflects the semantics of the implemented data type and the conflict resolution policy. Note that an update operation o is considered to be visible to a replica r either if o is directly applied by a client at r , or indirectly through merge with another replica r' on which o was visible. To specify MRDT correctness, we propose to use the notion of RA-linearizability [25]: the state at any replica during any execution must be achievable by applying a sequence (or linearization) of the update operations visible to the replica. Further, this linearization should obey the user-specified

197 conflict resolution policy for concurrent operations, and the local replica order for non-concurrent
 198 operations.


199 Our definition of RA-linearizability allows viewing the state of an MRDT replica as a sequence
 200 of update operations applied on the initial state, thus abstracting over the merge function and how
 201 it handles concurrent operations. Consequently, any formal reasoning (e.g. assertion checking,
 202 functional correctness, equivalence checking etc.) can now essentially forget about the presence of
 203 merges, and only focus on update operations, with the additional guarantee that operations would
 204 have been correctly linearized taking into account the conflict resolution policy and local replica
 205 ordering.

206 Proving RA-linearizability for MRDTs is straightforward when there is only a single replica on
 207 which all operations are performed, since there is no interleaving among operations on a single
 208 replica. Complexity arises when update operations happen concurrently across replicas, which
 209 are then merged. For a merge operation, we need to show that the output can be obtained by
 210 applying a linearization of update operations witnessed by both replicas being merged. However,
 211 the states being merged would have been obtained after an arbitrary number of update operations
 212 or even other merges. Further, the MRDT framework maintains only the states, but not the update
 213 operations leading to those states, thus requiring the verification technique to somehow infer the
 214 update operations leading to a state, and then show that merge constructs the correct linearization.

215 We break down this difficult problem gradually with a series of observations. We will start with
 216 an intuitively correct approach, show how it could be broken through examples, and gradually
 217 refine it to make it work. As a starting point, we first observe that we can leverage the following
 218 algebraic properties of the MRDT update operations and the merge function: (i) commutativity
 219 of merge and update operations, (ii) commutativity of merge, (iii) idempotence of merge, and (iv)
 220 commutativity of update operations. To motivate this, we first introduce some terminology. An
 221 event $e = \langle t, r, o \rangle$ is generated for every update operation instance, where t is the event's timestamp
 222 and r is the replica on which the update operation o is applied. Applying an event e on a replica with
 223 state σ changes the replica state to $e(\sigma) = \text{do}(\sigma, t, r, o)$ using the implementation of the operation o .
 224 Given a sequence of events $\pi = e_1 e_2 \dots e_n$, we use the notation $\pi(\sigma)$ to denote $e_n(\dots(e_2(e_1(\sigma))))$.
 225 Now, the properties described above can be formally defined as follows (forall $\sigma_\top, \sigma_1, \sigma_2, e, e'$):
 226

- 227 (P1) $\text{merge}(\sigma_\top, e(\sigma_1), \sigma_2) = e(\text{merge}(\sigma_\top, \sigma_1, \sigma_2))$
 228 (P2) $\text{merge}(\sigma_\top, \sigma_1, \sigma_2) = \text{merge}(\sigma_\top, \sigma_2, \sigma_1)$
 229 (P3) $\text{merge}(\sigma_\top, \sigma_\top, \sigma_\top) = \sigma_\top$
 230 (P4) $e(e'(\sigma)) = e'(e(\sigma))$
 231

232 As per our proposed definition of RA-linearizability, we need to show that there exists a lineariza-
 233 tion of events visible at the replica such that the state of the replica can be obtained by applying
 234 this linearization. As mentioned earlier, an event can become visible at a replica either by a direct
 235 client application, or by merging with another replica. To illustrate this, consider the scenario
 236 shown in Fig. 3 where two replicas with states σ_1 and σ_2 are be-
 237 ing merged. These states were obtained by applying a sequence
 238 of events π_1 and π_2 respectively on the LCA state σ_\top . We call the
 239 events in π_1 and π_2 as local to their respective replicas. Now, when
 240 the two states are merged to create a new state σ_m we would need to
 241 show that the state $\sigma_m (= \text{merge}(\sigma_\top, \sigma_1, \sigma_2))$ can be obtained by linear-
 242 izing all the events in π_1 and π_2 , and applying this linearization
 243 on the state σ_\top .
 244



245 Fig. 3. Linearizing a merge operation

246 To show that the merge function constructs a linearization, we
 247 can take advantage of properties (P1)-(P4). In particular, com-
 248 mutativity of merge and update operation application (P1) al-
 249 lows us to move an event from the second argument of merge
 250 to outside, and we can then repeatedly apply this property to peel off all the events in π_1 .
 251 More formally, by performing induction on the sequence π_1 and using (P1), we can show that
 252 $\text{merge}(\sigma_\top, \pi_1(\sigma_\top), \sigma_2) = \pi_1(\text{merge}(\sigma_\top, \sigma_\top, \sigma_2))$. We can then use commutativity of merge (P2) to
 253 swap the last two arguments of merge, and then apply (P1) again to peel off all the events in π_2 , thus
 254 establishing that $\text{merge}(\sigma_\top, \sigma_\top, \pi_2(\sigma_\top)) = \pi_2(\text{merge}(\sigma_\top, \sigma_\top, \sigma_\top))$. Finally, using merge idempot-
 255 ence (P3), and combining all the previous results, we can infer that $\text{merge}(\sigma_\top, \sigma_1, \sigma_2) = \pi_2(\pi_1(\sigma_\top))$.
 256 Commutativity of update operations (P4) ensures that all linearizations of events in π_1 and π_2 lead
 257 to the same state, thus ratifying the specific linearization order $\pi_1\pi_2$ that we constructed using
 258 properties P1-P3. We call this process as bottom-up linearization, since we built the sequence from
 259 end through property (P1), linearizing one event at a time.

260 It is also easy to see that the counter MRDT implementation in
 261 Fig. 1 satisfies (P1)-(P4). In particular, commutativity of integer addi-
 262 tion and subtraction essentially gives us (P1)-(P4) for free. While
 263 this strategy works for the counter MRDT, commutativity of all
 264 update operations is in general a very strong requirement, and
 265 would fail for other datatypes. For example, the OR-set MRDT of
 266 Fig. 2 does not satisfy (P4), as the add_a and rem_a operations do not
 267 commute.

268 In the presence of non-commutative update operations, the prop-
 269 erty (P1) now needs to be altered, as we need to consider the conflict
 270 resolution policy to decide the replica from which an event needs to be peeled off. To illustrate this,
 271 consider an OR-set execution depicted in Fig. 4. We show the version graph of the execution, where
 272 each oval represents a version. The state of the version is depicted inside the oval. The versions
 273 v_1 and v_2 are obtained by applying rem_a and add_a operations to the version v_\top on two different
 274 replicas (r_1 and r_2). Each edge is labeled with the event corresponding to the application of an
 275 operation. Let $\sigma_\top = \{\}$ denote the state of the LCA v_\top . The versions v_1 and v_2 are then merged at
 276 r_2 which gives rise to a new version v_m with state $\text{merge}(\sigma_\top, e_1(\sigma_\top), e_2(\sigma_\top))$. Now, since e_1 and
 277 e_2 do not commute, the conflict resolution policy of OR-set places e_1 (i.e. the remove operation)
 278 before e_2 (i.e. the add operation). Hence, we want the merged version to follow the linearization
 279 order $e_2(e_1(\sigma_\top))$. This requires us to first peel off the event e_2 from the third argument of merge.
 280 To achieve this, we can alter the property (P1) by making it aware of the conflict resolution policy
 281 as follows:

$$282 \text{ (P1')} \quad (e_1, e_2) \in \text{rc} \implies \text{merge}(\sigma_\top, e_1(\sigma_1), e_2(\sigma_2)) = e_2(\text{merge}(\sigma_\top, e_1(\sigma_1), \sigma_2))^{1}$$

283 Property (P1') would then allow us to establish the required linearization order. Property (P4)
 284 also needs to be altered due to the presence of non-commutative update operations. We modify
 285 (P4) to enforce commutativity for non-rc related events, which gives us flexibility to include such
 286 events in any order while constructing the linearization sequence:

$$287 \text{ (P4')} \quad (e_1, e_2) \notin \text{rc} \wedge (e_2, e_1) \notin \text{rc} \implies e_1(e_2(\sigma)) = e_2(e_1(\sigma))$$

288 However, we now face another major challenge: proving (P1') for the OR-set MRDT. For the
 289 counter MRDT, the operations and merge function used integer addition and subtraction, which
 290 commute with each other. But for the OR-set, add_a uses set union, while merge uses set difference
 291




Fig. 4. OR-set execution

292 ¹Note that we are abusing the rc notation slightly, since rc is a relation over operations O , but we are considering it over
 293 operation instances (i.e. events)
 294

and intersection, which do not commute in general. Hence, (P1') does not hold for arbitrary $\sigma_\tau, \sigma_1, \sigma_2$.

To illustrate this concretely, consider the same execution of Fig. 4, except assume that the state σ_τ of the LCA v_τ is $\{(a, 1)\}$. Let us try to establish (P1') for the merge of versions v_1 and v_2 . First, note that as per the OR-set rc , the antecedent of (P1') is satisfied, as $(e_1, e_2) \in rc$. Now, the RHS in the consequent must contain the tuple $(a, 1)$, since the event e_2 adds $(a, 1)$ to the result of the merge. Does the LHS also contain $(a, 1)$? Expanding the definition of merge in the LHS, $(a, 1)$ will not be present in $(\sigma_\tau \cap e_1(\sigma_\tau) \cap e_2(\sigma_\tau))$ (because $(a, 1) \notin e_1(\sigma_\tau)$, as e_1 removes a). Similarly, since $(a, 1)$ is in σ_τ , it will not be present in $e_2(\sigma_\tau) \setminus \sigma_\tau$. It will not be in $e_1(\sigma_\tau) \setminus \sigma_\tau$, as e_1 removes a . To conclude, $(a, 1)$ will not be present in the LHS, thus invalidating the consequent of (P1').

However, we note that this particular execution is actually spurious, because the tuple $(a, 1)$ in the LCA could only have been added by another add_a operation whose timestamp is the same as e_2 . But this is not possible as the data store ensures that timestamps are unique across all events. In the general case, we would not be able to show (P1') for OR-set because the tuple (a, t) being added by the add_a operation (event e_2) could also be present in the LCA state. However, this situation cannot occur.

Thus, it is possible to show (P1') for all *feasible* states $\sigma_\tau, \sigma_1, \sigma_2$ that may occur during an actual execution. In the case of OR-set, there are two arguments which are required to infer this: (i) timestamps are unique across all events and (ii) if a tuple (a, t) is present in the state σ , then there must have been an add_a operation with timestamp t in the history of events leading to σ . While the first argument is a property of the data store, the second argument is an invariant linking a state with the history of events leading to that state. Such arguments are in general hard to infer, and would also change across different MRDTs. We now present our second major observation which allows us to automatically verify (P1') for feasible states without requiring invariants like argument (ii) linking MRDT states and events.

2.3 Verification using Induction on Event Sequences

In order to show property (P1') for an MRDT implementation, we need to consider the feasible states which would be given as input to the merge function during an actual execution. We observe that we can leverage RA-linearizability of the MRDT implementation, and hence characterize these feasible states by sequences of MRDT update operations (more precisely, events corresponding to update operation instances). We can now use induction over these sequences to establish property (P1'). Note that the input states to merge may themselves have been obtained through prior merges, but we can inductively assume that these prior merges resulted in correct linearizations. Since merge takes as input three states $(\sigma_\tau, \sigma_1, \sigma_2)$, we need to consider three sequences which led to these states and induct on all the three separately.

Concretely, let π_τ be a sequence of events which when applied on the initial MRDT state σ_0 results in the state σ_τ . Since the LCA state always contains events which are common to the states σ_1 and σ_2 , π_τ will be the common prefix of the sequences leading to both σ_1 and σ_2 . We consider the sequences π_1 and π_2 that consist of the local events which when applied on σ_τ led to σ_1 and σ_2 respectively. Fig. 5 depicts the situation. Notice that the last two events on each replica before the merge are fixed to be e_1 and e_2 , which would be related by the rc relation, as per the requirement of property (P1').




Fig. 5. Induction on event sequences

$$\text{merge}(\sigma_0, e_1(\sigma_0), e_2(\sigma_0)) = e_2(\text{merge}(\sigma_0, e_1(\sigma_0), \sigma_0)) \quad (1)$$

$$\begin{aligned} \text{merge}(\sigma_\top, e_1(\sigma_\top), e_2(\sigma_\top)) &= e_2(\text{merge}(\sigma_\top, e_1(\sigma_\top), \sigma_\top)) \\ \implies \text{merge}(e(\sigma_\top), e_1(e(\sigma_\top)), e_2(e(\sigma_\top))) &= e_2(\text{merge}(e(\sigma_\top), e_1(e(\sigma_\top)), e(\sigma_\top))) \end{aligned} \quad (2)$$

We first induct on the sequence π_\top which leads to the state σ_\top . For this, we assume that $\pi_1 = \pi_2 = \epsilon$, and hence $\sigma_\top = \sigma_1 = \sigma_2 = \pi_\top(\sigma_0)$. We also assume the antecedent of property (P1'), i.e. $(e_1, e_2) \in \text{rc}$, and hence our goal is to show its consequent. For the OR-set, e_1 will be a rem_a event, while e_2 will be an add_a event (say with timestamp t).

Eqn. (1) is the base-case of the induction (where $\pi_\top = \epsilon$), and this can be now directly discharged since σ_0 is an empty set, and hence clearly won't contain (a, t) . Eqn. (2) is the inductive case, which assumes that (P1') is true for some LCA state σ_\top , and tries to prove the property when one more update operation (signified by the event e) is applied on the LCA (and also on both σ_1 and σ_2 , since LCA operations are common to both states to be merged). This can also be automatically discharged with the property that events e, e_1, e_2 have different timestamps. Intuitively, the inductive hypothesis establishes that $(a, t) \notin \sigma_\top$, and since the timestamp of event e is different from e_1 and e_2 , it cannot add (a, t) to the LCA, thus preserving the property that $(a, t) \notin e(\sigma_\top)$, thereby implying the consequent. This completes the proof for property (P1') for any arbitrary LCA state σ_\top that may be feasible in an actual execution. A similar inductive strategy is used for proving property (P1') for feasible states σ_1 and σ_2 (more details in §4).

2.4 Intermediate Merges

In our linearization strategy for merges (given by properties (P1'-P4')), we first considered the local update operations of each branch, linearized them according to the conflict-resolution policy, and then applied this sequence on the LCA. This effectively orders the update operations that led to the LCA before the update operations local to each branch.

However, in a Git-based execution model, due to a phenomenon known as intermediate merges, it may happen that update operations of the LCA may need to be linearized after update operations local to a branch. To illustrate this, consider an execution of the OR-set MRDT as shown in Fig. 6. There are 3 operations and 2 merges being performed in this execution, with the events e_1, e_3 at replica r_1 and event e_2 at replica r_2 .

Instead of merging with the latest version v_3 at replica r_1 , replica r_2 first merges with an intermediate version v_1 to generate the version v_4 . Next, this version v_4 is merged with the latest version v_3 of replica r_1 . However, note that for this merge, the LCA will be version v_1 . This is because the set of events associated with version v_3 is $\{e_1, e_3\}$, while for version v_4 , it is $\{e_1, e_2\}$. Hence, the set of common events among both versions would be $\{e_1\}$, which corresponds to the version v_1 . Indeed, in the version graph, both v_1 and v_0 are ancestors of v_3 and v_4 , but v_1 is the lowest common ancestor².

In Fig. 6, we have also provided the linearization of events associated with each version. Notice that for version v_4 , which is obtained through a merge of v_1 and v_2 , the conflict resolution policy of the OR-set linearizes e_2 before e_1 . Now, for the merge




Fig. 6. Intermediate merge

²in §3, we will formally prove that the LCA of two versions according to the version graph contains the intersection of events in both the versions.

of v_3 and v_4 , we have a situation where a local event (e_2 in v_4) needs to be linearized before an event of the LCA (e_1 in v_1). This does not fit our linearization strategy. Let us see why. If we were to try to apply (P1'), it would linearize e_1 after e_3 , since these are the last operations in the two states to be merged and the conflict resolution policy orders $\text{add}_a(e_1)$ after $\text{rem}_a(e_3)$. However, in the execution, e_1 and e_3 are causally related, i.e. e_1 occurs before e_3 on the same replica, and hence they should be linearized in that order. Intuitively, property (P1') does not work because it does not consider the possibility that the last event in one replica could be visible to the last event in another replica, and hence the linearization must obey the visibility relation.

In order to handle this situation, we consider another algebraic property (P1-1), which explicitly forces visibility relation among the last events by making one of them part of the LCA:

$$(P1-1) \text{ merge}(e_1(\sigma_0), e_3(\sigma_1), e_1(\sigma_2)) = e_3(\text{merge}(e_1(\sigma_0), \sigma_1, e_1(\sigma_2)))$$

Note that events in the LCA are visible to events on both replicas being merged. Hence, by having the same event e_1 in both the first and third argument to merge in the LHS, e_3 would have to be linearized after e_1 to respect the visibility order, thus over-riding the rc ordering among them. Property (P1-1) can be directly applied to the execution in Fig. 6 for the merge of v_3 and v_4 (with σ_0 as the state of v_0 , σ_1 as the state of v_1 and σ_2 as the state of v_2), constructing the correct linearization.

We will revisit the example in Fig. 6 and properties (P1') and (P1-1) in a more formal setting in §4, renaming them as BOTTOMUP-2-OP and BOTTOMUP-1-OP. We will also identify the conditions under which these properties can guarantee the existence of a correct linearization.

3 Problem Definition

In this section, we formally define the semantics of the replicated data store on top of which the MRDT implementations operate (§3.1), the notion of RA-linearizability for MRDTs (§3.2), and the process of bottom-up linearization (§3.3).

3.1 Semantics of the Replicated Data Store

[CREATEBRANCH]

$$\frac{N' = N[v \mapsto N(H(r))] \quad r \in \text{dom}(H) \quad r' \notin \text{dom}(H) \quad v \notin \text{dom}(N) \quad H' = H[r \mapsto v] \quad L' = L[v \mapsto L(H(r))] \quad G' = (\text{dom}(N) \cup \{v\}, E \cup \{(H(r), v)\})}{(N, H, L, G, \text{vis}) \xrightarrow{\text{createBranch}(r', r)} (N', H', L', G', \text{vis})}$$

[APPLY]

$$\frac{o \in O_\tau \quad \forall e' \in \bigcup \text{range}(L). \text{time}(e') \neq t \quad e = (t, r, o) \quad r \in \text{dom}(H) \quad v \notin \text{dom}(N) \quad N' = N[v \mapsto \text{do}(N(H(r)), e)] \quad H' = H[r \mapsto v] \quad L' = L[v \mapsto L(H(r)) \cup \{e\}] \quad G' = (\text{dom}(N'), E \cup \{(H(r), v)\}) \quad \text{vis}' = \text{vis} \cup (L(H(r)) \times \{e\})}{(N, H, L, G, \text{vis}) \xrightarrow{\text{apply}(t, r, o)} (N', H', L', G', \text{vis}')}$$

[MERGE]

$$\frac{r_1, r_2 \in \text{dom}(H) \quad v \notin \text{dom}(N) \quad v_\tau = \text{LCA}(H(r_1), H(r_2)) \quad N' = N[v \mapsto \text{merge}(N(v_\tau), N(H(r_1)), N(H(r_2)))] \quad H' = H[r_1 \mapsto v] \quad L' = L[v \mapsto L(H(r_1)) \cup L(H(r_2))] \quad G' = (\text{dom}(N'), E \cup \{(H(r_1), v), (H(r_2), v)\})}{(N, H, L, G, \text{vis}) \xrightarrow{\text{merge}(r_1, r_2)} (N', H', L', G', \text{vis})}$$

[QUERY]

$$\frac{r \in \text{dom}(H) \quad q \in Q_\tau \quad a = \text{query}(N(H(r)), q)}{(N, H, L, G, \text{vis}) \xrightarrow{\text{query}(r, q, a)} (N, H, L, G, \text{vis})}$$

Fig. 7. Semantics of the replicated datastore

The semantics of the replicated store defines all possible executions of an MRDT implementation. Formally, the semantics are parametrised by an MRDT implementation $\mathcal{D} = \langle \Sigma, \sigma_0, \text{do}, \text{merge}, \text{query}, \text{rc} \rangle$ of type $\tau = \langle O_\tau, Q_\tau, \text{Val}_\tau \rangle$ and are represented by a labeled transition system $\mathcal{S}_{\mathcal{D}} = (\Phi, \rightarrow)$. Each configuration in Φ maintains a set of versions, where each version is created either by applying an MRDT operation to an existing version, or by merging two versions. Each replica is associated with a head version, which is the most recent version seen at the replica. Formally, each configuration C in Φ is a tuple $\langle N, H, L, G, \text{vis} \rangle$, where:

- $N : \text{Version} \rightarrow \Sigma$ is a partial function that maps versions to their states (Version is the set of all possible versions).
- $H : \mathcal{R} \rightarrow \text{Version}$ is also a partial function that maps replicas to their head versions.
- $L : \text{Version} \rightarrow \mathbb{P}(\mathcal{E})$ maps a version to the set of events that led to this version. Each event $e \in \mathcal{E}$ is an update operation instance, uniquely identified by a timestamp value (we define $\mathcal{E} = \mathcal{T} \times \mathcal{R} \times O$).
- $G = (\text{dom}(N), E)$ is the version graph, whose vertices represent the versions in the configuration (i.e. those in the domain of N) and whose edges represent a relationship between different versions (we explain the different types of edges below).
- $\text{vis} \subseteq \mathcal{E} \times \mathcal{E}$ is a partial order over events.

Figure 7 gives a formal description of the transition rules. `CREATEBRANCH` forks a new replica r' from an existing replica r , installing a new version v at r' with the same state as the head version $H(r)$ of r , and adding an edge $(H(r), v')$ in the version graph. `APPLY` applies an update operation o on some replica r , generating a new event e with a timestamp different than all events generated so far. $\bigcup \text{range}(L)$ denotes the set of events witnessed across all versions. A new version v is also created whose state is obtained by applying o on the current state of the replica r . The version graph is updated by adding the edge $(H(r), v)$. The vis relation as well as the function L , which tracks events applied at each version, are also updated. In particular, each event e' already applied at r , i.e. $e' \in L(H(r))$, is made visible to e : $(e', e) \in \text{vis}$, while $L'(v)$ is obtained by adding e to $L(H(r))$.

`MERGE` takes two replicas r_1 and r_2 , applies the merge function on the states of their head versions to generate a new version v , which is installed as the new head version at r_1 . Edges are added in the version graph from the previous head versions of r_1 and r_2 to v . $L(v)$ is obtained by taking a union of $L(r_1)$ and $L(r_2)$, and there is no change in the visibility relation. `QUERY` takes a replica r and a query operation q and applies q to the state at the head version of r , returning an output value a . Note that the `QUERY` transition does not modify the configuration and the return value of the query is stored as part of the transition label. While our operational semantics is based on and inspired by previous works [11, 23], we note that it is more general and precisely captures the MRDT system model as opposed to previous works. In particular, Kaki et al. [11] places significant restrictions on the `MERGE` transition, disallowing arbitrary replicas to be merged to ensure that there is a total order on the merge transitions. While the semantics in Soundarapandian et al. [23] does allow arbitrary merges, it is more abstract and high-level, and does not even keep track of versions and the version graph.

Notation: We now introduce some notation that will be used throughout the paper. Given a configuration C , we use $X(C)$ to project the component X of C . For a relation R , we use $x \xrightarrow{R} y$ to signify that $(x, y) \in R$. We use $R|_S$ to indicate the relation as given by R but restricted to elements of the set S . Let R^* denote the reflexive-transitive closure of R , and let R^+ denote the transitive closure of R . For an event e , we use the projection functions `op`, `time`, `rep` to obtain the update operation, timestamp and replica resp. For a sequence of events π , $\pi|_S(\sigma)$ denotes application of the sub-sequence of π restricted to events in S . For a configuration C , we use $e_1 ||_C e_2$ to denote

491 that e_1 and e_2 are concurrent, that is $\neg(e_1 \xrightarrow{\text{vis}(C)} e_2 \vee e_2 \xrightarrow{\text{vis}(C)} e_1)$. Given a total order over a set of
 492 events \mathcal{E} , represented by a sequence π , and $\text{lo} \subseteq \mathcal{E} \times \mathcal{E}$, we say that π extends lo if $\text{lo} \subseteq \pi$. The
 493 relation rc orders update operations, but for convenience we sometime use it for ordering events,
 494 with the intention that it is actually being applied on the underlying update operations. We use
 495 $e_1 \neq e_2$ to indicate that $\text{time}(e_1) \neq \text{time}(e_2)$.

496 We define the initial configuration of $\mathcal{S}_{\mathcal{D}}$ as $C_0 = \langle N_0, H_0, L_0, G_0, \emptyset \rangle$, which consists of only one
 497 replica r_0 . Here, $H_0 = [r_0 \mapsto v_0]$, $N_0 = [v_0 \mapsto \sigma_0]$, where σ_0 is the initial state as given by \mathcal{D}_τ , while
 498 v_0 denotes the initial version and $L_0 = [v_0 \mapsto \emptyset]$. The graph $G_0 = (\{v_0\}, \emptyset)$ is the initial version graph.
 499 An execution of $\mathcal{S}_{\mathcal{D}}$ is defined to be a finite sequence of transitions, $C_0 \xrightarrow{t_1} C_1 \xrightarrow{t_2} C_2 \dots \xrightarrow{t_n} C_n$.
 500 Note that the label of a transition corresponds to its type. Let $\llbracket \mathcal{S}_{\mathcal{D}} \rrbracket$ denote the set of all possible
 501 executions of $\mathcal{S}_{\mathcal{D}}$.

502 Finally, as mentioned earlier, merge is a ternary function, taking as input the states of two
 503 versions to be merged, and the state of the lowest common ancestor (LCA) of the two versions.
 504 Version $v_1 \in V$ is defined to be a causal ancestor of version $v_2 \in V$ if and only if $(v_1, v_2) \in E^*$.
 505

506 *Definition 3.1 (LCA).* Given a version graph $G = (V, E)$ and versions $v_1, v_2 \in V$, $v_\top \in V$ is defined
 507 to be the lowest common ancestor of v_1 and v_2 (denoted by $LCA(v_1, v_2)$) if (i) $(v_\top, v_1) \in E^*$ and
 508 $(v_\top, v_2) \in E^*$, (ii) $\forall v \in V. (v, v_1) \in E^* \wedge (v, v_2) \in E^* \implies (v, v_\top) \in E^*$.
 509

510 Note that the version history graph at any point in any execution is guaranteed to be acyclic (i.e.
 511 a DAG), and hence the LCA (if it exists) is guaranteed to be unique. We now present an important
 512 property linking the LCA of two versions with events applied at each version.

513 *LEMMA 3.2.* Given a configuration $C = \langle N, H, L, G, \text{vis} \rangle$ reachable in some execution $\tau \in \llbracket \mathcal{S}_{\mathcal{D}} \rrbracket$ and
 514 two versions $v_1, v_2 \in \text{dom}(N)$, if v_\top is the LCA of v_1 and v_2 in G , then $L(v_\top) = L(v_1) \cap L(v_2)$ ³.
 515

516 Thus, the events of the LCA are exactly those applied at both the
 517 versions. This intuitively corresponds to the fact that $LCA(v_1, v_2)$
 518 is the most recent version from which the two versions v_1 and v_2
 519 diverged. Note that it is possible that the LCA may not exist for
 520 two versions. Fig. 8 depicts the version graph of such an execution.
 521 Vertices with in-degree 1 (i.e. v_1, v_2, v_3, v_4) have been generated by
 522 applying a new update operation (with the orange edges labeled by
 523 the corresponding events e_1, e_2, e_3, e_4), while vertices with in-degree
 524 2 have been obtained by merging two other versions (depicted by
 525 blue edges). The merge of v_1 and v_4 (leading to v_6) has a unique LCA
 526 v_0 , similarly, merge of v_2 and v_3 (leading to v_5) also has a unique
 527 LCA v_0 . However, if we now want to merge v_5 and v_6 , both v_1 and
 528 v_2 are ancestors, but there is no LCA. We note that this execution
 529 will actually be prohibited by the semantics of Kaki et al. [11], since
 530 the two merges leading to v_5 and v_6 are concurrent.

531 Notice that $L(v_5) = \{e_1, e_2, e_3\}$, while $L(v_6) = \{e_1, e_2, e_4\}$. Hence, by Lemma 3.2, $L(LCA(v_5, v_6)) =$
 532 $\{e_1, e_2\}$, but such a version is not generated during the execution. To resolve this issue, we introduce
 533 the notion of *potential LCAs*.

534 *Definition 3.3 (Potential LCAs).* Given a version graph $G = (V, E)$ and versions $v_1, v_2 \in V$,
 535 $v_\top \in V$ is defined to be a potential LCA of v_1 and v_2 if (i) $(v_\top, v_1) \in E^*$ and $(v_\top, v_2) \in E^*$, (ii)
 536 $\neg(\exists v. (v, v_1) \in E^* \wedge (v, v_2) \in E^* \wedge (v_\top, v) \in E^*)$.
 537




Fig. 8. Version Graph with no LCA for v_5 and v_6

538 ³All proofs are in the Appendix §A
 539

540 For merging v_1 and v_2 , we first find all the potential LCAs, and recursively merge them to obtain
 541 the actual LCA state. For the execution in Fig. 8, the potential LCAs of v_5 and v_6 would be v_1 and
 542 v_2 (with $L(v_1) = \{e_1\}$ and $L(v_2) = \{e_2\}$); merging them would get us the actual LCA. In §A.1, we
 543 prove that this recursive merge-based strategy is guaranteed to generate the actual LCA.

544 3.2 Replication-aware Linearizability for MRDTs

545 As mentioned in §2, our goal is to show that the state of every version v generated during an
 546 execution is a linearization of the events in $L(v)$. We use the notation lo to indicate the linearization
 547 relation, which is a binary relation over events. For an execution in $\mathcal{S}_{\mathcal{D}}$, we want lo between the
 548 events of the execution to satisfy certain desirable properties: (i) lo between two events should not
 549 change during an execution, (ii) lo should obey the conflict resolution policy for concurrent events
 550 and (iii) lo should obey the replica-local vis ordering for non-concurrent events. This would ensure
 551 that two versions which have observed the same set of events will have the same state (i.e. *strong*
 552 *eventual consistency*), and this state would also be a linearization of update operations of the data
 553 type satisfying the conflict resolution policy.

554 While the lo relation in classical linearizability literature is typically a total order, in our context,
 555 we take advantage of commutativity of update operations, and only define lo over non-commutative
 556 events. As we will see later, this flexibility allows us to have different sequences of events which
 557 extend the same lo relation between non-commutative events, and hence are guaranteed to lead
 558 to the same state. We use the notation $e \rightleftarrows e'$ to indicate that events e and e' commute with each
 559 other. Formally, this means that $\forall \sigma. e(e'(\sigma)) = e'(e(\sigma))$. Two update operations o, o' commute
 560 if $\forall e, e'. \text{op}(e) = o \wedge \text{op}(e') = o' \implies e \rightleftarrows e'$. As mentioned earlier, the rc relation is also only
 561 defined between non-commutative update operations.


562 LEMMA 3.4. *Given a set of events \mathcal{E} , if $\text{lo} \subseteq \mathcal{E} \times \mathcal{E}$ is defined over every pair of non-commutative
 563 events in \mathcal{E} , then for any two sequences π_1, π_2 which extend lo , for any state σ , $\pi_1(\sigma) = \pi_2(\sigma)$.*

564 Given a configuration $C = \langle N, H, L, G, \text{vis} \rangle$, let $\mathcal{E}_C = \bigcup \text{range}(L(C))$ denote the set of events
 565 witnessed across all versions in C . Then, our goal is to define an appropriate linearization relation
 566 $\text{lo}_C \subseteq \mathcal{E}_C \times \mathcal{E}_C$, which adheres to the rc relation for concurrent events, the vis relation for non-
 567 concurrent events, and for every version $v \in \text{dom}(N)$, $N(v)$ should be obtained by sequentializing
 568 the events in $L(v)$, with the sequence extending lo_C . Note that this requires lo^+ to be irreflexive⁴.

569 We now demonstrate that an lo relation with all the desirable
 570 properties may not exist for all executions. Suppose there are MRDT
 571 update operations o, o' such that $o \xrightarrow{\text{rc}} o'$. Fig. 9 contains a part of
 572 the version graph generated during some execution, containing
 573 two instances of both o and o' . We use $e_i : o_i$ to denote that event
 574 $\text{op}(e_i) = o_i$. Notice that e_1 and e_4 , e_2 and e_3 are concurrent, while
 575 e_1 and e_3 , e_2 and e_4 are non-concurrent. Applying the rc ordering
 576 on concurrent events, we would want $e_3 \xrightarrow{\text{lo}} e_2$ and $e_4 \xrightarrow{\text{lo}} e_1$, while
 577 applying vis ordering, we would want $e_1 \xrightarrow{\text{lo}} e_3$ and $e_2 \xrightarrow{\text{lo}} e_4$.
 578 However, this results in a lo -cycle, thus making it impossible to
 579 construct a sequence of update operations for the merge of v_5 and v_6 , which adheres to the lo
 580 ordering.

581 Notice that the above execution only requires the rc relation to be non-empty (i.e. there should
 582 exist some $(o, o') \in \text{rc}$). If the rc relation is empty, then all update operations would commute

583 ⁴ lo need not be transitive, as we only want to define lo between non-commutative events, and non-commutativity is not a
 584 transitive property



585 Fig. 9. Example demonstrating
 586 cycle in lo

with each other, and hence the lo relation would also be empty. If rc is non-empty, rc⁺ should be irreflexive to ensure irreflexivity of lo⁺. Note that rc⁺ being irreflexive means that for any MRDT update operation o , $(o, o) \notin rc$, and hence o must commute with itself, since rc relation is defined for all pairs of non-commutative update operations. Furthermore, Fig. 9 shows that even if rc⁺ is irreflexive, it may still not be possible to construct an lo relation which can be extended to a total order and which adheres to the rc relation between all pairs of concurrent events. To ensure existence of an lo relation such that lo⁺ is irreflexive when rc⁺ is irreflexive, we define it as follows:

Definition 3.5 (Linearization relation). Let C be a configuration reachable in some execution in $\llbracket \mathcal{S}_{\mathcal{D}} \rrbracket$. Let \mathcal{E}_C be the set of events in C . Then, lo_C is defined as:

$$\forall e_1, e_2 \in \mathcal{E}_C. e_1 \xrightarrow{lo_C} e_2 \Leftrightarrow (e_1 \xrightarrow{vis(C)} e_2 \wedge \neg e_1 \overset{\rhd}{\rhd} e_2) \\ \vee (e_1 \parallel_C e_2 \wedge e_1 \xrightarrow{rc} e_2 \wedge \neg(\exists e_3 \in \mathcal{E}. e_2 \xrightarrow{vis(C)} e_3 \wedge \neg e_2 \overset{\rhd}{\rhd} e_3))$$

lo_C follows the visibility relation only between non-commutative events. For concurrent non-commutative events e_1 and e_2 with $e_1 \xrightarrow{rc} e_2$, lo_C follows the rc relation only if there is no event e_3 such that e_2 is visible to e_3 and e_2 doesn't commute with e_3 . Applying this definition to the execution in Fig. 9, for the configuration obtained after merge, we would have neither $e_4 \xrightarrow{lo} e_1$, nor $e_3 \xrightarrow{lo} e_2$, thus avoiding the cycle in lo.

LEMMA 3.6. *For an MRDT \mathcal{D} such that rc⁺ is irreflexive, for any configuration C reachable in $\mathcal{S}_{\mathcal{D}}$, lo_C^+ is irreflexive.*

Going forward, we will assume that rc⁺ is irreflexive for any MRDT \mathcal{D} . We note that restricting lo to not always obey the rc relation by considering non-commutative update operations happening locally (and thus related by vis) is also sensible from a practical perspective. For example, in the case of OR-set, even though we have $rem_a \xrightarrow{rc} add_a$, if add_a is locally followed by another rem_a , it doesn't make sense to order a concurrent rem_a event before the add_a event. More generally, if an event e_2 is visible to another event e_3 with which it doesn't commute, then e_2 is effectively "overwritten" by e_3 , and hence there is no need to linearize a concurrent event e_1 before e_2 .

While lo_C is now guaranteed to be irreflexive for any configuration C , and hence can be extended to a sequence, it now no longer enforces an ordering among all non-commutative pairs of events. Thus, there could exist sequences π_1, π_2 extending an lo_C relation which may contain a pair of non-commutative events in different orders. For example, in Fig. 9, for the configuration C obtained after the merge, $lo_C = \{(e_1, e_3), (e_2, e_4)\}$, resulting in sequences $\pi_1 = e_1 e_2 e_3 e_4$ and $\pi_2 = e_1 e_3 e_2 e_4$ which both extend lo_C , but contain the non-commutative events e_2 and e_3 in different orders. Thus, Lemma 3.4 can no longer be applied, and it is not guaranteed that π_1 and π_2 would lead to the same state. Notice that in the sequences π_1 and π_2 above, even though e_2 and e_3 appear in different orders, e_4 always appears after both. Indeed, e_4 must appear after e_2 due to visibility relation, and since e_3 and e_4 commute with each other (since both correspond to the same operation o), it is enough to consider sequences where e_4 appears after e_3 . Based on the above observation, we now introduce a notion called conditional commutativity to ensure that sequences such as π_1, π_2 would lead to the same state:

Definition 3.7 (Conditional Commutativity). Events e and e' are said to conditionally commute with respect to event e'' (denoted by $e \overset{e''}{\rhd} e'$) if $\forall \sigma \in \Sigma. \forall \pi \in \mathcal{E}^*. e''(\pi(e(e'(\sigma)))) = e''(\pi(e'(e(\sigma))))$.

Update operations o and o' conditionally commute w.r.t. update operation o'' if $\forall e, e', e''. \text{op}(e) = o \wedge \text{op}(e') = o' \wedge \text{op}(e'') = o'' \Rightarrow e \stackrel{e''}{\rightleftarrows} e'$. For example, for the OR-set MRDT of Fig. 2, $\text{add}_a \stackrel{\text{rem}_a}{\rightleftarrows} \text{rem}_a$. Even though add and remove operations of the same element do not commute with each other, if there is guaranteed to be a future remove operation, then they do commute. For the execution in Fig. 9, if e_2 and e_3 conditionally commute w.r.t. e_4 , then both the sequences π_1 and π_2 will lead to the same state. For non-commutative update operations which are not ordered by lo , we enforce their conditional commutativity through the following property:

$$\text{COND-COMM}(\mathcal{D}) \triangleq \forall o_1, o_2, o_3 \in O. (o_1 \xrightarrow{\text{rc}} o_2 \wedge \neg o_2 \stackrel{o_3}{\rightleftarrows} o_3) \Rightarrow o_1 \stackrel{o_3}{\rightleftarrows} o_2$$

$\text{COND-COMM}(\mathcal{D})$ is a property of an MRDT \mathcal{D} , enforcing conditional commutativity of update operations o_1 and o_2 w.r.t. o_3 if o_2 does not commute with o_3 . Connecting this with the definition of linearization relation, if there are events e_1, e_2, e_3 performing operations o_1, o_2, o_3 resp., and if $e_1 \xrightarrow{\text{rc}} e_2, e_2 \xrightarrow{\text{vis}} e_3$ and $\neg e_2 \stackrel{o_3}{\rightleftarrows} e_3$, then there will not be a linearization relation between e_1 and e_2 . However, $\text{COND-COMM}(\mathcal{D})$ would then ensure that the ordering of e_1 and e_2 will not matter, due to the presence of the event e_3 . We also formalize the requirement of an rc relation between all pairs of non-commutative update operations:

$$\text{RC-NON-COMM}(\mathcal{D}) \triangleq \forall o_1, o_2 \in O. \neg o_1 \stackrel{\text{rc}}{\rightleftarrows} o_2 \Leftrightarrow o_1 \xrightarrow{\text{rc}} o_2 \vee o_2 \xrightarrow{\text{rc}} o_1$$

LEMMA 3.8. *For an MRDT \mathcal{D} which satisfies $\text{RC-NON-COMM}(\mathcal{D})$ and $\text{COND-COMM}(\mathcal{D})$, for any reachable configuration C in $\mathcal{S}_{\mathcal{D}}$, for any two sequences π_1, π_2 over \mathcal{E}_C which extend lo_C , for any state σ , $\pi_1(\sigma) = \pi_2(\sigma)$.*

Definition 3.9 (RA-linearizability of MRDT). Let \mathcal{D} be an MRDT which satisfies $\text{RC-NON-COMM}(\mathcal{D})$ and $\text{COND-COMM}(\mathcal{D})$. Then, a configuration $C = \langle N, H, L, G, \text{vis} \rangle$ of $\mathcal{S}_{\mathcal{D}}$ is RA-linearizable if, for every active replica $r \in \text{range}(H)$, there exists a sequence π consisting of all events in $L(H(r))$ such that $\text{lo}(C)|_{L(H(r))} \subseteq \pi$ and $N(H(r)) = \pi(\sigma_0)$. An execution $\tau \in \llbracket \mathcal{S}_{\mathcal{D}} \rrbracket$ is RA-linearizable if all of its configurations are RA-linearizable. Finally, \mathcal{D} is RA-linearizable if all of its executions are RA-linearizable.

For a configuration to be RA-linearizable, every active replica must have a state which can be obtained by applying a sequence of events witnessed at that replica, and that sequence must obey the linearization relation of the configuration. For an execution to be RA-linearizable, all of its configurations must be RA-linearizable. Lemma 3.6 ensures the existence of a sequence extending the linearization relation, while Lemma 3.8 ensures that two versions which have witnessed the same set of events will have the same state (i.e. strong eventual consistency). Further, we also show that if an MRDT is RA-linearizable, then for any query operation in any execution, the query result is derived from the state obtained by applying the update events seen at the corresponding replica right before the query:

LEMMA 3.10. *If MRDT \mathcal{D} is RA-linearizable, then for all executions $\tau \in \llbracket \mathcal{S}_{\mathcal{D}} \rrbracket$, for all transitions $C \xrightarrow{\text{query}(r,q,a)} C'$ in τ where $C = \langle N, H, L, G, \text{vis} \rangle$, there exists a sequence π consisting of all events in $L(H(r))$ such that $\text{lo}(C)|_{L(H(r))} \subseteq \pi$ and $a = \text{query}(\pi(\sigma_0), q)$.*

Compared to the definition of RA-linearizability in Wang et. al. [25], there is one major difference: Wang et. al. also consider a sequential specification in the form of a set of valid sequences of data-type operations, and requires the linearization sequence to belong to the specification. Our definition simply requires the state of a replica to be a linearization of the update operations applied to the replica, without appealing to a separate sequential specification. Once this is done, we can

separately show that a linearization of the MRDT operations obeys the sequential specification. For this, we can ignore the presence of the merge operation as well as the MRDT system model (which are taken care of by the RA-linearizability definition), thus boiling down to proving a specification over a sequential functional implementation, which is a well-studied problem.

3.3 Bottom-up Linearization

As demonstrated in §2, our approach to show RA-linearizability of an MRDT implementation is based on using algebraic properties of merge (specifically, commutativity of merge and update operation application) which allows us to show that the result of a merge operation is a linearization of the events in each of the versions being merged. We first describe a generic template for the algebraic properties which can be used to prove RA-linearizability:

$$\frac{\forall j. \pi_j \in \mathcal{E} \cup \{\epsilon\} \quad l, a, b \in \Sigma \quad \pi \in \{\pi_0, \pi_1, \pi_2\} \quad \forall j. \pi'_j = \pi_j - \pi}{\text{merge}(\pi_0(l), \pi_1(a), \pi_2(b)) = \pi(\text{merge}(\pi'_0(l), \pi'_1(a), \pi'_2(b))))} \quad [\text{BOTTOMUPTEMPLATE}]$$

The template for the algebraic property is given in the conclusion of the above rule, while the premises describe certain conditions. Each π_j for $j \in \{0, 1, 2\}$ is a sequence of 0 or 1 event (i.e. either ϵ or a single event e_j), while l, a, b are arbitrary states of the MRDT. Note that applying the ϵ event on a state leaves it unchanged (i.e. $\epsilon(s) = s$). Then, we can select one event π which has been applied to the arguments of merge on the LHS, and bring it outside, i.e. remove the event from each argument on which it was applied, and instead apply the event to the result of merge. Note that the notation $\pi'_j = \pi_j - \pi$ means that if $\pi = \pi_j$, then $\pi'_j = \epsilon$, else $\pi'_j = \pi_j$.

The rule (P1') given in §2.2 can be seen as an instantiation of the above template with $\pi_0 = \epsilon$, $\pi_1 = e_1$, $\pi_2 = e_2$ and $\pi = e_2$ where $e_1 \xrightarrow{\text{rc}} e_2$. Similarly, (P1-1) is another instantiation with $\pi_0 = \pi_2 = e_1$, $\pi_1 = e_3$ and $\pi = e_3$ where $e_3 \neq e_1$. Assuming that the input arguments to merge are obtained through sequences of events τ_0, τ_1, τ_2 , the template rule builds the linearization sequence $\tau = \tau' e$ where e is the last event in one of the τ_i s, and τ' is recursively generated by applying the rule on $\tau' = \tau - e$. We call this procedure as *bottom-up linearization*. The event e should be chosen in such a way that the sequence τ is an extension of the linearization relation (Def. 3.5).

However, bottom-up linearization might fail if the last event in the merge output is not the last event in any of the three arguments to merge. For example, consider the execution shown in Fig. 10, where there exists an rc-chain: $o_2 \xrightarrow{\text{rc}} o_3 \xrightarrow{\text{rc}} o_1$, and o_1 and o_2 are non-commutative. e_1 is visible to e_2 , while event e_3 is concurrent to e_1 and e_2 . Now, for the version obtained after merging v_3 and v_4 , the linearization relation would be $e_1 \xrightarrow[\text{vis}]{\text{lo}} e_2$ and $e_2 \xrightarrow[\text{rc}]{\text{lo}} e_3$. Notably, even though e_1 and e_3 are also concurrent, and rc orders o_3 before o_1 , this will not result in a linearization relation from e_3 to e_1 , due to the presence of a non-commutative update operation e_2 to which e_1 is visible. The bottom-up linearization for the merge of v_3 and v_4 , will result in the sequence $e_1 e_2 e_3$, which is an extension of the linearization order.

However, suppose we first merge versions v_2 and v_4 , to obtain the version v_5 , where the linearization relation is $e_3 \xrightarrow[\text{rc}]{\text{lo}} e_1$. Merging v_3 and v_5 (with LCA v_2) would have the same linearization relation as merging v_3 and v_4 . However, the sequences leading to v_3 and v_5 are $e_1 e_2$ and $e_3 e_1$ respectively, while the only sequence which extends the linearization relation for their merge is $e_1 e_2 e_3$. Bottom-up linearization will then be constrained to pick either e_1 or e_2 to appear at the end, but such a




Fig. 10. Example demonstrating the failure of bottom-up linearization in the presence of an rc-chain

687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735

sequence will not extend the linearization relation resulting in failure of bottom-up linearization. To avoid such cases, we place an additional constraint which prohibits the presence of an rc-chain:

$$\text{NO-RC-CHAIN}(\mathcal{D}) \triangleq \neg(\exists o_1, o_2, o_3 \in O. o_1 \xrightarrow{\text{rc}} o_2 \xrightarrow{\text{rc}} o_3)$$

If there is an rc-chain, executions such as Fig. 10 are possible, resulting in infeasibility of bottom-up linearization. However, we will show that if an MRDT satisfies $\text{NO-RC-CHAIN}(\mathcal{D})$, then we can use bottom-up linearization to prove that \mathcal{D} is linearizable. We note that NO-RC-CHAIN is a pragmatic restriction and consistent with standard conflict-resolution strategies such as add/remove-wins, enable/disable-wins, update/delete-wins, etc. which are typically used in MRDT implementations.

4 Verifying RA-linearizability of MRDTs

In this section, we present our verification strategy for proving RA-linearizability of MRDTs using bottom-up linearization. According to Def. 3.9, in order to prove that an MRDT \mathcal{D} is linearizable, we need to consider every configuration C reachable in any execution, and show that all replicas in C have states which can be obtained by linearizing the events applied to the replica, i.e. finding a sequence which obeys the linearization relation (Def. 3.5). We will assume that \mathcal{D} satisfies the three constraints (RC-NON-COMM , COND-COMM and NO-RC-CHAIN) necessary for an MRDT to be linearizable, and for bottom-up linearization to succeed.

Our overall proof strategy is to use induction on the length of the execution and to extract generic verification conditions (VCs) which help us to discharge the inductive case. These VCs would essentially be instantiations of the BOTTOMUPTEMPLATE rule, proving that the merge operation results in a linearization of the events of the two versions being merged. Proving these VCs for arbitrary MRDTs is not straightforward (as discussed in §2.3), and hence we propose another induction scheme over event sequences. We first discuss the instantiations of the BOTTOMUPTEMPLATE rule required for linearizing merges.

4.1 Linearizing Merge Operations

Consider an execution $\tau \in \llbracket \mathcal{S}_{\mathcal{D}} \rrbracket$ such that all configurations in τ are linearizable. Suppose τ ends in the configuration C . Now, we extend τ by one more transition, resulting in the new configuration C' ; we need to prove that C' is also linearizable. Let $C = \langle N, H, L, G, \text{vis} \rangle$, $C' = \langle N', H', L', G', \text{vis}' \rangle$. It is easy to see if that this transition is caused due to CREATEBRANCH or APPLY rules, then C' will be linearizable. For example, in the $[\text{APPLY}]$ transition, where a new update operation o is applied on a replica r (generating a new event e), only the state at r changes, and this new state is obtained by directly applying e on the original state σ at r . Since σ was assumed to be linearizable, there exists a sequence π which extends $\text{lo}(C)|_{L(H(r))}$, with $\sigma = \pi(\sigma_0)$ (recall that $L(H(r))$ denotes the set of events applied at r). Then, the new state $e(\sigma)$ is clearly linearizable through the sequence πe which extends $\text{lo}(C')|_{L'(H'(r))}$.

We focus on the difficult case when there is a MERGE transition from C to C' which merges the replicas r_1 and r_2 . Let σ_1 and σ_2 be the states of the head versions v_1 and v_2 at r_1 and r_2 respectively. Let σ_{\top} be the state of the LCA version v_{\top} of v_1 and v_2 . Recall that $L(v_{\top}) = L(v_1) \cap L(v_2)$. The transition will install a new version with state $\sigma_m = \text{merge}(\sigma_{\top}, \sigma_1, \sigma_2)$ at the replica r_1 , leaving the other replicas unchanged. Also, $L'(v_m) = L(v_1) \cup L(v_2)$. We need to show that there exists a sequence π of events in $L'(v_m)$ such that π extends $\text{lo}(C')|_{L'(v_m)}$ and $\sigma_m = \pi(\sigma_0)$.

We first describe the structure of a sequence π which extends $\text{lo}(C')|_{L'(v_m)}$. For ease of readability, we use L_1 for $L(v_1)$, L_2 for $L(v_2)$ and L_{\top} for $L(v_{\top})$, and lo_m for $\text{lo}(C')|_{L'(v_m)}$. We define the following

sets of events:

$$L'_1 = L_1 \setminus L_\top \quad L'_2 = L_2 \setminus L_\top$$

$$L_1^b = \{e \in L'_1 \mid \exists e_\top \in L_\top. (e \xrightarrow{\text{lo}_m} e_\top \vee \exists e' \in L'_1. e \xrightarrow{\text{lo}_m} e' \xrightarrow{\text{lo}_m} e_\top)\}$$

$$L_2^b = \{e \in L'_2 \mid \exists e_\top \in L_\top. (e \xrightarrow{\text{lo}_m} e_\top \vee \exists e' \in L'_2. e \xrightarrow{\text{lo}_m} e' \xrightarrow{\text{lo}_m} e_\top)\}$$

$$L_\top^a = \{e_\top \in L_\top \mid \exists e \in L_1^b \cup L_2^b. e \xrightarrow{\text{lo}_m} e_\top\} \quad L_1^a = L'_1 \setminus L_1^b \quad L_2^a = L'_2 \setminus L_2^b \quad L_\top^b = L_\top \setminus L_\top^a$$

L'_1 and L'_2 are the local events in each version. Note that any pair of events $e_1 \in L'_1, e_2 \in L'_2$ will necessarily be concurrent. This is because in any reachable configuration, any version v is always **causally closed**, which means that if $e_1 \xrightarrow{\text{vis}} e_2$ and $e_2 \in L(v)$, then $e_1 \in L(v)$. Hence, for events $e_1 \in L'_1, e_2 \in L'_2$, if $e_1 \xrightarrow{\text{vis}} e_2$ then $e_1 \in L'_2$, which would make e_1 a non-local event (i.e. part of the LCA). Bottom-up linearization first linearizes the local events across the two versions using the rc relation for non-commutative events, and then linearizes events of the LCA. However, as demonstrated by the example in §2.4, local events may also need to be linearized before events of the LCA (due to possible intermediate merges), and these events are collected in the sets L_1^b and L_2^b . Specifically, L_i^b ($i = 1, 2$) contains those local events e in L'_i which either occur lo_m before some event in the LCA, or which occur lo_m before another local event e' which occurs lo_m before an LCA event. The events of the LCA which need to be linearized after local events are collected in L_\top^a . Finally, L_1^a and L_2^a contain local events which do not occur lo_m before an LCA event.

Example 4.1. Consider the execution in Fig. 6, and the merge of versions v_3 and v_4 , for which the LCA is v_1 . For this merge, $L'_1 = \{e_3\}, L'_2 = \{e_2\}, L_1^b = \emptyset, L_2^b = \{e_2\}, L_\top^a = \{e_1\}$. For the merge of versions v_1 and v_2 (whose LCA is v_0), $L'_1 = \{e_1\}, L'_2 = \{e_2\}$, while L_1^b, L_2^b, L_\top^a will all be empty (since no local event comes lo -before an LCA event).

We now show that there exists a sequence π which extends lo_m and which has events in $S_1 = L_\top^b$ followed by $S_2 = L_\top^a \cup L_1^b \cup L_2^b$ followed by $S_3 = L_1^a \cup L_2^a$ (later, we will discuss the ordering of events inside each set S_i). To prove this, we will demonstrate that there is no lo_m from events in S_i to events in S_{i-1} . Based on the definitions of the S_i sets, we can deduce some obvious facts: (i) there cannot be events $e \in S_3, e' \in L_\top$ such that $e \xrightarrow{\text{lo}_m} e'$, because otherwise, such an event e would be in $L_1^b \cup L_2^b$ (and hence not in S_3), (ii) there cannot be events $e \in L_1^b \cup L_2^b, e' \in L_\top^b$ such that $e \xrightarrow{\text{lo}_m} e'$, because otherwise, such an event e' would be in L_\top^a . In addition, using NO-RC-CHAIN and RC-NON-COMM, we also prove the following:

LEMMA 4.2. (1) For events $e \in L_1^a \cup L_2^a, e' \in L_1^b \cup L_2^b, \neg(e \xrightarrow{\text{lo}_m} e')$.

(2) For events $e \in L_\top^a, e' \in L_\top^b, \neg(e \xrightarrow{\text{lo}_m} e')$.

(1) from the above lemma ensures that there is no lo_m relation from S_3 to S_2 , while (2) ensures the same from S_2 to S_1 . Hence a sequence with the structure $S_1 S_2 S_3$ would extend lo_m . Let us now consider the ordering among events in each set. First, for S_3 , this set contains local events which are guaranteed to not come lo_m before any event of the LCA. An event in L_1^a will be concurrent with an




Fig. 11. Structure of sequence π extending lo_m

event in L_2^a , and the linearization relation between them will depend upon the rc relation between the underlying operations (if the events don't commute). We now instantiate `BOTTOMUPTEMPLATE` for the case where both L_1^a and L_2^a are non-empty in the rule `BOTTOMUP-2-OP` in Fig. 12, so that the linearization needs to consider the rc relation between events in the two sets.

$$\begin{array}{l}
\text{[BOTTOMUP-2-OP]} \\
\frac{e_1 \neq e_2 \quad e_1 \xrightarrow{\text{rc}} e_2 \vee e_1 \not\Rightarrow e_2}{\text{merge}(l, e_1(a), e_2(b)) = e_2(\text{merge}(l, e_1(a), b))} \\
\text{[BOTTOMUP-1-OP]} \\
\frac{(e_\top \neq \epsilon \wedge e_1 \neq e_\top) \vee (e_\top = \epsilon \wedge l = b)}{\text{merge}(e_\top(l), e_1(a), e_\top(b)) = e_1(\text{merge}(e_\top(l), a, e_\top(b)))} \\
\text{[BOTTOMUP-0-OP]} \quad \text{[MERGEIDEMPOTENCE]} \quad \text{[MERGECOMMUTATIVITY]} \\
\text{merge}(e_\top(l), e_\top(a), e_\top(b)) = e_\top(\text{merge}(l, a, b)) \quad \text{merge}(a, a, a) = a \quad \text{merge}(l, a, b) = \text{merge}(l, b, a)
\end{array}$$

Fig. 12. Bottom-up Linearization

Note that e_1, e_2 and l, a, b are all universally quantified. The `BOTTOMUP-2-OP` rule is an algebraic property of merge which needs to be separately shown for each MRDT implementation. For our case where we are trying to linearize $\text{merge}(\sigma_\top, \sigma_1, \sigma_2)$, we can apply `BOTTOMUP-2-OP` with $l = \sigma_\top$, $e_1(a) = \sigma_1$ and $e_2(b) = \sigma_2$. Note that since L_1^a and L_2^a are both non-empty, $e_1 \in L_1^a$, $e_2 \in L_2^a$ (in fact, e_1 and e_2 would be the maximal events in L_1^a and L_2^a according to lo_m). `BOTTOMUP-2-OP` would then linearize e_2 at the end of the sequence. If $e_1 \xrightarrow{\text{rc}} e_2$, then $e_1 \xrightarrow{\text{lo}_m} e_2$, and thus linearizing e_2 at the end obeys the lo_m ordering. Note that due to the `NO-RC-CHAIN` constraint, e_2 cannot come lo_m before another concurrent event e_3 . `BOTTOMUP-2-OP` can now be recursively applied on $\text{merge}(l, e_1(a), b)$, by considering e_1 and the last event leading to the state b . By repeatedly applying `BOTTOMUP-2-OP` all the remaining events in L_1^a and L_2^a can be linearized until one of the sets becomes empty.

Let us now consider the scenario where exactly one of L_1^a and L_2^a is empty. WLOG, let L_1^a be non-empty. We instantiate `BOTTOMUPTEMPLATE` for the case where L_1^a is non-empty and L_2^a is empty in the rule `BOTTOMUP-1-OP` in Fig. 12, so that the linearization orders all events of L_1^a after events of S_2 .

Let us consider the first clause in the premise where $e_\top \neq \epsilon$. To understand `BOTTOMUP-1-OP`, note that if L_2^a is empty, then all local events in L_2^b are linearized before the LCA events. In this case, the last event which leads to the state σ_2 must be an LCA event. `BOTTOMUP-1-OP` uses this observation, with $e_\top(l) = \sigma_\top$, $e_1(a) = \sigma_1$ and $e_\top(b) = \sigma_2$. Notice that the last event in both the LCA and the second argument to merge are exactly the same. e_\top will be the maximal event (according to lo_m relation) in L_\top^a , while e_1 will be the maximal event in L_1^a . `BOTTOMUP-1-OP` then linearizes e_1 at the end of the sequence, thus ensuring that all L_1^a events are linearized after events in S_1 and S_2 . It is possible that L_\top^a is empty, in which case L_2^b will be empty, which is covered by the second clause where $e_\top = \epsilon$ and $l = b$ since there is no local event in the second state.

Example 4.3. Referring to Example 4.1 for the execution in Fig. 6, recall that for the merge of v_3 and v_4 , we have $L_1^a = \{e_3\}$, $L_2^a = \emptyset$ and $L_\top = \{e_1\}$. `BOTTOMUP-1-OP` can be applied in this scenario to linearize e_3 at the end of the sequence.

`BOTTOMUP-2-OP` and `BOTTOMUP-1-OP` can thus be used to linearize all events in S_3 . Let us now consider S_2 , which contains both local events in $L_1^b \cup L_2^b$ and LCA events in L_\top^a . We first provide a more fine-grained structure of lo_m among events in the set S_2 . Let $L_\top^a = \{e_1^\top, \dots, e_m^\top\}$. For each e_i^\top , we collect all local events from L_1^b and L_2^b which need to be linearized before e_i^\top . For local events which need to be linearized before multiple e_i^\top s, we associate them with the smallest such i . We

883 use $L_1^b(e_i^\top)$ and $L_2^b(e_i^\top)$ to denote these sets. Formally:

$$884 \quad \forall e_i^\top \in L_\top^a. L_1^b(e_i^\top) = \{e \in L_1^a \mid (\forall j. j < i \implies e \notin L_1^b(e_j^\top)) \wedge e \xrightarrow{\text{lo}_m} e_i^\top \vee \exists e' \in L_1^a. e \xrightarrow{\text{lo}_m} e' \xrightarrow{\text{lo}_m} e_i^\top\}$$

886 $L_2^b(e_i^\top)$ is defined in a similar manner. We now prove the following lemma using NO-RC-CHAIN
887 and RC-NON-COMM:

889 LEMMA 4.4. (1) For all events $e_i^\top, e_j^\top \in L_\top^a$, where $L_\top^a = \{e_1^\top, \dots, e_m^\top\}$, $\neg(e_i^\top \xrightarrow{\text{lo}_m} e_j^\top)$

890 (2) For events $e \in L_1^b(e_i^\top) \cup L_2^b(e_i^\top)$, $e' \in L_1^b(e_j^\top) \cup L_2^b(e_j^\top)$ where $j < i$, $\neg(e \xrightarrow{\text{lo}_m} e')$.

892 From (1) in the above lemma, since there is no lo_m relation among events in L_\top^a , consider the
893 sequence $e_1^\top e_2^\top \dots e_m^\top$ as a starting point for the sequence of events in S_2 which extends lo_m . We
894 then inject $L_1^b(e_i^\top) \cup L_2^b(e_i^\top)$ before each e_i^\top in the sequence $e_1^\top e_2^\top \dots e_m^\top$, as shown in Fig. 11. Note
895 that in Fig. 11, we have only presented various segments of the sequence, with the ordering within
896 those segments determined by vis and rc . By (2) in Lemma 4.4, we can show that such a sequence
897 will extend lo_m among the events in S_2 .

898 To show that merge follows the sequence π for S_2 , we now instantiate BOTTOMUPTEMPLATE for
899 the case where L_1^a and L_2^a are empty (i.e. S_3 has already been linearized) in the rule BOTTOM-0-OP
900 in Fig. 12. Following the structure of π in Fig. 11, e_\top would be the event $e_m^\top \in L_\top^a$. Note that since
901 e_m^\top is an LCA event, it will be present in both states being merged. BOTTOMUP-0-OP then allows
902 this event to be linearized first at the end.

904 *Example 4.5.* Following on from Example 4.3 for the execution in Fig. 6 for the merge of v_3 and
905 v_4 , after BOTTOMUP-1-OP is applied to linearize e_3 , the states to be merged would be the versions
906 v_1 and v_4 (with LCA v_1), both of whose last operation is e_1 . Hence, BOTTOMUP-0-OP would be
907 applicable, which would linearize e_1 .

908 After applying BOTTOMUP-0-OP to linearize the LCA event e_m^\top , we then need to linearize events
909 in $L_1^b(e_m^\top) \cup L_2^b(e_m^\top)$. However, the event e_m^\top has already been linearized, so none of the events
910 in $L_1^b(e_m^\top) \cup L_2^b(e_m^\top)$ appear lo_m after an LCA event. This scenario can now be handled using
911 BOTTOMUP-2-OP (if both $L_1^b(e_m^\top)$ and $L_2^b(e_m^\top)$ are non-empty) or BOTTOMUP-1-OP (if one of 2 sets
912 is empty). These rules will appropriately linearize the events in $L_1^b(e_m^\top) \cup L_2^b(e_m^\top)$ taking into
913 account the rc relation for concurrent events and vis relation for non-concurrent events. Once
914 $L_1^b(e_m^\top) \cup L_2^b(e_m^\top)$ becomes empty, we then encounter the next LCA event in L_\top^a , which can again be
915 linearized using BOTTOMUP-0-OP.

916 The three instantiations of BOTTOMUPTEMPLATE can thus be repeatedly applied to linearize the
917 rest of the events in S_2 . Following this, all the local events would have been linearized, leaving only
918 the LCA events in S_1 . This would result in all three arguments to merge being equal, in which case
919 we can use the MERGEIDEMPOTENCE rule in Fig. 12. Using MERGEIDEMPOTENCE, we can equate the
920 output of merge to it's argument, which has already been assumed to be appropriately linearized.

921 In order to avoid mirrored versions of BOTTOMUP-2-OP and BOTTOMUP-1-OP where the second
922 and third arguments are swapped, we also require the MERGECOMMUTATIVITY property in Fig. 12.
923 We now state our soundness theorem linking the various properties with RA-linearizability of
924 MRDT:

925 THEOREM 4.6. *If an MRDT \mathcal{D} satisfies BOTTOMUP-2-OP, BOTTOMUP-1-OP, BOTTOMUP-0-OP,
926 MERGEIDEMPOTENCE and MERGECOMMUTATIVITY, then \mathcal{D} is linearizable.*

927 The proof closely follows the informal arguments that we have presented in this sub-section,
928 using induction on the size of the various sets $L_1^a, L_2^a, L_1^b \cup L_2^b, L_\top^a$.

4.2 Automated Verification

While we have identified the sufficient conditions to show RA-linearizability of an MRDT using bottom-up linearization, proving these conditions for arbitrary MRDTs is not straightforward. Further, while the BOTTOMUP-X-OP properties as shown in the previous sub-section had universal quantification over MRDT states l, a, b , in general, for proving RA-linearizability, we only need to show these properties for feasible states that may arise during an actual execution.

We now leverage the fact that the feasible states would have been obtained through linearization of the visible events at the respective versions. In particular, we can characterize the states on which merge can be invoked through the various events sets $L_1^a, L_2^a, L_1^b, L_2^b, L_T^a, L_T^b$ that we had defined in the previous sub-section. We only need to prove the BOTTOMUP-X-OP properties for states which have been obtained through linearizations of events in these event sets. For this purpose, we propose an induction scheme which establishes the required properties while traversing the event sets as depicted in Fig. 11 in a top-down fashion.

Table 1. Induction scheme for BOTTOMUPTEMPLATE. For clarity, we use \cdot for function composition, and μ for merge.

VC Name	Pre-condition	Post-condition
$\psi_{\text{base}}^{L_T^b}$		$\mu(\pi_0(\sigma_0), \pi_1(\sigma_0), \pi_2(\sigma_0)) = \pi \cdot \mu(\pi'_0(\sigma_0), \pi'_1(\sigma_0), \pi'_2(\sigma_0))$
$\psi_{\text{ind}}^{L_T^b}$	$\mu(\pi_0(l), \pi_1(l), \pi_2(l)) = \pi \cdot \mu(\pi'_0(l), \pi'_1(l), \pi'_2(l))$	$\mu(\pi_0 \cdot e_T, \pi_1 \cdot e_T, \pi_2 \cdot e_T) = \pi \cdot \mu(\pi'_0 \cdot e_T, \pi'_1 \cdot e_T, \pi'_2 \cdot e_T)$
$\psi_{\text{ind}}^{L_T^a}$	$\exists e. e \xrightarrow{rc} e_T$ $\mu(\pi_0(l), \pi_1(a), \pi_2(b)) = \pi \cdot \mu(\pi'_0(l), \pi'_1(a), \pi'_2(b))$	$\mu(\pi_0 \cdot e_T, \pi_1 \cdot e_T(a), \pi_2 \cdot e_T(b)) = \pi \cdot \mu(\pi'_0 \cdot e_T, \pi'_1 \cdot e_T(a), \pi'_2 \cdot e_T(b))$
$\psi_{\text{ind1}}^{L_1^b}$	$e_b \xrightarrow{rc} e_T$ $\mu(\pi_0 \cdot e_T(l), \pi_1 \cdot e_T(a), \pi_2 \cdot e_T(b)) = \pi \cdot \mu(\pi'_0 \cdot e_T(l), \pi'_1 \cdot e_T(a), \pi'_2 \cdot e_T(b))$	$\mu(\pi_0 \cdot e_T(l), \pi_1 \cdot e_T \cdot e_b(a), \pi_2 \cdot e_T(b)) = \pi \cdot \mu(\pi'_0 \cdot e_T(l), \pi'_1 \cdot e_T \cdot e_b(a), \pi'_2 \cdot e_T(b))$
$\psi_{\text{ind2}}^{L_1^b}$	$e_b \xrightarrow{rc} e_T \wedge \neg e \rightleftharpoons e_b$ $\mu(\pi_0 e_T(l), \pi_1 e_T e_b(a), \pi_2 e_T(b)) = \pi \cdot \mu(\pi'_0 e_T(l), \pi'_1 e_T e_b(a), \pi'_2 e_T(b))$	$\mu(\pi_0 \cdot e_T(l), \pi_1 \cdot e_T \cdot e_b \cdot e(a), \pi_2 \cdot e_T(b)) = \pi \cdot \mu(\pi'_0 \cdot e_T(l), \pi'_1 \cdot e_T \cdot e_b \cdot e(a), \pi'_2 \cdot e_T(b))$
$\psi_{\text{ind1}}^{L_1^a}$	$e_b \xrightarrow{rc} e_T$ $\mu(\pi_0 \cdot e_T(l), \pi_1 \cdot e_T(a), \pi_2 \cdot e_T(b)) = \pi \cdot \mu(\pi'_0 \cdot e_T(l), \pi'_1 \cdot e_T(a), \pi'_2 \cdot e_T(b))$	$\mu(\pi_0 \cdot e_T(l), \pi_1 \cdot e_T(a), \pi_2 \cdot e_T \cdot e_b(b)) = \pi \cdot \mu(\pi'_0 \cdot e_T(l), \pi'_1 \cdot e_T(a), \pi'_2 \cdot e_T \cdot e_b(b))$
$\psi_{\text{ind2}}^{L_2^b}$	$e_b \xrightarrow{rc} e_T \wedge \neg e \rightleftharpoons e_b$ $\mu(\pi_0 e_T(l), \pi_1 e_T e_b(a), \pi_2 e_T(b)) = \pi \cdot \mu(\pi'_0 e_T(l), \pi'_1 e_T e_b(a), \pi'_2 e_T(b))$	$\mu(\pi_0 \cdot e_T(l), \pi_1 \cdot e_T \cdot e_b(a), \pi_2 \cdot e_T \cdot e_b \cdot e(b)) = \pi \cdot \mu(\pi'_0 \cdot e_T(l), \pi'_1 \cdot e_T \cdot e_b(a), \pi'_2 \cdot e_T \cdot e_b \cdot e(b))$
$\psi_{\text{ind}}^{L_1^a}$	$\mu(\pi_0(l), \pi_1(a), \pi_2(b)) = \pi \cdot \mu(\pi'_0(l), \pi'_1(a), \pi'_2(b))$	$\mu(\pi_0(l), \pi_1 \cdot e(a), \pi_2(b)) = \pi \cdot \mu(\pi'_0(l), \pi'_1 \cdot e(a), \pi'_2(b))$
$\psi_{\text{ind}}^{L_2^a}$	$\mu(\pi_0(l), \pi_1(a), \pi_2(b)) = \pi \cdot \mu(\pi'_0(l), \pi'_1(a), \pi'_2(b))$	$\mu(\pi_0(l), \pi_1(a), \pi_2 \cdot e(b)) = \pi \cdot \mu(\pi'_0(l), \pi'_1(a), \pi'_2 \cdot e(b))$

Here, we present the induction scheme for the generic BOTTOMUPTEMPLATE rule. The scheme can then be instantiated for all the three BOTTOMUP-X-OP rules. Table 1 contains the verification conditions corresponding to the base case and inductive case over the different event sets. Every VC has the form (pre-condition \implies post-condition), and all variables are universally quantified. Our goal is to show the BOTTOMUPTEMPLATE rule for all feasible MRDT states l, a, b , where l is the state of the LCA of a and b . Let L_T, L_1, L_2 be the event sets corresponding to l, a, b respectively. We define the event sets $L_1^a, L_2^a, L_1^b, L_2^b, L_T^a, L_T^b$ in exactly the same manner as the previous sub-section, based on the linearization relation of the configuration obtained by the merge(l, a, b) transition. Note that the events in π_0, π_1, π_2 (used in the BOTTOMUPTEMPLATE rule) would also come from the above event sets, but in the following discussion, we freeze these events, i.e. all our assertions about the events sets will be modulo these events.

We start with the VC $\psi_{\text{base}}^{L_T^b}$, which corresponds to the case where every event set is empty. There is no pre-condition, and the post-condition requires BOTTOMUPTEMPLATE to hold on the initial MRDT

981 state σ_0 . For example, for the BOTTOMUP-2-OP rule, $\psi_{\text{base}}^{L_\top^b}$ VC would be $\text{merge}(\sigma_0, e_1(\sigma_0), e_2(\sigma_0)) =$
 982 $e_2(\text{merge}(\sigma_0, e_1(\sigma_0), \sigma_0))$, where $e_1 \xrightarrow{\text{rc}} e_2$ or e_1 and e_2 commute. Notice that e_1 and e_2 would be
 983 events in L_1^a and L_2^a , and our assertion about all event sets being empty is modulo these events.
 984

985 Next, the VC $\psi_{\text{ind}}^{L_\top^b}$ corresponds to the inductive case on L_\top^b , where we assume every event set
 986 except L_\top^b to be empty. The pre-condition corresponds to the inductive hypothesis, where we
 987 assume the property to hold for some event set L_\top^b , and the post-condition asserts that the property
 988 holds while adding another event e_\top to L_\top^b . Recall that L_\top^b corresponds to the LCA events which
 989 come lo before all local events. Since all the other event sets are empty, this translates to the same
 990 state l for all the three arguments to merge in the pre-condition, and applying the LCA event e_\top to
 991 all three arguments in the post-condition.

992 Next, we induct on the set L_\top^a , i.e. the set of LCA events which occur lo after a local event. The
 993 base case, where $|L_\top^a| = \emptyset$ exactly corresponds to the result of the induction on L_\top^b . The inductive
 994 case is covered by the VC $\psi_{\text{ind}}^{L_\top^a}$, which adds an LCA event e_\top to all three arguments of merge from
 995 pre-condition to post-condition. Notice that we also have another pre-condition which requires the
 996 existence of some event e which should come rc-before e_\top , which is necessary for e_\top to be in L_\top^a .
 997 The post-condition just adds a new LCA event e_\top . The events in $L_1^b(e_\top)$ and $L_2^b(e_\top)$ will be added
 998 by the next 4 VCs.

999 $\psi_{\text{ind1}}^{L_1^b}$ and $\psi_{\text{ind2}}^{L_1^b}$ add an event in L_1^b from the pre-condition to the post-condition. $\psi_{\text{ind1}}^{L_1^b}$ considers an
 1000 event e_b which occurs rc-before the LCA event e_\top . Notice that the pre-condition of $\psi_{\text{ind1}}^{L_1^b}$ is exactly
 1001 the same as the post-condition of $\psi_{\text{ind}}^{L_\top^a}$. In the post-condition of $\psi_{\text{ind1}}^{L_1^b}$, the event e_b is applied before
 1002 e_\top on the argument a to merge, thus reflecting that this is an event in L_1^b . $\psi_{\text{ind2}}^{L_1^b}$ adds an event $e \in L_1^b$
 1003 which does not commute with an existing event $e_b \in L_1^b$ (see the definition of L_1^b). $\psi_{\text{ind1}}^{L_2^b}$ and $\psi_{\text{ind2}}^{L_2^b}$
 1004 are analogous and do the same thing for the argument b to merge.
 1005

1006 Finally, $\psi_{\text{ind}}^{L_1^a}$ and $\psi_{\text{ind}}^{L_2^a}$ add events from L_1^a and L_2^a . The base cases for the two sets would exactly
 1007 correspond to the result of the induction carried out so far on the rest of the event sets. For the inductive
 1008 case, in $\psi_{\text{ind}}^{L_1^a}$ (resp. $\psi_{\text{ind}}^{L_2^a}$), a new event e is added on the second argument a (resp. third argument
 1009 b) from the pre-condition to the post-condition. This establishes the rule BOTTOMUPTEMPLATE for
 1010 any feasible input arguments to merge during any execution. We denote the set of VCs in Table 1
 1011 by $\psi^*(\text{BOTTOMUPTEMPLATE})$.
 1012
 1013

1014 **THEOREM 4.7.** *If an MRDT \mathcal{D} satisfies the VCs $\psi^*(\text{BOTTOMUP-2-OP})$, $\psi^*(\text{BOTTOMUP-1-OP})$,
 1015 $\psi^*(\text{BOTTOMUP-0-OP})$, MERGEIDEMPOTENCE and MERGECOMMUTATIVITY, then \mathcal{D} is linearizable.*
 1016

1017 5 Experimental Evaluation

1018 We have implemented our verification technique in the F^{*} programming language and verified
 1019 several MRDTs using it. We also extracted OCaml code from the verified implementations and ran
 1020 them as part of Irmin [9], a Git-like distributed database which follows the MRDT system model
 1021 described in §3. This distinguishes our work from prior works in automated RDT verification [16]
 1022 which focuses on verifying abstract models rather than actual implementations.
 1023

1024 Our framework in F^{*} consists of an F^{*} interface that defines signatures for an MRDT implemen-
 1025 tation (Fig. 2) and the VCs described in Table 1; these are encoded as F^{*} lemmas. This interface
 1026 contains 200 lines of F^{*} code. An MRDT developer instantiates the interface with their specific
 1027 MRDT implementation and calls upon F^{*} to prove the lemmas (i.e., the VCs). Once this is done, our
 1028 metatheory (Theorem 4.7) guarantees that the MRDT implementation is linearizable.
 1029

Table 2. Verified MRDTs. * denotes MRDT implementations not present in prior work.

MRDT	rc Policy	#LOC	Verification Time (s)
Increment-only counter [12]	none	6	0.72
PN counter [23]	none	10	1.64
Enable-wins flag*	disable \xrightarrow{rc} enable	30	29.80
Disable-wins flag*	enable \xrightarrow{rc} disable	30	37.91
Grows-only set [12]	none	6	0.45
Grows-only map [23]	none	11	4.65
OR-set [23]	rem _a \xrightarrow{rc} add _a	20	4.53
OR-set (efficient)*	rem _a \xrightarrow{rc} add _a	34	660.00
Remove-wins set*	add _a \xrightarrow{rc} rem _a	22	9.60
Set-wins map*	del _k \xrightarrow{rc} set _k	20	5.06
Replicated Growable Array [1]	none	13	1.51
Optional register*	unset \xrightarrow{rc} set	35	200.00
Multi-valued Register*	none	7	0.65
JSON-style MRDT*	Fig. 13	26	148.84

We instantiate the interface with MRDT implementations of several datatypes such as counter, flag, set, map, and list (Table 2). All the results were obtained on a Intel®Xeon®Gold 5120 x86-64 machine running Ubuntu 22.04 with 64GB of main memory. While some of the MRDTs have been taken from previous works [1, 12, 23] or translated from their CRDT counterparts, we also develop some new implementations, denoted by * in Table 2. We also uncovered bugs in previous MRDT implementations (Enable-wins flag and Efficient OR-set) from [23], which we fixed (more details in §5.2). We note that in all our experiments, all the VCs were automatically discharged by F* in a reasonable amount of time.

While our approach ensures that the MRDT implementations are verified in the F* framework, it is important to note that the user is obligated to trust the F* language implementation, the extraction mechanism, the OCaml language implementation, the OCaml runtime, and the hardware.

We now highlight several notable features about our verified MRDTs. We have designed and developed the first correct implementations of both an enable-wins and disable-wins flag MRDT. Our implementation of efficient OR-set maintains a per-replica, per-element counter instead of adding different versions of the same element (as done by the OR-set implementation of Fig. 2), thus matching the theoretical lower bound in terms of space-efficiency for any OR-set CRDT implementation (as proved in [4]). We have developed the first known MRDT implementation of a remove-wins set datatype. Finally, as a demonstration of vertical compositionality, we have developed a JSON MRDT which is composed of several component MRDTs, with its correctness guarantee being directly derived from the correctness of the component MRDTs.

5.1 Case study: A verified polymorphic JSON-style MRDT

JSON is a notable example of a data type which is composed of several other datatypes. JSON is widely used as a data interchange format in many databases and web services [10]. Our JSON MRDT is modeled as an unordered collection of key/value pairs, where the values can be any primitive types, such as counter, list, etc., or they can be JSON type themselves. We assume that keys are update-only; that is, key-value mappings can be added and modified, but once a key is added, it cannot be deleted. Previous works, such as Automerge [2], have developed similar JSON-style CRDT models. However, these models are monomorphic, which means that the data type of the values must be known in advance. Our goal is to develop a more generic JSON-style MRDT that

1079 supports polymorphic values, i.e. we leave the value data type as an abstract type which can be
 1080 instantiated with any concrete MRDT.

1081 1: $\Sigma_{\text{json}} : (k : (\text{string} \times \Omega)) \rightarrow \Sigma_{\text{snd}(k)}$
 1082 2: $O_{\text{json}} = \{\text{set}(k, o) \mid o \in O_{\text{snd}(k)}\}$
 1083 3: $Q_{\text{json}} = \{\text{get}(k, q) \mid q \in Q_{\text{snd}(k)}\}$
 1084 4: $\sigma_{0_{\text{json}}} = \lambda(k : \text{string} \times \Omega). \sigma_{0_{\text{snd}(k)}}$
 1085 5: $\text{do}(\sigma, t, r, \text{set}(k, o)) = \sigma[k \mapsto o(\sigma(k), t, r)]$
 1086 6: $\text{merge}_{e_{\text{json}}}(\sigma_{\top}, \sigma_1, \sigma_2) = \lambda(k : \text{string} \times \Omega). \text{merge}_{e_{\text{snd}(k)}}(\sigma_{\top}(k), \sigma_1(k), \sigma_2(k))$
 1087 7: $\text{query}_{j_{\text{json}}}(\sigma, \text{get}(k, q)) = \text{query}_{\text{snd}(k)}(\sigma(k), q)$
 1088 8: $\text{rc}_{\text{json}} = \{(\text{set}(k_1, o_1), \text{set}(k_2, o_2)) \in O_{\text{json}} \times O_{\text{json}} \mid k_1 = k_2 \wedge (o_1, o_2) \in \text{rc}_{\text{snd}(k_1)}\}$

Fig. 13. JSON-style MRDT implementation

1091 Fig. 13 shows the implementation of the JSON MRDT. It uses a map to maintain association
 1092 between keys and values. Notice that the key is a tuple consisting of the identifier string and an
 1093 MRDT type $\alpha \in \Omega$ which denotes the type of the value. The type α can be any arbitrary MRDT
 1094 with implementation $\mathcal{D}_{\alpha} = (\Sigma_{\alpha}, \sigma_{0_{\alpha}}, \text{merge}_{e_{\alpha}}, \text{query}_{\alpha}, \text{rc}_{\alpha})$. Different key strings can now map to
 1095 different value MRDT types. We also allow overloading: the same key string can be associated
 1096 with multiple values of different types. The JSON MRDT allows update operations of the form
 1097 $\text{set}(k, o)$ where o is an operation of the underlying value MRDT associated with the key k . $\text{set}(k, o)$
 1098 simply applies the operation o on the value associated with k , leaving the other key-value pairs
 1099 unchanged. The JSON merge calls the underlying MRDT merge on the values associated with
 1100 each key. The query operation of the form $\text{get}(k, q)$ retrieves the value associated with k in σ and
 1101 applies the query operation q of the underlying data type to it. The conflict resolution policy of
 1102 JSON operations (rc_{json}) depends on the conflict resolution of the value types when two operations
 1103 update the same key (i.e. same identifier and value type). Every other pair of JSON operations
 1104 commute with each other.

1105 Notably, the proof of RA-linearizability of the JSON MRDT is directly derived from the proofs of
 1106 the underlying value MRDT types. If all the MRDTs in Ω are linearizable, then the JSON MRDT
 1107 is also linearizable. We have proved all the VCs for the JSON MRDT in F^* by using the VCs of
 1108 the underlying value MRDTs. We can now instantiate Ω with any set of verified MRDTs, thereby
 1109 obtaining the verified JSON MRDT for free.

1111 5.2 Buggy MRDT Implementation in [23]

1112 We now present some details of one of the buggy MRDTs, Enable-wins flag, that we discovered using our framework in
 1113 Soundarapandian et al. [23]. The state of the enable-wins flag MRDT consists of a pair: a counter and a flag. The counter
 1114 tracks the number of the enable events, while the flag is set to true on an enable event. The desired specification for this
 1115 flag is that it should be true when there is at least one enable event not visible to any disable event. In our framework, we
 1116 can express this specification as $\text{disable} \xrightarrow{\text{rc}} \text{enable}$, linearizing the enable operation after a concurrent disable. When we
 1117 attempted to verify this implementation in our framework,
 1118 we discovered that one of the VCs, $\psi_{\text{ind2-top}}^{L^b}$, was failing. Our
 1119 investigation revealed that the implementation violated the
 1120 specification, with the bug appearing in an execution with intermediate merges.

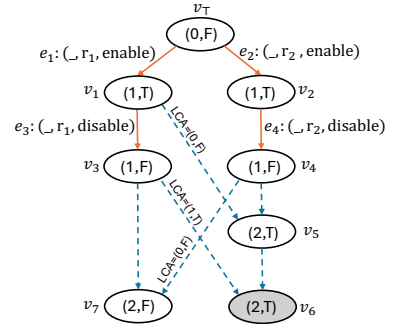


Fig. 14. An enable-wins flag execution

1128 Consider the execution depicted in Fig. 14. When merging versions v_3 and v_5 (with LCA v_1),
 1129 since the counter value of v_5 is greater than v_1 , the flag in the merged version v_6 is set to true.
 1130 However, this contradicts the Enable-wins flag specification, which states that the flag should
 1131 be true only when there is an enable event that is not visible to any disable event. All enable
 1132 events in the execution are disabled by subsequent disable events on their individual replicas,
 1133 yet the flag is true at v_6 . Notice that the version v_5 is obtained due to an intermediate merge. We
 1134 discovered that Soundarapandian et al. [23] had an implementation bug in the framework. The
 1135 framework expects a simulation relation from the MRDT developer, in addition to the specification
 1136 and the implementation. This simulation relation serves as a proof artefact. Soundarapandian et al.
 1137 [23] checks whether the developer-provided simulation relation is valid and the bug occurred
 1138 in the validity-checking procedure. Due to this, Soundarapandian et al. [23] admitted the buggy
 1139 enable-wins flag implementation⁵.

1140 We further note that this buggy implementation does not even satisfy strong eventual consistency.
 1141 In Fig. 14, merging v_3 and v_4 results in v_7 , where the flag is false. Note that both versions v_6 and
 1142 v_7 have observed the same set of updates on both replicas, yet they lead to divergent states. This
 1143 violates strong eventual consistency. We fixed this implementation by maintaining a counter-flag
 1144 pair for every replica, i.e. changing the state to a map from replica-IDs to counter-flag pair.

1145

1146 5.3 Verifying state-based CRDTs

1147 Although the development in the paper so far has focused on verifying MRDTs, we note that our
 1148 framework can also directly verify state-based CRDTs. The only difference between the two is that
 1149 state-based CRDTs do not maintain the LCA, and merge is a binary function. Our VCs (Table 1)
 1150 can be directly applied on state-based CRDTs, by simply ignoring the LCA argument for all merges.
 1151 Note that while the merge function in state-based CRDTs does not use the LCA, our VCs still use
 1152 the LCA to determine whether an event is local or common to both replicas, and appropriately
 1153 linearize events taking into account both rc and vis relations. The entire set of VCs retrofitted for
 1154 state-based CRDTs can be found in Table 4. We have also successfully implemented and verified 7
 1155 state-based CRDTs in our framework: Increment-only counter, PN counter, Observed-Remove set,
 1156 Two-Phase set, Grows-only set, Grows-only map and Multi-valued register.

1157

1158 5.4 Limitations

1159 Our framework is currently unable to verify some MRDT implementations such as Queue from
 1160 previous works [12, 23]. The Queue MRDT follows at-least-once semantics for dequeues, which
 1161 allows concurrent dequeue operations to return the same element from the queue, thereby having
 1162 the effect of a single dequeue. Such an implementation is clearly not linearizable as per our definition,
 1163 since we cannot omit any event while constructing the linearization. It would be possible to modify
 1164 our notion of linearization to also allow events to be omitted; we leave this investigation as part of
 1165 future work. Our verification technique is also not complete, but in practice we have been able to
 1166 successfully verify all MRDT implementations (except Queue) from earlier works.

1167

1168 6 Related Work and Conclusion

1169 Reconciling concurrent updates is a challenging problem in distributed systems. CRDTs [3, 20, 21]
 1170 (and more recently MRDTs) have emerged as a principled approach for building correct and
 1171 efficient replicated implementations. Numerous works have focused on specifying and verifying
 1172 CRDTs [1, 4, 7, 8, 13, 15–18, 25, 26]. Op-based CRDTs have a considerably different system model
 1173

1174

1175 ⁵Buggy implementation can be found in §A.3

1176

1177 than MRDTs, where every operation instance at a replica is individually sent to other replicas.
1178 Hence, verification efforts targeting them [7, 15–17, 25] are mostly orthogonal to our work.

1179 The system model of state-based CRDTs bears a close resemblance to the MRDT model, as it also
1180 requires a merge function to be implemented for reconciling concurrent updates. However, there
1181 are stringent restrictions to ensure convergence and consistency for state-based CRDTs, requiring
1182 the CRDT states to form a join-semilattice, every update to be monotonic and the merge function
1183 to be the join operation of the lattice. The three algebraic properties of a semi-lattice: idempotence,
1184 commutativity, and associativity guarantee convergence.

1185 Some CRDT works focus solely on ensuring convergence without addressing functional correct-
1186 ness. For instance, Porre et al. [18] does not fully capture the user intent when verifying state-based
1187 CRDTs. Consider a Counter CRDT with only an increment operation and an *incorrect* merge func-
1188 tion that ignores its input states and always returns 0. Such an implementation is still convergent.
1189 However, it clearly does not capture the developer intent, which is that the value of the counter
1190 should be equal to the number of increment operations. Functional correctness is as important as
1191 convergence for replicated data types. Our framework addresses both by couching both in terms of
1192 RA-linearizability. We will flag the above implementation as incorrect, since the state after merge
1193 cannot be obtained by linearizing the operations performed on both the replicas.

1194 In the context of CRDTs, Wang et al. [25] proposed the notion of replication-aware linearizability,
1195 which requires all replicas to have a state which can be obtained by linearizing the update operations
1196 visible to the replica according to the sequential specification. However, they do not propose any
1197 automated verification methodology for RA-linearizability. Further, though the main paper Wang
1198 et al. [25] focuses on op-based CRDTs, the extended version Enea et al. [5] does address state-based
1199 CRDTs, but they also require a semi-lattice-based formulation of the CRDT states for proving
1200 RA-linearizability.

1201 Few works [11, 23] have explored the problem of verifying MRDT implementations. Kaki et al. [11]
1202 only focus on verifying convergence, but not functional correctness. Moreover, they significantly
1203 restrict the underlying system model by synchronizing all merge operations, which as mentioned
1204 in the paper itself could lead to longer convergence times. Soundarapandian et al. [23] verify both
1205 convergence and functional correctness, and their system model does not require merges to be
1206 synchronized. However, their approach is not fully automated, and requires developers to provide
1207 a simulation relation linking concrete MRDT states with an abstract state which is based on a
1208 event-based declarative model. Their specification language is also based on an event-based model
1209 and is not very intuitive or developer-friendly. A few MRDT implementations from [23] were found
1210 to be buggy, and these errors were due to faulty simulation relations.

1211 To conclude, in this work, we present the first, fully-automated verification methodology for
1212 MRDTs. We introduce the notion of replication-aware linearizability for MRDTs, as well as a simple
1213 specification framework based on ordering non-commutative update operations. We identify certain
1214 restrictions on the specification to ensure existence of a consistent linearization. We then leverage
1215 the definition of replication-aware linearizability to propose an automated verification methodology
1216 based on induction on operation sequences. We have successfully applied the technique on a number
1217 of complex MRDTs. While the foundations have been laid in this work, we believe there is a lot
1218 of scope for enriching the technique even further by considering more complex linearization
1219 strategies.

1221 References

- 1222 [1] Hagit Attiya, Sebastian Burckhardt, Alexey Gotsman, Adam Morrison, Hongseok Yang, and Marek Zawirski. 2016.
1223 Specification and Complexity of Collaborative Text Editing. In *Proceedings of the 2016 ACM Symposium on Principles*
1224 *of Distributed Computing, PODC 2016, Chicago, IL, USA, July 25–28, 2016*, George Giakkoupis (Ed.). ACM, 259–268.

- 1226 <https://doi.org/10.1145/2933057.2933090>
- 1227 [2] Automerge. 2022. Automerge. <https://automerge.org/>
- 1228 [3] Annette Bieniusa, Marek Zawirski, Nuno M. Pregoia, Marc Shapiro, Carlos Baquero, Valter Balegas, and Sérgio
1229 Duarte. 2012. An optimized conflict-free replicated set. *CoRR* abs/1210.3368 (2012). arXiv:1210.3368 <http://arxiv.org/abs/1210.3368>
- 1230 [4] Sebastian Burckhardt, Alexey Gotsman, Hongseok Yang, and Marek Zawirski. 2014. Replicated data types: specification,
1231 verification, optimality. In *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages,*
1232 *POPL '14, San Diego, CA, USA, January 20-21, 2014*, Suresh Jagannathan and Peter Sewell (Eds.). ACM, 271–284.
1233 <https://doi.org/10.1145/2535838.2535848>
- 1234 [5] Constantin Enea, Suha Orhun Mutluergil, Gustavo Petri, and Chao Wang. 2019. Replication-Aware Linearizability.
1235 *CoRR* abs/1903.06560 (2019). arXiv:1903.06560 <http://arxiv.org/abs/1903.06560>
- 1236 [6] Git. 2021. Git: A distributed version control system. <https://git-scm.com/>
- 1237 [7] Victor B. F. Gomes, Martin Kleppmann, Dominic P. Mulligan, and Alastair R. Beresford. 2017. Verifying strong
1238 eventual consistency in distributed systems. *Proc. ACM Program. Lang.* 1, OOPSLA (2017), 109:1–109:28. <https://doi.org/10.1145/3133933>
- 1239 [8] Alexey Gotsman, Hongseok Yang, Carla Ferreira, Mahsa Najafzadeh, and Marc Shapiro. 2016. 'Cause I'm strong enough:
1240 reasoning about consistency choices in distributed systems. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT*
1241 *Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, Rastislav
1242 Bodik and Rupak Majumdar (Eds.). ACM, 371–384. <https://doi.org/10.1145/2837614.2837625>
- 1243 [9] Irmin. 2021. Irmin: A distributed database built on the principles of Git. <https://irmin.org/>
- 1244 [10] Json. [n. d.]. Json: A lightweight data-interchange format. <https://www.json.org/>
- 1245 [11] Gowtham Kaki, Prasanth Prahladan, and Nicholas V. Lewchenko. 2022. RunTime-assisted convergence in replicated data
1246 types. In *PLDI '22: 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation,*
1247 *San Diego, CA, USA, June 13 - 17, 2022*, Ranjit Jhala and Isil Dillig (Eds.). ACM, 364–378. <https://doi.org/10.1145/3519939.3523724>
- 1248 [12] Gowtham Kaki, Swarn Priya, K. C. Sivaramakrishnan, and Suresh Jagannathan. 2019. Mergeable replicated data types.
1249 *Proc. ACM Program. Lang.* 3, OOPSLA (2019), 154:1–154:29. <https://doi.org/10.1145/3360580>
- 1250 [13] Shadaj Laddad, Conor Power, Mae Milano, Alvin Cheung, and Joseph M. Hellerstein. 2022. Katara: synthesizing CRDTs
1251 with verified lifting. *Proc. ACM Program. Lang.* 6, OOPSLA2 (2022), 1349–1377. <https://doi.org/10.1145/3563336>
- 1252 [14] Leslie Lamport. 1978. Time, Clocks, and the Ordering of Events in a Distributed System. *Commun. ACM* 21, 7 (1978),
1253 558–565. <https://doi.org/10.1145/359545.359563>
- 1254 [15] Yiyun Liu, James Parker, Patrick Redmond, Lindsey Kuper, Michael Hicks, and Niki Vazou. 2020. Verifying replicated
1255 data types with typeclass refinements in Liquid Haskell. *Proc. ACM Program. Lang.* 4, OOPSLA (2020), 216:1–216:30.
1256 <https://doi.org/10.1145/3428284>
- 1257 [16] Kartik Nagar and Suresh Jagannathan. 2019. Automated Parameterized Verification of CRDTs. In *Computer Aided*
1258 *Verification - 31st International Conference, CAV 2019, New York City, NY, USA, July 15-18, 2019, Proceedings, Part*
1259 *II (Lecture Notes in Computer Science, Vol. 11562)*, Isil Dillig and Serdar Tasiran (Eds.). Springer, 459–477. https://doi.org/10.1007/978-3-030-25543-5_26
- 1260 [17] Sreeja S. Nair, Gustavo Petri, and Marc Shapiro. 2020. Proving the Safety of Highly-Available Distributed Objects.
1261 In *Programming Languages and Systems - 29th European Symposium on Programming, ESOP 2020, Held as Part of the*
1262 *European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020, Proceedings*
1263 *(Lecture Notes in Computer Science, Vol. 12075)*, Peter Müller (Ed.). Springer, 544–571. https://doi.org/10.1007/978-3-030-44914-8_20
- 1264 [18] Kevin De Porre, Carla Ferreira, and Elisa Gonzalez Boix. 2023. VeriFx: Correct Replicated Data Types for the Masses. In
1265 *37th European Conference on Object-Oriented Programming, ECOOP 2023, July 17-21, 2023, Seattle, Washington, United*
1266 *States (LIPIcs, Vol. 263)*, Karim Ali and Guido Salvaneschi (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik,
1267 9:1–9:45. <https://doi.org/10.4230/LIPICS.ECOOP.2023.9>
- 1268 [19] Riak. 2021. Resilient NoSQL Databases. <https://riak.com/>
- 1269 [20] Hyun-Gul Roh, Myeongjae Jeon, Jinsoo Kim, and Joonwon Lee. 2011. Replicated abstract data types: Building blocks
1270 for collaborative applications. *J. Parallel Distributed Comput.* 71, 3 (2011), 354–368. <https://doi.org/10.1016/J.JPDC.2010.12.006>
- 1271 [21] Marc Shapiro, Nuno M. Pregoia, Carlos Baquero, and Marek Zawirski. 2011. Conflict-Free Replicated Data Types. In
1272 *Stabilization, Safety, and Security of Distributed Systems - 13th International Symposium, SSS 2011, Grenoble, France,*
1273 *October 10-12, 2011. Proceedings (Lecture Notes in Computer Science, Vol. 6976)*, Xavier Défago, Franck Petit, and Vincent
1274 Villain (Eds.). Springer, 386–400. https://doi.org/10.1007/978-3-642-24550-3_29
- [22] Marc Shapiro, Nuno Pregoia, Carlos Baquero, and Marek Zawirski. 2011. *A comprehensive study of Convergent and Commutative Replicated Data Types*. Technical Report RR-7506. Inria – Centre ParisRocquencourt; INRIA. 50 pages.

- 1275 <https://hal.inria.fr/fr-00555588>
- 1276 [23] Vimala Soundarapandian, Adharsh Kamath, Kartik Nagar, and K. C. Sivaramakrishnan. 2022. Certified mergeable
1277 replicated data types. In *PLDI '22: 43rd ACM SIGPLAN International Conference on Programming Language Design
1278 and Implementation, San Diego, CA, USA, June 13 - 17, 2022*, Ranjit Jhala and Isil Dillig (Eds.). ACM, 332–347. <https://doi.org/10.1145/3519939.3523735>
- 1279 [24] Nikhil Swamy, Catalin Hritcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan
1280 Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean Karim Zinzindohoue, and Santiago Zanella
1281 Béguelin. 2016. Dependent types and multi-monadic effects in F. In *Proceedings of the 43rd Annual ACM SIGPLAN-
1282 SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*,
1283 Rastislav Bodik and Rupak Majumdar (Eds.). ACM, 256–270. <https://doi.org/10.1145/2837614.2837655>
- 1284 [25] Chao Wang, Constantin Enea, Suha Orhun Mutluergil, and Gustavo Petri. 2019. Replication-aware linearizability.
1285 In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI
1286 2019, Phoenix, AZ, USA, June 22-26, 2019*, Kathryn S. McKinley and Kathleen Fisher (Eds.). ACM, 980–993. <https://doi.org/10.1145/3314221.3314617>
- 1287 [26] Peter Zeller, Annette Bieniusa, and Arnd Poetzsch-Heffter. 2014. Formal Specification and Verification of CRDTs. In
1288 *Formal Techniques for Distributed Objects, Components, and Systems - 34th IFIP WG 6.1 International Conference, FORTE
1289 2014, Held as Part of the 9th International Federated Conference on Distributed Computing Techniques, DisCoTec 2014,
1290 Berlin, Germany, June 3-5, 2014. Proceedings (Lecture Notes in Computer Science, Vol. 8461)*, Erika Abraham and Catuscia
1291 Palamidessi (Eds.). Springer, 33–48. https://doi.org/10.1007/978-3-662-43613-4_3

1292 A Appendix

1293 A.1 Proofs of §3

1294 **Lemma 3.2** Given a configuration $C = \langle N, H, L, G, vis \rangle$ reachable in some execution $\tau \in \llbracket \mathcal{S}_{\mathcal{D}} \rrbracket$ and
1295 two versions $v_1, v_2 \in \text{dom}(N)$, if v_{\top} is the LCA of v_1 and v_2 in G , then $L(v_{\top}) = L(v_1) \cap L(v_2)$.

1296 **PROOF.** If $(v, v') \in E$, then $L(v) \subseteq L(v')$. This is because either $L(v') = L(v) \cup \{e\}$ for some event
1297 e due to the apply transition, or $L(v') = L(v) \cup L(v'')$ due to the merge transition.

1298 Hence, if $(v, v') \in E^*$, then $L(v) \subseteq L(v')$.

1299 Since $(v_{\top}, v_1) \in E^*$ and $(v_{\top}, v_2) \in E^*$, hence $L(v_{\top}) \subseteq L(v_1)$ and $L(v_{\top}) \subseteq L(v_2)$. Hence, $L(v_{\top}) \subseteq$
1300 $L(v_1) \cap L(v_2)$.

1301 Consider vertices u, w and event e such that $(u, w) \in E$, $e \notin L(u)$, $e \in L(w)$ and in-degree of
1302 w is 1. Then w is called the *generator* vertex of event e . Note that there will always be a unique
1303 generator vertex for each event.

1304 **PROPOSITION A.1.** For all versions v , events e , if $e \in L(v)$, and w is the generator version of e , then
1305 $(w, v) \in E^*$.

1306 Consider $e \in L(v_1) \cap L(v_2)$. Then if w is the generator version of e , by Proposition A.1 $(w, v_1) \in E^*$
1307 and $(w, v_2) \in E^*$. Then, by definition of LCA, $(w, v_{\top}) \in E^*$. Hence, $L(w) \subseteq L(v_{\top})$. This implies that
1308 $e \in L(v_{\top})$. Thus, $L(v_1) \cap L(v_2) \subseteq L(v_{\top})$.

1309 We now prove Proposition A.1. If v has in-degree 2, then suppose $(w_1, v) \in E$, $(w_2, v) \in E$ and
1310 $L(v) = L(w_1) \cup L(w_2)$. Then either $e \in L(w_1)$ or $e \in L(w_2)$. WLOG, suppose $e \in L(w_1)$. We now
1311 recursively apply Proposition A.1 on w_1 . Then, $(w, w_1) \in E^*$, which implies $(w, v) \in E^*$.

1312 If v has in-degree 1, then suppose $(u, v) \in E$. If $e \in L(u)$, we recursively apply Proposition A.1
1313 on u . If $e \notin L(u)$, then v itself is the generator version of e , and clearly, $(v, v) \in E^*$.

1314 Note that everytime we move backwards along an edge by recursively applying Proposition A.1,
1315 we are either decreasing the number of events in the source vertex, or the number of unvisited
1316 vertices in the graph while still retaining e . Since the graph is acyclic and finite, and the number of
1317 events are also finite, eventually, we will hit the generator version. \square

1318 **Recursive Merge Strategy:** For a given version graph $G = (V, E)$, for versions v_1, v_2 , if the LCA
1319 does not exist, then our strategy is to find potential LCAs. For each potential LCA v_p , $(v_p, v_1) \in E^*$,
1320 \square

1324 $(v_p, v_2) \in E^*$ and $\nexists v. (v, v_1) \in E^* \wedge (v, v_2) \in E^* \wedge (v_p, v) \in E^*$. Note that since the version graph is
 1325 rooted at the initial version v_0 , a common ancestor of any two versions v_1 and v_2 always exist. Let
 1326 V_c be the set of all common ancestors of v_1 and v_2 .

$$1327 \quad V_c = \{v \in V \mid (v, v_1) \in E^* \wedge (v, v_2) \in E^*\}$$

1328 For two common ancestors $v, v' \in V_c$, either there is a path between them or there isn't. If there
 1329 is a path, say $(v, v') \in E^*$, then v can neither be a potential or actual LCA. In this way, we eliminate
 1330 all common ancestors which cannot be potential or actual LCAs. Finally, we are left with the set
 1331 of potential LCAs V_p . Hence, for any $v, v' \in V_p$, $(v, v') \notin E^*$ and $(v', v) \notin E^*$. It is then clear to see
 1332 that if $V_p = \{v_\top\}$, i.e. V_p is singleton, then v_\top must be the actual LCA, because every other common
 1333 ancestor v must have been eliminated due to $(v, v_\top) \in E^*$.

1334 Otherwise, if V_p is not singleton, we pairwise invoke merge on every pair of versions in V_p . Note
 1335 that we would have to repeat the same merge strategy while merging any two versions in V_p . We now
 1336 show that if v_m is the version obtained by merging all the versions in V_p , then $L(v_m) = L(v_1) \cap L(v_2)$.
 1337 Since every version $v \in V_p$ is a common ancestor of v_1 and v_2 , $L(v) \subseteq L(v_1) \cap L(v_2)$, and hence
 1338 $L(v_m) \subseteq L(v_1) \cap L(v_2)$. Consider $e \in L(v_1) \cap L(v_2)$. Now, consider the generator version w of e . By
 1339 Proposition A.1, w is a common ancestor of v_1 and v_2 . Either $w \in V_p$, in which case by merging
 1340 w to get v_m , we would have $e \in L(v_m)$. Or else, w would have been eliminated, in which case
 1341 there will exist some version $v \in V_p$ such that $(w, v) \in E^*$. Hence, $e \in L(v)$, which implies $e \in L(v_m)$.

1342 **Lemma 3.4** Given a set of events \mathcal{E} , if $\text{lo} \subseteq \mathcal{E} \times \mathcal{E}$ is defined over every pair of non-commutative
 1343 events in \mathcal{E} , then for any two sequences π_1, π_2 which extend lo , for any state σ , $\pi_1(\sigma) = \pi_2(\sigma)$.

1344 **PROOF.** If $\pi_1 = \pi_2$, then the result trivially holds. Consider the first point of difference between
 1345 π_1 and π_2 .

$$1346 \quad \pi_1 = \tau.e_1.\tau_1, \pi_2 = \tau.e_2.\tau_2.$$

1347 Then e_1 must appear somewhere in τ_2 .

$$1348 \quad \pi_2 = \tau.e_2.\tau_3.e_1.\tau_4.$$

1349 We consider two cases here:

1350 **Case 1:** ($\tau_3 = \phi$)

1351 Since e_1 and e_2 are in different orders in π_1 and π_2 , neither $e_1 \xrightarrow{\text{lo}} e_2$ nor $e_2 \xrightarrow{\text{lo}} e_1$. Since lo is
 1352 defined over every pair of non-commutative events, but is not defined between e_1 and e_2 , they must
 1353 commute. Hence, we can flip e_2 and e_1 in π_2 , leading to the same state.

1354 **Case 2:** ($\tau_3 \neq \phi$)

1355 $\pi_2 = \tau.e_2.\tau_5.e_3.e_1.\tau_4$. Then in π_1 , e_3 is not present in τ , hence it must be present after e_1 . Now e_1 and
 1356 e_3 are in different orders in π_1 and π_2 , hence neither $e_1 \xrightarrow{\text{lo}} e_3$ nor $e_3 \xrightarrow{\text{lo}} e_1$.

1357 By the same argument as above applied on e_1 and e_2 , we can flip e_1 and e_3 in π_2 . We keep doing
 1358 this for all events in τ_5 until e_2 is adjacent to e_1 after which we can flip them. Thus we can change
 1359 π_2 such that e_1 will appear in the same position in π_1 . We can keep doing this until π_1 and π_2 are
 1360 identical. \square

1361 **Lemma 3.6** For an MRDT \mathcal{D} such that rc^+ is irreflexive, for any configuration C reachable in $\mathcal{S}_{\mathcal{D}}$,
 1362 lo_C^+ is irreflexive.

1363 To prove that lo_C^+ is irreflexive, we need to prove that there cannot be cycles formed out of lo_C
 1364 edges.

1365 **PROOF.** A cycle cannot be formed using only vis edges, as vis^+ is irreflexive. Similarly, a cycle
 1366 cannot be formed using only rc edges, as rc^+ is irreflexive. Therefore, any potential cycle must

1367

1368

1369

1370

1371

1372

1373 consist of adjacent $\xrightarrow{\text{rc}}$ and $\xrightarrow{\text{vis}}$ edges. Consider three events e_1, e_2, e_3 such that $e_1 \xrightarrow{\text{lo}} e_2 \xrightarrow{\text{lo}} e_3$. Since
 1374 $e_1 \xrightarrow{\text{lo}} e_2$, this implies $e_1 \xrightarrow{\text{rc}} e_2 \wedge e_1 \parallel_C e_2$. Given that $e_2 \xrightarrow{\text{vis}} e_3$, the relation $e_1 \xrightarrow{\text{lo}} e_2$ is not possible.
 1375 Thus, this case is also not feasible. Hence, there cannot be cycles formed out of lo_C edges. \square
 1376
 1377

1378 **Lemma 3.8** For an MRDT \mathcal{D} which satisfies $\text{RC-NON-COMM}(\mathcal{D})$ and $\text{COND-COMM}(\mathcal{D})$, for any
 1379 reachable configuration C in $\mathcal{S}_{\mathcal{D}}$, for any two sequences π_1, π_2 over E_C which extend lo_C , for any
 1380 state σ , $\pi_1(\sigma) = \pi_2(\sigma)$.
 1381

1382 **PROOF.** Consider the first point of difference between π_1 and π_2 .

1383 $\pi_1 = \tau.e_1.\tau_1, \pi_2 = \tau.e_2.\tau_2$.

1384 Then e_1 must appear somewhere in τ_2 .

1385 $\pi_2 = \tau.e_2.\tau_3.e_1.\tau_4$.

1386 We consider two cases here:

1387 **Case 1:** ($\tau_3 = \phi$)

1388 Since e_1 and e_2 are in different orders in π_1 and π_2 , neither $e_1 \xrightarrow{\text{lo}} e_2$ nor $e_2 \xrightarrow{\text{lo}} e_1$. If either $e_1 \xrightarrow{\text{vis}} e_2$
 1389 or $e_2 \xrightarrow{\text{vis}} e_1$, it must be the case that $e_1 \not\Rightarrow e_2$. In this case, we can flip the order of e_1 and e_2 in π_2
 1390 leading to the same state. Suppose $e_1 \parallel_C e_2$, if neither $e_1 \xrightarrow{\text{rc}} e_2$ nor $e_2 \xrightarrow{\text{rc}} e_1$, $e_1 \not\Rightarrow e_2$. In this case,
 1391 we can again flip them in π_2 . Suppose $e_1 \xrightarrow{\text{rc}} e_2$, since $\neg(e_1 \xrightarrow{\text{lo}} e_2)$, by definition of lo , $\exists e_3.e_2 \xrightarrow{\text{lo}} e_3$.
 1392 Then $\neg(e_2 \not\Rightarrow e_3)$. By COND-COMM , it must be the case that $e_1 \xrightarrow{e_3} e_2$. Since $e_2 \xrightarrow{\text{lo}} e_3$, e_3 must be
 1393 present in τ_4 . By definition of COND-COMM , we can flip e_2 and e_1 in π_2 , leading to the same state.
 1394 Similar argument can be applied to $e_2 \xrightarrow{\text{rc}} e_1$.
 1395

1396 **Case 2:** ($\tau_3 \neq \phi$)

1397 $\pi_2 = \tau.e_2.\tau_5.e_3.e_1.\tau_4$. Then in π_1 , e_3 is not present in τ , hence it must be present after e_1 . Now e_1 and
 1398 e_2 are in different orders in π_1 and π_2 , hence neither $e_1 \xrightarrow{\text{lo}} e_3$ nor $e_3 \xrightarrow{\text{lo}} e_1$.
 1399

1400 By the same argument as above applied on e_1 and e_2 , we can flip e_1 and e_3 in π_2 . We keep doing
 1401 this for all events in τ_5 until e_2 is adjacent to e_1 after which we can flip them. Thus we can change
 1402 π_2 such that e_1 will appear in the same position in π_1 . We can keep doing this until π_1 and π_2 are
 1403 identical. \square
 1404

1405 **Lemma 3.10** If MRDT \mathcal{D} is RA-linearizable, then for all executions $\tau \in \llbracket \mathcal{S}_{\mathcal{D}} \rrbracket$, and for all transitions
 1406 $C \xrightarrow{\text{query}(r,q,a)} C'$ in τ , where $C = \langle N, H, L, G, \text{vis} \rangle$, there exists a sequence π consisting of all events
 1407 in $L(H(r))$ such that $\text{lo}(C)|_{L(H(r))} \subseteq \pi$ and $a = \text{query}(\pi(\sigma_0), q)$.
 1408

1409 **PROOF.** Consider an MRDT \mathcal{D} that is RA-linearizable. Let $\tau = C_0 \xrightarrow{t_1} C_1 \xrightarrow{t_2} C_2 \dots \xrightarrow{t_n} C$ be
 1410 an execution of $\mathcal{S}_{\mathcal{D}}$, where $\{t_1, \dots, t_n\}$ are the labels of the transition system. For a transition
 1411 $C \xrightarrow{\text{query}(r,q,a)} C'$ in τ , where $C = \langle N, H, L, G, \text{vis} \rangle$, we know that C is RA-linearizable from Def. 3.9.
 1412 That is, for every active replica $r \in \text{range}(H)$, there exists a sequence π such that $\text{lo}(C)|_{L(H(r))} \subseteq \pi$
 1413 and $N(H(r)) = \pi(\sigma_0)$. According to the semantics, we have $a = \text{query}(N(H(r)), q)$. Thus $a =$
 1414 $\text{query}(\pi(\sigma_0), q)$. \square
 1415

1416 A.2 Proofs of §4

1417 **Lemma 4.2** (1) For events $e \in L_1^a \cup L_2^a, e' \in L_1^b \cup L_2^b, \neg(e \xrightarrow{\text{lo}_m} e')$.
 1418

1419 **PROOF.** Suppose $e \xrightarrow{\text{lo}_m} e'$ is true. There are 2 possibilities:
 1420
 1421

- 1422 (1) $e \xrightarrow[\text{vis}]{\text{lo}} e'$: By definition of L_i^b , there are 2 cases:
 1423
 1424 (a) $\exists e_\top \in L_\top, e' \xrightarrow{\text{lo}_m} e_\top$: But this would require e to be in $L_1^b \cup L_2^b$.
 1425 (b) $\exists e_\top \in L_\top, e'' \in L'_1 \cup L'_2, e' \xrightarrow{\text{lo}_m} e'' \xrightarrow{\text{lo}_m} e_\top$:
 1426 (i) $e' \xrightarrow[\text{vis}]{\text{lo}} e''$: Due to transitivity of vis, $e \xrightarrow{\text{vis}} e''$. This would require $e \in L_1^b \cup L_2^b$.
 1427 (ii) $e'' \xrightarrow[\text{vis}]{\text{lo}} e_\top$ is not possible as L_\top^a is causally closed.
 1428 (iii) $e' \xrightarrow[\text{rc}]{\text{lo}} e'' \xrightarrow[\text{rc}]{\text{lo}} e_\top$ is not possible due to NO-RC-CHAIN restriction.
 1429
 1430 (2) $e \xrightarrow[\text{rc}]{\text{lo}} e'$: By definition of L_i^b , there are 2 cases:
 1431
 1432 (a) $\exists e_\top \in L_\top, e' \xrightarrow{\text{lo}_m} e_\top$:
 1433 (i) $e' \xrightarrow[\text{vis}]{\text{lo}} e_\top$ is not possible as L_\top^a is causally closed.
 1434 (ii) $e' \xrightarrow[\text{rc}]{\text{lo}} e_\top$ is not possible due to NO-RC-CHAIN restriction. Since $e \parallel_C e_\top$, we have
 1435 $e \xrightarrow[\text{rc}]{\text{lo}} e_\top$ which requires $e \in L_1^b \cup L_2^b$.
 1436 (b) $\exists e_\top \in L_\top, e'' \in L'_1 \cup L'_2, e' \xrightarrow{\text{lo}_m} e'' \xrightarrow{\text{lo}_m} e_\top$:
 1437 (i) $e'' \xrightarrow[\text{vis}]{\text{lo}} e_\top$ is not possible as L_\top^a is causally closed.
 1438 (ii) $e' \xrightarrow[\text{rc}]{\text{lo}} e''$ is not possible due to NO-RC-CHAIN restriction.
 1439 (iii) $e' \xrightarrow[\text{vis}]{\text{lo}} e'' \xrightarrow[\text{rc}]{\text{lo}} e_\top$: $e' \xrightarrow{\text{rc}} e''$ creates RC-chain. Since $e' \parallel_C e_\top$, we have $e' \xrightarrow[\text{rc}]{\text{lo}} e_\top$
 1440 which violates the NO-RC-CHAIN restriction. $e'' \xrightarrow{\text{rc}} e'$ would require e and e' to
 1441 conditionally commute with each other. So $e \xrightarrow[\text{rc}]{\text{lo}} e'$ does not hold true.
 1442
 1443
 1444
 1445
 1446
 1447
 1448
 1449

□

1450
 1451
 1452 (2) For events $e \in L_\top^a, e' \in L_\top^b, \neg(e \xrightarrow{\text{lo}_m} e')$.

1453
 1454 PROOF. By definition of $L_\top^a, \exists e'' \in L_1^b \cup L_2^b, e'' \xrightarrow{\text{lo}_m} e, e'' \xrightarrow{\text{vis}} e$ is not possible as L_\top^a is causally
 1455 closed. Suppose $e \xrightarrow{\text{lo}_m} e'$ is true. There are 3 possibilities:
 1456

- 1457 (1) $e \xrightarrow[\text{vis}]{\text{lo}} e'$:
 1458 (a) $e'' \xrightarrow[\text{rc}]{\text{lo}} e : e \xrightarrow{\text{rc}} e'$ causes RC-chain. Since $e'' \parallel_C e'$, we have $e'' \xrightarrow[\text{rc}]{\text{lo}} e'$ which requires
 1459 $e' \in L_\top^a, e' \xrightarrow{\text{rc}} e$ cause e'' and e to conditionally commute with each other. So this case
 1460 does not hold true.
 1461 (2) $e \xrightarrow[\text{rc}]{\text{lo}} e' : e'' \xrightarrow{\text{vis}} e$ is not possible as L_\top^a is causally closed.
 1462 (a) $e'' \xrightarrow[\text{rc}]{\text{lo}} e$: causes RC-chain.
 1463
 1464
 1465
 1466

□

1467
 1468 **Lemma 4.4** (1) For events $e_i^\top, e_j^\top \in L_\top^a$, where $L_\top^a = \{e_1^\top, \dots, e_m^\top\}, \neg(e_i^\top \xrightarrow{\text{lo}_m} e_j^\top)$.
 1469
 1470

PROOF. By definition of L_\top^a , $\exists e \in L_1^b(e_i^\top) \cup L_2^b(e_i^\top). e \xrightarrow{\text{lo}_m} e_i^\top. e \xrightarrow{\text{vis}} e_i^\top$ is not possible as L_\top^a is causally closed. Suppose $e_i^\top \xrightarrow{\text{lo}_m} e_j^\top$. There are 3 possibilities.

$$(1) e_i^\top \xrightarrow{\text{lo}_{\text{vis}}} e_j^\top :$$

(a) $e \xrightarrow{\text{lo}_{\text{rc}}} e_i^\top : e_i^\top \xrightarrow{\text{rc}} e_j^\top$ causes RC-chain. Since $e \parallel_C e_j^\top$, we have $e \xrightarrow{\text{lo}_{\text{rc}}} e_j^\top$ which requires $e \in L_1^b(e_j^\top) \cup L_2^b(e_j^\top)$. But e belongs to $L_1^b(e_i^\top) \cup L_2^b(e_i^\top)$. $e_j^\top \xrightarrow{\text{rc}} e_i^\top$ cause e and e_i^\top to conditionally commute with each other. So this case does not hold true.

$$(2) e_i^\top \xrightarrow{\text{lo}_{\text{rc}}} e_j^\top :$$

(a) $e \xrightarrow{\text{lo}_{\text{rc}}} e_i^\top$: By NO-RC-CHAIN restriction, this case cannot happen.

□

(2) For events $e \in L_1^b(e_i^\top) \cup L_2^b(e_i^\top)$, $e' \in L_1^b(e_j^\top) \cup L_2^b(e_j^\top)$ where $j < i$, $\neg(e \xrightarrow{\text{lo}_m} e')$.

PROOF. Suppose $e \xrightarrow{\text{lo}_m} e'$, $\neg(e \rightleftarrows e')$. By definition of $L_1^b(e_i^\top)$ and $L_2^b(e_i^\top)$, we know that $e \xrightarrow{\text{lo}} e_i^\top$ and $e' \xrightarrow{\text{lo}} e_j^\top$. We consider several possibilities based on this:

(1) Neither $e \xrightarrow{\text{lo}_{\text{vis}}} e_i^\top$ nor $e' \xrightarrow{\text{lo}_{\text{vis}}} e_j^\top$ is true because L_\top^a is causally closed.

(2) $e \xrightarrow{\text{lo}_{\text{rc}}} e_i^\top \wedge e' \xrightarrow{\text{lo}_{\text{rc}}} e_j^\top$:

(a) $e \xrightarrow{\text{rc}} e' \vee e' \xrightarrow{\text{rc}} e$ creates RC chain.

□

Theorem 4.6 If an MRDT \mathcal{D} satisfies BOTTOMUP-2-OP, BOTTOMUP-1-OP, BOTTOMUP-0-OP, MERGEIDEMPOTENCE and MERGECOMMUTATIVITY, then \mathcal{D} is linearizable.

PROOF. To prove that \mathcal{D} is linearizable, we will prove that any execution $\tau \in \llbracket \mathcal{S}_{\mathcal{D}} \rrbracket$ is linearizable, for which we will show that all of its configurations are linearizable. Let $\tau = C_0 \xrightarrow{t_1} C_1 \xrightarrow{t_2} C_2 \dots \xrightarrow{t_n} C$ be an execution of $\mathcal{S}_{\mathcal{D}}$, where $\{t_1, \dots, t_n\}$ are labels of the transition system. We prove by induction on the length of τ . Base case of C_0 which consists of only one replica r_0 is trivially satisfied, as no operations are applied on the head version v_0 at r_0 . Assuming the required result holds in the execution $C_0 \rightarrow^* C$, and suppose there is a new transition $C \rightarrow C'$, we need to prove that C is linearizable. There are four cases corresponding to the four transition rules given in Fig. 8.

A.2.1 Case (CREATEBRANCH): Assume that a new replica r' is forked off from the origin replica r . Let $C = \langle N, H, L, G, \text{vis} \rangle$ and $C' = \langle N', H', L', G', \text{vis} \rangle$ be the configurations of the replica before and after the branch creation. According to the semantics, we have $L(H(r)) = L'(H'(r'))$ and $N(H(r)) = N'(H'(r'))$. We need to prove that Def. 3.9 holds for C' . This is obvious since Def. 3.9 holds for C by the induction assumption.

A.2.2 Case (APPLY): Assume that an event e is applied on a replica r . Let $C = \langle N, H, L, G, \text{vis} \rangle$ and $C' = \langle N', H', L', G', \text{vis}' \rangle$ be the configurations of the replica before and after the apply operation. By semantics we have $L'(H'(r)) = L(H(r)) \cup \{e\}$. We need to prove that Def. 3.9 holds for C' . By induction assumption, $\exists \pi. \text{lo}(C)|_{L(H(r))} \subseteq \pi \wedge N(H(r)) = \pi(\sigma_0)$. Here $\text{lo}(C')|_{L'(H'(r))}$ is the linearization order $\text{lo}(C)|_{L(H(r))}.e$ and $\pi' = \pi.e$. We need to show that π' extends $\text{lo}(C')|_{L'(H'(r))}$. We have $N'(H'(r)) = e(\pi(\sigma_0))$. Event e is visible to all events in π according to the semantics of

1520 apply. Since $\forall e' \in \pi. e' \xrightarrow{\text{vis}} e, e \xrightarrow{\text{lo}} e'$ is not possible due to anti-symmetry of vis. $e \xrightarrow{\text{lo}} e'$ is also
 1521 not possible as it would require e and e' to be concurrent events. Hence, π' is a total order which
 1522 extends $\text{lo}(C')|_{L'(H'(r))}$. This proves the required result.
 1523

1524 **A.2.3 Case (MERGE):** Consider there is a merge(r_1, r_2) transition to C' where r_2 merges with r_1 .
 1525 Let $C = \langle N, H, L, G, \text{vis} \rangle$, $C' = \langle N', H', L', G', \text{vis} \rangle$, and let $H(r_1) = v_1, H(r_2) = v_2$. Let v_\top be the
 1526 LCA of v_1 and v_2 in G . Let $N(v_1) = a, N(v_2) = b, N(v_\top) = l$. The transition will install a version
 1527 v_m with state $m = \text{merge}(l, a, b)$ at the replica r_1 , leaving the other replicas unchanged. Also,
 1528 $L'(v_m) = L(v_1) \cup L(v_2)$. We need to show that there exists a sequence π_m of events in $L'(v_m)$ such
 1529 that π_m extends $\text{lo}(C')|_{L'(v_m)}$ and $m = \pi(\sigma_0)$. For ease of readability, we use L_1 for $L(v_1)$, L_2 for
 1530 $L(v_2)$ and L_\top for $L(v_\top)$, and lo_m for $\text{lo}(C')|_{L'(v_m)}$.
 1531

1532 We repeat the definitions of various event sets below:

$$1533 \quad L'_1 = L_1 \setminus L_\top \quad L'_2 = L_2 \setminus L_\top$$

$$1534 \quad L_1^b = \{e \in L_1 \mid \exists e_\top \in L_\top. (e \xrightarrow{\text{lo}_m} e_\top \vee \exists e' \in L'_1. e \xrightarrow{\text{lo}_m} e' \xrightarrow{\text{lo}_m} e_\top)\}$$

$$1535 \quad L_2^b = \{e \in L_2 \mid \exists e_\top \in L_\top. (e \xrightarrow{\text{lo}_m} e_\top \vee \exists e' \in L'_2. e \xrightarrow{\text{lo}_m} e' \xrightarrow{\text{lo}_m} e_\top)\}$$

$$1536 \quad L_\top^a = \{e_\top \in L_\top \mid \exists e \in L_1^b \cup L_2^b. e \xrightarrow{\text{lo}_m} e_\top\}$$

$$1537 \quad L_1^a = L'_1 \setminus L_1^b \quad L_2^a = L'_2 \setminus L_2^b \quad L_\top^b = L_\top \setminus L_\top^a$$

1538 Let $\text{lo}_i = \text{lo}(C')|_{L_i}$ for $i = 1, 2$.

1539 First we will prove that lo between two events should remain the same in all versions. $\forall e, e' \in$
 1540 $L_i. e \xrightarrow{\text{lo}_i} e' \Leftrightarrow e \xrightarrow{\text{lo}_m} e'$. Note that vis and rc ordering between events remains same in L_i and
 1541 $L'(v_m)$.
 1542

- 1543 • If $e \xrightarrow{\text{rc}} e', e \parallel_C e'$ and $\neg(\exists e'' \in L(v_i). e' \xrightarrow{\text{vis}} e'' \wedge \neg e' \not\Rightarrow e'')$, then these constraints will
 1544 continue to hold in L_m . Because it is not possible that $e' \in L'_1, e'' \in L'_2$ such that $e' \xrightarrow{\text{vis}} e''$.
 1545 Because otherwise $e' \in L'_2 \Rightarrow e' \in L_\top$.
- 1546 • If $e \xrightarrow{\text{vis}} e' \wedge \neg e \not\Rightarrow e'$ in L_i , then it continues to hold in L_m .
 1547

1548 By induction assumption, we know that

$$1549 \quad \exists \pi_a. \text{lo}(C)|_{L(v_1)} \subseteq \pi_a \wedge a = \pi_a(\sigma_0)$$

$$1550 \quad \exists \pi_b. \text{lo}(C)|_{L(v_2)} \subseteq \pi_b \wedge b = \pi_b(\sigma_0)$$

$$1551 \quad \exists \pi_\top. \text{lo}(C)|_{L(v_\top)} \subseteq \pi_\top \wedge l = \pi_\top(\sigma_0)$$

1552 To start off, let's consider the set $L_1^a \cup L_2^a$. These are all local events of v_1 and v_2 , which are not
 1553 linearized before events of the LCA. We consider different cases depending on the size of this set.
 1554

1555 **CASE 1:** ($|L_1^a \cup L_2^a| = 0$)

1556 We note that in this case, a, b can be defined as follows: $a = \pi_a|_{(L_\top^b \cup L_1^b \cup L_\top^a)}(\sigma_0)$, $b = \pi_b|_{(L_\top^b \cup L_2^b \cup L_\top^a)}(\sigma_0)$.
 1557 We need to show that there exists a sequence π_m that extends lo_m such that $\text{merge}(l, a, b) = \pi_m(\sigma_0)$.
 1558 Here, we induct on the size of the set L_\top^a .
 1559

1560 **BASE CASE 1:** ($|L_\top^a| = 0$)

1561 Then $L_1^b \cup L_2^b = \phi$. So $l = a = b$. $\text{merge}(l, l, l) = l$ is inferred by MERGEIDEMPOTENCE. We know that
 1562 l is correctly linearized, hence the required result follows.
 1563
 1564
 1565
 1566
 1567
 1568

1569 INDUCTIVE CASE 1: ($|L_{\top}^a| > 0$)

1570 Let $L_{\top}^a = \{e_1^{\top}, \dots, e_{m-1}^{\top}, e_m^{\top}\}$. Let $S = \{e_1^{\top}, \dots, e_{m-1}^{\top}\}$. By IH, for the set S , we have the required
 1571 result. We define l', a', b' based on the above set S : $l' = \pi_{l|L_{\top}^b \cup S}(\sigma_0)$, $a' = \pi_{a|L_{\top}^b \cup \bigcup_{e \in S} L_1^b(e) \cup S}(\sigma_0)$,

1572 $b' = \pi_{b|L_{\top}^b \cup \bigcup_{e \in S} L_2^b(e) \cup S}(\sigma_0)$. Note that in this case, all the LCA events which are linearized after local
 1573 events are already taken as part of the states l', a', b' . Now, suppose we add one more LCA event
 1574 e_m^{\top} to all states. We define a'', b'' such that $a'' = \pi_{a|L_1^b(e_m^{\top})}(a')$, $b'' = \pi_{b|L_2^b(e_m^{\top})}(b')$.

1575 Then, $l = e_m^{\top}(l')$, $a = e_m^{\top}(a'')$, $b = e_m^{\top}(b'')$. e_m^{\top} is not linearized before any of the events in
 1576 $L_{\top}^b \cup L_1^b \cup L_2^b \cup S$ based on the definition of L_{\top}^a .

1577 Now, by BOTTOMUP-0-OP rule,

$$1578 \text{merge}(e_m^{\top}(l'), e_m^{\top}(a''), e_m^{\top}(b'')) = e_m^{\top}(\text{merge}(l', a'', b'')) \quad (3)$$

1580 Now that we have linearized e_m^{\top} , we need to linearize the events that led to $\text{merge}(l', a'', b'')$.
 1581 Let's denote $L_1^b(e_m^{\top})$ as M_1^a and $L_2^b(e_m^{\top})$ as M_2^a . Now we induct on the size of the set $M_1^a \cup M_2^a$.

1584 BASE CASE 1.1: ($|M_1^a \cup M_2^a| = 0$)

1585 $a'' = a', b'' = b'$. By induction assumption, $\exists \pi. \text{lo}(C)_{|(L_{\top}^b \cup \bigcup_{e \in S} L_1^b(e) \cup \bigcup_{e \in S} L_2^b(e) \cup S)} \subseteq \pi$
 1586 and $\text{merge}(l', a', b') = \pi(\sigma_0)$. Hence, $\pi_m = \pi.e_m^{\top}$.

1588 INDUCTIVE CASE 1.1: ($|M_1^a \cup M_2^a| > 0$)

1589 We have 2 cases here:

1590 (1.1.1) Either of M_1^a or M_2^a is ϕ

1591 (1.1.2) Both M_1^a and M_2^a are not ϕ .

1593 CASE 1.1.1: ($M_1^a \neq \phi \wedge M_2^a = \phi$)

1594 Consider $e_1 \in M_1^a$ such that there does not exist $e \in M_1^a$ and $e_1 \xrightarrow{\text{lo}_m} e$, i.e. e_1 is the maximal event
 1595 according to lo_m . Since lo ordering between events remains the same in all versions, and since
 1596 versions v_1 and v_2 (which are being merged) were already linearizable, there would exist sequences
 1597 leading to the states a such that e_1 would appear at the end. Hence, there exists a''' such that
 1598 $a'' = e_1(a''')$. Since M_2^a is empty, all local events in L_2 are linearized before the rest of the LCA
 1599 events. Suppose $L_{\top}^a \setminus e_m^{\top} \neq \phi$ or $L_{\top}^b \neq \phi$, the last event which leads to the state l', b'' must be an
 1600 LCA event. Let's consider e_{\top} to be the maximal event in L_{\top} according to lo_m . Hence there exists
 1601 states l'', b''' such that $l' = e_{\top}(l'')$, $b'' = e_{\top}(b''')$. By BOTTOMUP-1-OP rule

$$1602 \text{merge}(e_{\top}(l''), e_1(a'''), e_{\top}(b''')) = e_1(\text{merge}(e_{\top}(l''), a''', e_{\top}(b'''))) \quad (4)$$

1604 If both $L_{\top}^a \setminus e_m^{\top} = \phi$ and $L_{\top}^b = \phi$, then $l' = b'' = \sigma_0$. By BOTTOMUP-1-OP

$$1605 \text{merge}(\sigma_0, e_1(a'''), \sigma_0) = e_1(\text{merge}(\sigma_0, a''', \sigma_0))$$

1606 From the induction assumption, we get that $\text{merge}(e_{\top}(l''), a''', e_{\top}(b'''))$ is already obtained
 1607 by the linearization of events applied on the initial state σ_0 . That is, there exists a sequence π'
 1608 over events in $L_{\top}^b \cup \bigcup_{e \in S} L_1^b(e) \cup \bigcup_{e \in S} L_2^b(e) \cup S \cup M_1^a \setminus e_1$ which extends lo_m relation such that
 1609 $\text{merge}(e_{\top}(l''), a''', e_{\top}(b''')) = \pi'(\sigma_0)$. Now, $\pi = \pi'.e_1$ is the required linearization.

1610 Let lo_1 be the linearization relation for $\text{merge}(e_{\top}(l''), a''', e_{\top}(b'''))$ (i.e. from the RHS in Eq. (4),
 1611 without the event e_1) and let lo_2 be the linearization relation for $\text{merge}(l', a'', b'')$ (i.e. the LHS in
 1612 Eq. (4)). Then π' according to the IH extends lo_1 . We will show that for any pair of events e, e'
 1613 in $\text{merge}(e_{\top}(l''), a''', e_{\top}(b'''))$, $e \xrightarrow{\text{lo}_2} e' \implies e \xrightarrow{\text{lo}_1} e'$. This ensures that if π extends lo_2 . Now,
 1614 the vis and rc relation between e' and e remains the same while determining both lo_1 and lo_2 . If

1618 $e \xrightarrow[\text{rc}]{\text{lo}_2} e'$, then e' cannot be visible to any non-commutative event while calculating lo_2 , but then
 1619
 1620 the same should be true for lo_1 as well. If $e \xrightarrow[\text{vis}]{\text{lo}_2} e'$, then clearly $e \xrightarrow[\text{vis}]{\text{lo}_1} e'$. This concludes the proof
 1621 that $\pi = \pi'.e_1$ must extend lo_2 .
 1622

1623 CASE 1.1.2: ($M_1^a \neq \phi \wedge M_2^a \neq \phi$)

1624 Consider $e_1 \in M_1^a, e_2 \in M_2^a$ such that there does not exist $e \in M_1^a$ and $e_i \xrightarrow{\text{lo}_m} e$ (for $i = 1, 2$), i.e.
 1625 each of the e_i s are maximal events according to lo_m . Since lo ordering between events remains
 1626 the same in all versions, and since versions v_1 and v_2 (which are being merged) were already
 1627 linearizable, there would exist sequences leading to the states a'' and b'' such that e_1 and e_2 would
 1628 appear at the end resp. Hence, there exists a''' and b''' such that $a'' = e_1(a''')$ and $b'' = e_1(b''')$.
 1629 Since $e_1 \parallel_C e_2$, they are related to each other by rc relation or they commute with each other i.e.,
 1630 $e_1 \xrightarrow{\text{rc}} e_2 \vee e_2 \xrightarrow{\text{rc}} e_1 \vee e_1 \rightleftharpoons e_2$. We will consider the case when $e_2 \xrightarrow{\text{rc}} e_1 \vee e_1 \rightleftharpoons e_2$. $e_1 \xrightarrow{\text{rc}} e_2$ is
 1631 handled by MERGECOMMUTATIVITY. The equation becomes
 1632

$$1633 \quad \text{merge}(l', e_1(a'''), e_2(b''')) = e_1(\text{merge}(l', a''', e_2(b'''))) \quad (5)$$

1634 which is the BOTTOMUP-2-OP rule.
 1635
 1636

1637 From the induction assumption, we get that $\text{merge}(l', a''', e_2(b'''))$ is already obtained by the
 1638 linearization of events applied on the initial state σ_0 . If π' is the linearization for this merge, then
 1639 $\pi = \pi'.e_1$ is the required linearization.

1640 For this, we prove that e_1 is not linearized before any of the events in $M_1^a \setminus \{e_1\} \cup M_2^a$. Clearly,
 1641 e_1 is not linearized before any event in $M_1^a \setminus \{e_1\}$ because it is the maximal event on that branch.
 1642 Since $e_2 \xrightarrow{\text{rc}} e_1, e_1 \xrightarrow{\text{vis}} e_2$ is not possible. $e_1 \xrightarrow{\text{rc}} e_2$ is not possible as rc^+ is irreflexive. So $e_1 \xrightarrow{\text{lo}} e_2$
 1643 is not possible. Let's assume there is some event e in $M_2^a \setminus \{e_2\}$ that comes lo after e_1 . There are 2
 1644 possibilities.
 1645

- 1646 • $e_1 \xrightarrow{\text{rc}} e$: Since $e_2 \xrightarrow{\text{rc}} e_1$, this case is not possible due to NO-RC-CHAIN restriction.
- 1647 • $e_1 \xrightarrow{\text{vis}} e$: This is not possible as events in $M_2^a \setminus \{e_2\}$ are concurrent with e_1 . This is because
 1648 every version is causally closed.

1649 CASE 2: ($|L_1^a \cup L_2^a| > 0$)

1650 The proof here will be identical to the proof of Inductive Case 1.1, substituting L_1^a and L_2^a for M_1^a
 1651 and M_2^a , and using the rules BOTTOMUP-1-OP, MERGECOMMUTATIVITY and BOTTOMUP-2-OP.
 1652

1653 A.2.4 Case (QUERY): Assume that a query operation is applied on a replica r . Let $C = \langle N, H, L, G, \text{vis} \rangle$
 1654 be the configuration of the replica before the operation. According to the semantics, the configura-
 1655 tion of the replica remains same after the query operation. By the induction hypothesis, Def. 3.9
 1656 holds for the configuration C . \square
 1657

1658 **Theorem 4.7** If an MRDT \mathcal{D} satisfies the VCs $\psi^*(\text{BOTTOMUP-2-OP}), \psi^*(\text{BOTTOMUP-1-OP}),$
 1659 $\psi^*(\text{BOTTOMUP-0-OP}), \text{MERGEIDEMPOTENCE}$ and $\text{MERGECOMMUTATIVITY}$, then \mathcal{D} is linearizable.
 1660

1661 **PROOF.** To prove that \mathcal{D} is linearizable, we will prove that any execution $\tau \in \llbracket \mathcal{S}_{\mathcal{D}} \rrbracket$ is linearizable,
 1662 for which we will show that all of its configurations are linearizable. Let $\tau = C_0 \xrightarrow{t_1} C_1 \xrightarrow{t_2} C_2 \dots \xrightarrow{t_n} C$
 1663 be an execution of $\mathcal{S}_{\mathcal{D}}$, where $\{t_1, \dots, t_n\}$ are labels of the transition system. We prove by
 1664 induction on the length of τ . Base case of C_0 which consists of only one replica r_0 is trivially
 1665 satisfied, as no operations are applied on the head version v_0 at r_0 . Assuming the required result
 1666

holds in the execution $C_0 \rightarrow^* C$, and suppose there is a new transition $C \rightarrow C'$, we need to prove that C is linearizable. There are four cases corresponding to the four transition rules given in Fig. 8.

A.2.5 Case (CREATEBRANCH): Assume that a new replica r' is forked off from the origin replica r . Let $C = \langle N, H, L, G, vis \rangle$ and $C' = \langle N', H', L', G', vis \rangle$ be the configurations of the replica before and after the branch creation. According to the semantics, we have $L(H(r)) = L'(H'(r'))$ and $N(H(r)) = N'(H'(r'))$. We need to prove that Def. 3.9 holds for C' . This is obvious since Def. 3.9 holds for C by the induction assumption.

A.2.6 Case (APPLY): Assume that an event e is applied on a replica r . Let $C = \langle N, H, L, G, vis \rangle$ and $C' = \langle N', H', L', G', vis' \rangle$ be the configurations of the replica before and after the apply operation. By semantics we have $L'(H'(r)) = L(H(r)) \cup \{e\}$. We need to prove that Def. 3.9 holds for C' . By induction assumption, $\exists \pi. \text{lo}(C)_{|L(H(r))} \subseteq \pi \wedge N(H(r)) = \pi(\sigma_0)$. Here $\text{lo}(C')_{|L'(H'(r))}$ is the linearization order $\text{lo}(C)_{|L(H(r))}.e$ and $\pi' = \pi.e$. We need to show that π' extends $\text{lo}(C')_{|L'(H'(r))}$. We have $N'(H'(r)) = e(\pi(\sigma_0))$. Event e is visible to all events in π according to the semantics of apply. Since $\forall e' \in \pi. e' \xrightarrow[\text{vis}]{\text{lo}} e, e \xrightarrow[\text{vis}]{\text{lo}} e'$ is not possible due to anti-symmetry of vis. $e \xrightarrow[\text{rc}]{\text{lo}} e'$ is also not possible as it would require e and e' to be concurrent events. Hence, π' is a total order which extends $\text{lo}(C')_{|L'(H'(r))}$. This proves the required result.

A.2.7 Case (MERGE): Consider there is a merge(r_1, r_2) transition to C' where r_2 merges with r_1 . Let $C = \langle N, H, L, G, vis \rangle, C' = \langle N', H', L', G', vis \rangle$, and let $H(r_1) = v_1, H(r_2) = v_2$. Let v_\top be the LCA of v_1 and v_2 in G . Let $N(v_1) = a, N(v_2) = b, N(v_\top) = l$. The transition will install a version v_m with state $m = \text{merge}(l, a, b)$ at the replica r_1 , leaving the other replicas unchanged. Also, $L'(v_m) = L(v_1) \cup L(v_2)$. We need to show that there exists a sequence π_m of events in $L'(v_m)$ such that π_m extends $\text{lo}(C')_{|L'(v_m)}$ and $m = \pi(\sigma_0)$. For ease of readability, we use L_1 for $L(v_1), L_2$ for $L(v_2)$ and L_\top for $L(v_\top)$, and lo_m for $\text{lo}(C')_{|L'(v_m)}$.

By induction assumption, we know that

$$\begin{aligned} \exists \pi_a. \text{lo}(C)_{|L(v_1)} &\subseteq \pi_a \wedge a = \pi_a(\sigma_0) \\ \exists \pi_b. \text{lo}(C)_{|L(v_2)} &\subseteq \pi_b \wedge b = \pi_b(\sigma_0) \\ \exists \pi_\top. \text{lo}(C)_{|L(v_\top)} &\subseteq \pi_\top \wedge l = \pi_\top(\sigma_0) \end{aligned}$$

To start off, let's consider the set $L_1^a \cup L_2^a$. These are all local events of v_1 and v_2 , which are not linearized before events of the LCA. We consider different cases depending on the size of this set.

CASE 1: ($|L_1^a \cup L_2^a| = 0$)

We note that in this case, a, b can be defined as follows: $a = \pi_{a|(L_1^b \cup L_1^b \cup L_\top^a)}(\sigma_0), b = \pi_{b|(L_1^b \cup L_2^b \cup L_\top^a)}(\sigma_0)$. We need to show that there exists a sequence π_m that extends lo_m such that $\text{merge}(l, a, b) = \pi_m(\sigma_0)$. Here, we induct on the size of the set L_\top^a .

BASE CASE 1: ($L_\top^a = \phi$)

Then $L_1^b \cup L_2^b = \phi$. So $l = a = b$. $\text{merge}(l, l, l) = l$ is handled by MERGEIDEMPOTENCE. We know that l is correctly linearized, hence the required result follows.

INDUCTIVE CASE 1: ($|L_\top^a| > 0$)

Let $L_\top^a = \{e_1^\top, \dots, e_{m-1}^\top, e_m^\top\}$. Let $S = \{e_1^\top, \dots, e_{m-1}^\top\}$. By IH, for the set S , we have the required result. We define l', a', b' based on the above set S : $l' = \pi_{l|L_\top^b \cup S}(\sigma_0), a' = \pi_{a|L_1^b \cup \cup_{e \in S} L_1^b(e) \cup S}(\sigma_0), b' = \pi_{b|L_\top^b \cup \cup_{e \in S} L_2^b(e) \cup S}(\sigma_0)$. Note that in this case, all the LCA events which are linearized after local

events are already taken as part of the states l', a', b' . Now, suppose we add one more LCA event e_m^\top to all states. We define a'', b'' such that $a'' = \pi_{a|L_1^b(e_m^\top)}(a')$, $b'' = \pi_{b|L_2^b(e_m^\top)}(b')$.

Then, $l = e_m^\top(l')$, $a = e_m^\top(a'')$, $b = e_m^\top(b'')$. e_m^\top is not linearized before any of the events in $L_\top^b \cup L_1^b \cup L_2^b \cup S$ based on the definition of L_\top^a .

Now, by BOTTOMUP-0-OP rule,

$$\text{merge}(e_m^\top(l'), e_m^\top(a''), e_m^\top(b'')) = e_m^\top(\text{merge}(l', a'', b'')) \quad (6)$$

We will now show prove BOTTOMUP-0-OP rule, i.e. Eqn. (6):

PROOF OF EQ. (6):

Let $l_b = \pi_{l|L_\top^b}(\sigma_0)$.

We first induct on $|L_\top^b|$ to show that $\text{merge}(e_m^\top(l_b), e_m^\top(l_b), e_m^\top(l_b)) = e_m^\top(\text{merge}(l_b, l_b, l_b))$

For the base case, we use $\psi_{\text{base-0op}}^{L_\top^b}$. For the inductive case, we use $\psi_{\text{ind-0op}}^{L_\top^b}$, whose pre-condition will be satisfied by the IH.

Next, we induct on $|L_\top^a \setminus \{e_m^\top\}|$ to show Eqn. (6).

For the base case, we have $|L_\top^a \setminus \{e_m^\top\}| = 0$. In this case, the set $S = \emptyset$. Also, $l' = a' = b' = l_b$. Hence, we need to show the following:

$$\text{merge}(e_m^\top(l_b), e_m^\top(\pi_{a|L_1^b(e_m^\top)}(a')), e_m^\top(\pi_{b|L_2^b(e_m^\top)}(b'))) = e_m^\top(\text{merge}(l', \pi_{a|L_1^b(e_m^\top)}(a'), \pi_{b|L_2^b(e_m^\top)}(b'))) \quad (7)$$

We will now induct on $|L_1^b(e_m^\top) \cup L_2^b(e_m^\top)|$ to show Eqn. (7).

For the base case where $|L_1^b(e_m^\top) \cup L_2^b(e_m^\top)| = 0$, it directly follows from the outcome of the induction on $|L_\top^b|$.

For the inductive case, we use one of $\psi_{\text{ind1-0op}}^{L_1^b}$, $\psi_{\text{ind2-0op}}^{L_1^b}$, $\psi_{\text{ind1-0op}}^{L_2^b}$ or $\psi_{\text{ind2-0op}}^{L_2^b}$ depending on the event e_b or e to be added to $L_1^b(e_m^\top)$ or $L_2^b(e_m^\top)$, with the pre-condition of these VCs being inferred from the IH.

This completes the proof of Eqn. (7).

Now, we consider the inductive case for $|L_\top^a \setminus \{e_m^\top\}|$ to show Eqn. (??). By IH, we get the following:

$$\text{merge}(e_m^\top(l'''), e_m^\top(a'''), e_m^\top(b''')) = e_m^\top(\text{merge}(l''', a''', b''')) \quad (8)$$

where for the set $S' = S \setminus e_{m-1}^\top$, $l''' = \pi_{l|L_\top^b \cup S'}(\sigma_0)$, $a''' = \pi_{a|L_\top^b \cup \cup_{e \in S'} L_1^b(e) \cup S'}(\sigma_0)$, $b''' = \pi_{b|L_\top^b \cup \cup_{e \in S'} L_2^b(e) \cup S'}(\sigma_0)$.

That is, we consider the effects of all event in S except e_{m-1}^\top .

Now, we first use $\psi_{\text{ind-0op}}^{L_\top^a}$ to apply e_{m-1}^\top to l''', a''' and b''' . Note that the pre-condition for $\psi_{\text{ind-0op}}^{L_\top^a}$ is satisfied due to Eqn. (8).

Next, we use induct on $|L_1^b(e_{m-1}^\top) \cup L_2^b(e_{m-1}^\top)|$ using the VCs $\psi_{\text{ind1-0op}}^{L_1^b}$, $\psi_{\text{ind2-0op}}^{L_1^b}$, $\psi_{\text{ind1-0op}}^{L_2^b}$ or $\psi_{\text{ind2-0op}}^{L_2^b}$ to add all events in these sets. Finally, we induct on $|L_1^b(e_m^\top) \cup L_2^b(e_m^\top)|$ to again add all these events, thereby proving Eqn. (??).

Now that we have linearized e_m^\top using Eqn. (??), we need to linearize the events that led to $\text{merge}(l', a'', b'')$. Let's denote $L_1^b(e_m^\top)$ as M_1^a and $L_2^b(e_m^\top)$ as M_2^a . Now we induct on the size of the set $M_1^a \cup M_2^a$.

BASE CASE 1.1: ($|M_1^a \cup M_2^a| = 0$)

$a'' = a', b'' = b'$. By induction assumption, $\exists \pi. \text{lo}(C)_{|(L_\top^b \cup \cup_{e \in S} L_1^b(e) \cup \cup_{e \in S} L_2^b(e) \cup S)} \subseteq \pi$

1765 and $\text{merge}(l', a', b') = \pi(\sigma_0)$. Hence, $\pi_m = \pi \cdot e_m^\top$.

1766

1767 **INDUCTIVE CASE 1.1:**($|M_1^a \cup M_2^a| > 0$)

1768 We have 2 cases here:

1769 (1.1.1) Either of M_1^a or M_2^a is ϕ

1770 (1.1.2) Both M_1^a and M_2^a are not ϕ .

1771

1772 **CASE 1.1.1:**($M_1^a \neq \phi \wedge M_2^a = \phi$)

1773

1774

1775

1776

1777

1778

1779

1780

1781

1782

1783

1784

1785

1786

1787

1788

1789

1790

1791

1792

1793

1794

1795

1796

1797

1798

1799

1800

1801

1802

1803

1804

1805

1806

1807

1808

1809

1810

1811

1812

1813

Consider $e_1 \in M_1^a$ such that there does not exist $e \in M_1^a$ and $e_1 \xrightarrow{\text{lo}_m} e$, i.e. e_1 is the maximal event according to lo_m . Since lo ordering between events remains the same in all versions, and since versions v_1 and v_2 (which are being merged) were already linearizable, there would exist sequences leading to the states a such that e_1 would appear at the end. Hence, there exists a''' such that $a'' = e_1(a''')$. Since M_2^a is empty, all local events in L_2 are linearized before the rest of the LCA events. Suppose $L_\top^a \setminus \{e_m^\top\} \neq \phi$ or $L_\top^b \neq \phi$, the last event which leads to the state l', b'' must be an LCA event. Let's consider e_\top to be the maximal event in L_\top according to lo_m . Hence there exists states l'', b''' such that $l' = e_\top(l'')$, $b'' = e_\top(b''')$. By **BOTTOMUP-1-OP** rule

$$\text{merge}(e_\top(l''), e_1(a'''), e_\top(b''')) = e_1(\text{merge}(e_\top(l''), a''', e_\top(b'''))) \quad (9)$$

Again, we prove **BOTTOMUP-1-OP** rule using the same induction scheme that we showed for **BOTTOMUP-0-OP**. Briefly, we use $\psi_{\text{base-1op}}^{L_\top^b}$ and $\psi_{\text{ind-1op}}^{L_\top^b}$ for induction on $|L_\top^b|$. Then, we use $\psi_{\text{ind-1op}}^{L_\top^a}$, $\psi_{\text{ind1-1op}}^{L_1^b}$, $\psi_{\text{ind2-1op}}^{L_1^b}$, $\psi_{\text{ind1-1op}}^{L_2^b}$ and $\psi_{\text{ind2-1op}}^{L_2^b}$ to build the event sets $L_\top^a \setminus \{e_m^\top\}$ and $\sqcup_{e \in L_\top^a \setminus \{e_m^\top\}} L_1^b(e) \cup \sqcup_{e \in L_\top^a \setminus \{e_m^\top\}} L_2^b(e)$.

From the induction assumption, we get that $\text{merge}(e_\top(l''), a''', e_\top(b'''))$ is already obtained by the linearization of events applied on the initial state σ_0 . That is, there exists a sequence π' over events in $L_\top^b \cup \sqcup_{e \in S} L_1^b(e) \cup \sqcup_{e \in S} L_2^b(e) \cup S \cup M_1^a \setminus e_1$ which extends lo_m relation such that $\text{merge}(e_\top(l''), a''', e_\top(b''')) = \pi'(\sigma_0)$. Now, $\pi = \pi' e_1$ is the required linearization.

1793

1794

1795

1796

1797

1798

1799

1800

1801

1802

1803

1804

1805

1806

1807

1808

1809

1810

1811

1812

1813

CASE 1.1.2:($M_1^a \neq \phi \wedge M_2^a \neq \phi$)

Consider $e_1 \in M_1^a, e_2 \in M_2^a$ such that there does not exist $e \in M_i^a$ and $e_i \xrightarrow{\text{lo}_m} e$ (for $i = 1, 2$), i.e. each of the e_i s are maximal events according to lo_m . Since lo ordering between events remains the same in all versions, and since versions v_1 and v_2 (which are being merged) were already linearizable, there would exist sequences leading to the states a'' and b'' such that e_1 and e_2 would appear at the end resp. Hence, there exists a''' and b''' such that $a'' = e_1(a''')$ and $b'' = e_1(b''')$. Since $e_1 \parallel_C e_2$, they are related to each other by rc relation or they commute with each other i.e., $e_1 \xrightarrow{\text{rc}} e_2 \vee e_2 \xrightarrow{\text{rc}} e_1 \vee e_1 \rightleftarrows e_2$. We will consider the case when $e_2 \xrightarrow{\text{rc}} e_1 \vee e_1 \rightleftarrows e_2$. $e_1 \xrightarrow{\text{rc}} e_2$ is handled by **MERGE**COMMUTATIVITY. The equation becomes

$$\text{merge}(l', e_1(a'''), e_2(b''')) = e_1(\text{merge}(l', a''', e_2(b'''))) \quad (10)$$

which is the **BOTTOMUP-2-OP** rule.

Again, we prove **BOTTOMUP-2-OP** rule using the same induction scheme that we showed for **BOTTOMUP-1-OP**. Briefly, we use $\psi_{\text{base-2op}}^{L_\top^b}$ and $\psi_{\text{ind-2op}}^{L_\top^b}$ for induction on $|L_\top^b|$. Then, we use $\psi_{\text{ind-2op}}^{L_\top^a}$, $\psi_{\text{ind1-2op}}^{L_1^b}$, $\psi_{\text{ind2-2op}}^{L_1^b}$, $\psi_{\text{ind1-2op}}^{L_2^b}$ and $\psi_{\text{ind2-2op}}^{L_2^b}$ to build the event sets $L_\top^a \setminus \{e_m^\top\}$ and $\sqcup_{e \in L_\top^a \setminus \{e_m^\top\}} L_1^b(e) \cup \sqcup_{e \in L_\top^a \setminus \{e_m^\top\}} L_2^b(e)$.

1814 From the induction assumption, we get that $\text{merge}(l', a''', e_2(b'''))$ is already obtained by the
 1815 linearization of events applied on the initial state σ_0 . If π' is the linearization for this merge, then
 1816 $\pi = \pi' e_1$ is the required linearization.

1817

1818 CASE 2: ($|L_1^a \cup L_2^a| > 0$)

1819 The proof here will be identical to the proof of Inductive Case 1.1, substituting L_1^a and L_2^a for M_1^a
 1820 and M_2^a , and using the rules BOTTOMUP-1-OP, MERGECOMMUTATIVITY and BOTTOMUP-2-OP.

1821 A.2.8 Case (QUERY): Assume that a query operation is applied on a replica r . Let $C = \langle N, H, L, G, vis \rangle$
 1822 be the configuration of the replica before the operation. According to the semantics, the configura-
 1823 tion of the replica remains same after the query operation. By the induction hypothesis, Def. 3.9
 1824 holds for the configuration C .
 1825
 1826 □

1827

1828

1828 A.3 Buggy MRDT implementation in [23]

1829

1830

1: $\Sigma = (\mathbb{N} \times \text{bool})$

1831

2: $O = \{\text{enable, disable}\}$

1832

3: $Q = \{\text{rd}\}$

1833

4: $\sigma_0 = (0, \text{false})$

1834

5: $\text{do}(\sigma, _ _ _, \text{enable}) = (\text{fst}(\sigma) + 1, \text{true})$

1835

6: $\text{do}(\sigma, _ _ _, \text{disable}) = (\text{fst}(\sigma), \text{false})$

1836

$$7: \text{merge_flag}((lc, lf), (ac, af), (bc, bf)) = \begin{cases} \text{true,} & \text{if } af = \text{true} \ \&\& \ bf = \text{true} \\ \text{false,} & \text{else if } af = \text{false} \ \&\& \ bf = \text{false} \\ ac > lc, & \text{else if } af = \text{true} \\ bc > lc, & \text{otherwise} \end{cases}$$

1840

8: $\text{merge}(\sigma_{\top}, \sigma_a, \sigma_b) = (\text{fst}(\sigma_a) + \text{fst}(\sigma_b) - \text{fst}(\sigma_{\top}), \text{merge_flag}(\sigma_{\top}, \sigma_a, \sigma_b))$

1841

9: $\text{query}(\sigma, rd) = \text{snd}(\sigma)$

1842

10: $\text{rc} = \{\{\text{disable, enable}\}\}$

1843

Fig. 15. Enable-wins flag MRDT implementation from [23]

1844

1845

1846

1847

1848

1849

1850

1851

1852

1853

1854

1855

1856

1857

1858

1859

1860

1861

1862

VC Name	Pre-condition	Post-condition
MERGE _{COMMUTATIVITY}		$\mu(a, b) = \mu(b, a)$
MERGE _{IDEMPOTENCE}		$\mu(s, s) = s$
$\psi_{\text{base-2op}}^{L^b_T}$	$e_2 \xrightarrow{rc} e_1 \vee e_2 \xrightarrow{rc} e_1$	$\mu(e_1(\sigma_0), e_2(\sigma_0)) = e_1(\mu(\sigma_0, e_2(\sigma_0)))$
$\psi_{\text{ind-2op}}^{L^b_T}$	$e_2 \xrightarrow{rc} e_1 \vee e_2 \xrightarrow{rc} e_1$	$\mu(e_1(l), e_2(l)) = e_1(\mu(l, e_2(l)))$ $\mu(e_1 \cdot e_T(l), e_2 \cdot e_T(l)) = e_1(\mu(e_T(l), e_2 \cdot e_T(l)))$
$\psi_{\text{ind-2op}}^{L^a_T}$	$(e_2 \xrightarrow{rc} e_1 \vee e_2 \xrightarrow{rc} e_1) \wedge (\exists e.e \xrightarrow{rc} e_T)$	$\mu(e_1(a), e_2(b)) = e_1(\mu(a, e_2(b)))$ $\mu(e_1 \cdot e_T(a), e_2 \cdot e_T(b)) = e_1(\mu(e_T(a), e_2 \cdot e_T(b)))$
$\psi_{\text{ind1-2op}}^{L^b_1}$	$(e_2 \xrightarrow{rc} e_1 \vee e_2 \xrightarrow{rc} e_1) \wedge e_b \xrightarrow{rc} e_T$	$\mu(e_1 \cdot e_T(a), e_2 \cdot e_T(b)) = e_1(\mu(e_T(a), e_2 \cdot e_T(b)))$ $\mu(e_1 \cdot e_T \cdot e_b(a), e_2 \cdot e_T(b)) = e_1(\mu(e_T \cdot e_b(a), e_2 \cdot e_T(b)))$
$\psi_{\text{ind2-2op}}^{L^b_1}$	$(e_2 \xrightarrow{rc} e_1 \vee e_2 \xrightarrow{rc} e_1) \wedge e_b \xrightarrow{rc} e_T$	$\mu(e_1 \cdot e_T \cdot e_b(a), e_2 \cdot e_T(b)) = e_1(\mu(e_T \cdot e_b(a), e_2 \cdot e_T(b)))$ $\mu(e_1 \cdot e_T \cdot e_b \cdot e(a), e_2 \cdot e_T(b)) = e_1(\mu(e_T \cdot e_b \cdot e(a), e_2 \cdot e_T(b)))$
$\psi_{\text{ind1-2op}}^{L^b_2}$	$(e_2 \xrightarrow{rc} e_1 \vee e_2 \xrightarrow{rc} e_1) \wedge e_b \xrightarrow{rc} e_T$	$\mu(e_1 \cdot e_T(a), e_2 \cdot e_T(b)) = e_1(\mu(e_T(a), e_2 \cdot e_T(b)))$ $\mu(e_1 \cdot e_T(a), e_2 \cdot e_T \cdot e_b(b)) = e_1(\mu(e_T(a), e_2 \cdot e_T \cdot e_b(b)))$
$\psi_{\text{ind2-2op}}^{L^b_2}$	$(e_2 \xrightarrow{rc} e_1 \vee e_2 \xrightarrow{rc} e_1) \wedge e_b \xrightarrow{rc} e_T$	$\mu(e_1 \cdot e_T(a), e_2 \cdot e_T(b)) = e_1(\mu(e_T(a), e_2 \cdot e_T(b)))$ $\mu(e_1 \cdot e_T \cdot e_b(a), e_2 \cdot e_T(b)) = e_1(\mu(e_T \cdot e_b(a), e_2 \cdot e_T(b)))$
$\psi_{\text{ind-2op}}^{L^a_1}$	$e_2 \xrightarrow{rc} e_1 \vee e_2 \xrightarrow{rc} e_1$	$\mu(e_1(a), e_2(b)) = e_1(\mu(a, e_2(b)))$ $\mu(e_1 \cdot e'_1(a), e_2(b)) = e_1(\mu(e'_1(a), e_2(b)))$
$\psi_{\text{ind-2op}}^{L^a_2}$	$e_2 \xrightarrow{rc} e_1 \vee e_2 \xrightarrow{rc} e_1$	$\mu(e_1(a), e_2(b)) = e_1(\mu(a, e_2(b)))$ $\mu(e_1(a), e_2 \cdot e'_2(b)) = e_1(\mu(a, e_2 \cdot e'_2(b)))$
$\psi_{\text{base-1op}}^{L^b_T}$		$\mu(e_1(\sigma_0), \sigma_0) = e_1(\mu(\sigma_0, \sigma_0))$
$\psi_{\text{ind-1op}}^{L^b_T}$		$\mu(e_1(l), l) = e_1(\mu(l, l))$ $\mu(e_1 \cdot e_T(l), e_T(l)) = e_1(\mu(e_T(l), e_T(l)))$
$\psi_{\text{ind-1op}}^{L^a_T}$	$\exists e.e \xrightarrow{rc} e_T$	$\mu(e_1(a), e'_1(b)) = e_1(\mu(a, e'_1(b)))$ $\mu(e_1 \cdot e_T(a), e_T \cdot e'_1(b)) = e_1(\mu(e_T(a), e_T \cdot e'_1(b)))$
$\psi_{\text{ind1-1op}}^{L^b_1}$	$e_b \xrightarrow{rc} e_T$	$\mu(e_1 \cdot e_T(a), e_T(b)) = e_1(\mu(e_T(a), e_T(b)))$ $\mu(e_1 \cdot e_T \cdot e_b(a), e_T(b)) = e_1(\mu(e_T \cdot e_b(a), e_T(b)))$
$\psi_{\text{ind2-1op}}^{L^b_1}$	$e_b \xrightarrow{rc} e_T \wedge (\neg e \xrightarrow{rc} e_b \vee e \xrightarrow{rc} e_T)$	$\mu(e_1 \cdot e_T \cdot e_b(a), e_T(b)) = e_1(\mu(e_T \cdot e_b(a), e_T(b)))$ $\mu(e_1 \cdot e_T \cdot e_b \cdot e(a), e_T(b)) = e_1(\mu(e_T \cdot e_b \cdot e(a), e_T(b)))$
$\psi_{\text{ind1-1op}}^{L^b_2}$	$e_b \xrightarrow{rc} e_T$	$\mu(e_1 \cdot e_T(a), e_T(b)) = e_1(\mu(e_T(a), e_T(b)))$ $\mu(e_1 \cdot e_T(a), e_T \cdot e_b(b)) = e_1(\mu(e_T(a), e_T \cdot e_b(b)))$
$\psi_{\text{ind2-1op}}^{L^b_2}$	$e_b \xrightarrow{rc} e_T \wedge (\neg e \xrightarrow{rc} e_b \vee e \xrightarrow{rc} e_T)$	$\mu(e_1 \cdot e_T(a), e_T \cdot e_b(b)) = e_1(\mu(e_T(a), e_T \cdot e_b(b)))$ $\mu(e_1 \cdot e_T(a), e_T \cdot e_b \cdot e(b)) = e_1(\mu(e_T(a), e_T \cdot e_b \cdot e(b)))$
$\psi_{\text{ind-1op}}^{L^a_1}$		$\mu(e_1(a), e_T(b)) = e_1(\mu(a, e_T(b)))$ $\mu(e_1 \cdot e'_1(a), e_T(b)) = e_1(\mu(e'_1(a), e_T(b)))$
$\psi_{\text{base-0op}}^{L^b_T}$		$\mu(e_1(\sigma_0), e_1(\sigma_0)) = e_1(\mu(\sigma_0, \sigma_0))$
$\psi_{\text{ind-0op}}^{L^b_T}$		$\mu(e_1(l), e_1(l)) = e_1(\mu(l, l))$ $\mu(e_1 \cdot e_T(l), e_1 \cdot e_T(l)) = e_1(\mu(e_T(l), e_T(l)))$
$\psi_{\text{ind-0op}}^{L^a_T}$	$\exists e.e \xrightarrow{rc} e_T$	$\mu(e_1(a), e_1(b)) = e_1(\mu(a, b))$ $\mu(e_1 \cdot e_T(a), e_1 \cdot e_T(b)) = e_1(\mu(e_T(a), e_T(b)))$
$\psi_{\text{ind1-0op}}^{L^b_1}$		$\mu(e_1(a), e_1(b)) = e_1(\mu(a, b))$ $\mu(e_1 \cdot e_b(a), e_1(b)) = e_1(\mu(e_b(a), b))$
$\psi_{\text{ind2-0op}}^{L^b_1}$	$\neg e \xrightarrow{rc} e_b \vee e \xrightarrow{rc} e_1$	$\mu(e_1 \cdot e_b(a), e_1(b)) = e_1(\mu(e_b(a), b))$ $\mu(e_1 \cdot e_b \cdot e(a), e_1(b)) = e_1(\mu(e_b \cdot e(a), b))$
$\psi_{\text{ind1-0op}}^{L^b_2}$		$\mu(e_1(a), e_1(b)) = e_1(\mu(a, b))$ $\mu(e_1(a), e_1 \cdot e_b(b)) = e_1(\mu(a, e_b(b)))$
$\psi_{\text{ind2-0op}}^{L^b_2}$	$\neg e \xrightarrow{rc} e_b \vee e \xrightarrow{rc} e_1$	$\mu(e_1(a), e_1 \cdot e_b(b)) = e_1(\mu(a, e_b(b)))$ $\mu(e_1(a), e_1 \cdot e_b \cdot e(b)) = e_1(\mu(a, e_b \cdot e(b)))$

Table 4. Complete set of Verification Conditions for CRDTs