

# Ordering Operations for Generic Replicated Data Types using Version Trees

Nazmus Saquib  
Univ. of California, Santa Barbara  
USA  
nazmus@cs.ucsb.edu

Chandra Krintz  
Univ. of California, Santa Barbara  
USA  
ckrintz@cs.ucsb.edu

Rich Wolski  
Univ. of California, Santa Barbara  
USA  
rich@cs.ucsb.edu

## Abstract

Data replication facilitates availability and recovery in a distributed environment. However, concurrent updates to multiple replicas result in divergence of data. Conflict-Free Replicated Data Types (CRDTs) are abstract data types that provide a principled approach to asynchronously reconcile this divergence. We propose a different perspective on the divergence of data whereby we treat data divergences as *versions* of the data. That is, instead of treating it only as a problem that needs to be solved, we consider it also to be a feature that provides a way to track versioning and evolution of data. Versioning information is helpful in multiple scenarios, such as provenance tracking and system debugging. Doing so allows us to leverage concepts such as the *version tree* found in the literature for persistent (versioned) data structures. We show that many techniques used in CRDTs to order elements can be derived from version trees, which predates CRDTs by more than 20 years. Using version trees for maintaining order and append-only logs for storage, we propose a method to ensure convergence of arbitrary data types, while maintaining information related to the evolution of data.

**CCS Concepts:** Computing methodologies ! Distributed algorithms.

**Keywords:** CRDT, data structure, replication, data versioning

## ACM Reference Format:

Nazmus Saquib, Chandra Krintz, and Rich Wolski. 2022. Ordering Operations for Generic Replicated Data Types using Version Trees. In *9th Workshop on Principles and Practice of Consistency for Distributed Data (PaPoC'22)*, April 5–8, 2022, RENNES, France. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3517209.3524038>



5IJTXPSLJTMJDFOTFEVDFEFSBSSFBUJWF\$PNNPOT\*UUSJCVUJPO\*OUFSOBUJPOBMJDFOTF  
PaPoC'22, April 5–8, 2022, RENNES, France  
©2022 Copyright held by the owner/author(s).  
ACM ISBN 978-1-4503-9256-3/22/04.  
<https://doi.org/10.1145/3517209.3524038>

## 1 Introduction

Distributed systems often replicate data for low latency and high availability. The shared state among the replicas is maintained according to different system models. For example, strong consistency requires a replica to coordinate with other replicas to execute an operation. Coordination increases latency and an operation may fail as a result of a network partition, preventing the necessary communication between replicas. Weaker consistency models such as eventual consistency enable replicas to execute an operation locally and asynchronously propagate the operation to other replicas. This results in lower latency but with a temporary divergence in replica views that must be eventually reconciled. Conflict-Free Replicated Data Types (CRDTs) [14, 20, 21] are abstract data types that provide a principled approach for this asynchronous reconciliation. CRDTs support a weaker model than strong consistency, namely, Strong Eventual Consistency (SEC) [1]. SEC guarantees that whenever two replicas receive the same set of updates, possibly in a different order, they reach the same state.

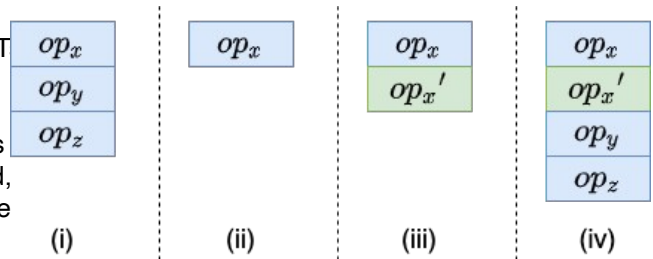
Broadly, there are two types of CRDTs: state-based and operation-based (or op-based) [1]. In state-based CRDTs, an operation is executed on the local replica state. A replica periodically propagates its state to other replicas to achieve consistency. A disadvantage of this approach is the communication overhead associated with shipping the full state, which at times can be large. Apart from this, state-based CRDTs require data type-specific *merge* function that provides a *join* for any pair of replica states, such that the set of all states forms a *semilattice*. In op-based CRDTs, an operation is executed on the local replica and the operation is asynchronously propagated to other replicas. Although op-based CRDTs do not communicate state, they require exactly-once causal broadcast. Moreover, op-based CRDTs require all operations on a data type to be commutative. Delta State Conflict-Free Replicated Data Types (CRDTs) [1] combine the advantages of state-based and op-based CRDTs. Like the state-based CRDTs, they can tolerate unreliable networks and, in particular, do not require exactly-once causal broadcast as a communication network property. Moreover, like the op-based approach, they do not require full replica state to be communicated, but rather, they communicate only state changes or “deltas”.

Due to the commutativity requirement of op-based CRDTs and the join semilattice requirement of the state-based CRDT often we must resort to using a restricted form of a data type rather than the conventional one. For example, the CRDT literature describes a number of variations for set, such as a grow-only set (G-set) where elements can only be added, two-phase set (2P-set) where a removed element cannot be added later, etc. In general, CRDTs for custom data types [15] have not been extensively studied in the literature, and the use of such data types often requires ad-hoc solutions.

One way to guarantee consistency among replicas of an arbitrary data type is to ensure all the replicas execute the same set of operations in the same order. In fact, this is the principle used in protocols such as Raft [16], which provides strong consistency. However, in a weaker consistency model where we allow the log of operations to diverge among the replicas, this might entail extra work. Specifically, we might need to rollback the log of operations up to a certain point, introduce new operation(s), and re-execute the previously rolled back operation(s). Figure 1 illustrates one such scenario. We represent a replica with  $B$  and its log of operations as  $(op_x, op_y, op_z)$ . Let us assume  $B$  executed  $op_x$ ,  $op_y$ , and  $op_z$  in order (first column). Later on,  $B$  came to know of operation  $op_x'$  which must be executed after  $op_x$ . Therefore,  $B$  first rolls back the log of operations up to  $op_x$  (second column), then executes the new operation  $op_x'$  (third column), and finally re-executes  $op_y$  and  $op_z$  (fourth column). During the log rollback, the state changes resulting from the rolled back operations (in this case  $op_y$  and  $op_z$ ) must be rolled back as well.

Merely applying a seemingly inverse operation might not have the intended outcome of rolling back operations. For example, roll back of an insertion into a binary search tree cannot be performed by a corresponding deletion as this might result in a different tree structure than the one resulting from the insertion operation not being executed in the first place. Therefore, we need a mechanism to revert the underlying data structure to a state where the rolled back operations were not even applied, rather than executing operations that appear to negate the effect of previous operations. One generic way to achieve this is to represent the state of the underlying data type using append-only logs as well. Then to ensure that the state of the underlying data type is rolled back along with the log of operations, the replica need only roll back entries from the tails of the logs used to represent the state of the data type. A similar technique is used for the move operation in a replicated tree data model as well [8]. While representing states for simple data types such as registers using logs is simple; more involved data types such as linked lists and binary search trees require complex algorithms [18].

The CRDT literature contains multiple works on maintaining order in list or sequence such as a causal tree (6,16) [



**Figure 1.** Log rollback to reach correct order of operations (from  $(op_x, op_y, op_z)$  to  $(op_x, op_x', op_y, op_z)$ ).

replicated growable array (RGA) [7, 20], LSEQ [11], Tree-doc [15], and WOOT [13]. Most of these works were originally developed for collaborative text-editing environments. However, we can use these algorithms to maintain an ordered list of operations (i.e.,  $(op_x, op_x')$ ) for arbitrary data types as well. We show that RGA shares similarities with version trees [4]. Version trees are tree data structures used to record causal relationship in versioned data structures (also known as *persistent data structures*). This implies that it is possible to incorporate data versioning while maintaining the order of operations. Data versioning has multiple use cases, such as facilitating system debugging, efficiently answering temporal query, and tracking evolution of data [8].

In this work, we first show how maintaining order of operations can lead to data versioning. At the same time, we explain the principle which allows us to track versions of arbitrary data types. Then, we show that version lists (text representations of version trees) for versioned data structures can use the same algorithm as RGA. A popular algorithm to maintain order in list or sequence CRDTs. This in turn implies any algorithm based on RGA has its root in version trees. In essence, the principle used for maintaining order in many list or sequence CRDTs lies in a 30 year old work on versioned data structures; which, interestingly, was formulated for a single-machine system with volatile memory. As versioning information can be captured while maintaining order of operations, this allows us to incorporate data versioning in our system without additional work; thus facilitating system debugging, data provenance, and efficient temporal query.

## 2 Persistent Data Structures

Conventional data structures are *ephemeral*, i.e., an update operation mutates the current state of the underlying data structure, resulting in a new state. However, many applications in computational geometry [9, 22] and text editing [16] can benefit from data versioning. Moreover, versioned data structures can facilitate system debugging and efficient temporal query [18].

In versioned data structures, an update operation results in a new version of the data structure while keeping records

of all the previous versions. Versioned data structures are also known as persistent data structures [4] in the literature. A versioned data structure is *partially persistent* if all versions can be accessed but only the latest can be updated. A versioned data structure is *fully persistent* if all versions can be both accessed and updated. Although the literature contains detailed descriptions of both specific [10] and generic [4] methods to make data structures persistent, in this work we are more concerned with how a *version stamp* is represented and ordered.


In persistent data structures, each version is tagged with a monotonically increasing version stamp (an integer value). The first version of a data structure has the version stamp 1. Whenever an update operation is applied to the data structure resulting in a new version, the version stamp is incremented by 1 and the new version is tagged with this incremented value. Throughout this paper, we use version stamp to refer to both the version of the underlying data structure and the operation that resulted in that version. The intended use will be clear from the context.

Although for partially persistent data structures this versioning scheme results in a natural linear order, fully persistent data structures only have a partial order over the version stamps. This partial ordering is defined by a rooted *version tree*. Each node in a version tree contains a version stamp. A directed edge from node  $D$  to node  $E$  denotes that version  $E$  was obtained by updating version  $D$ . The *youngest* child, i.e., the latest version among the children of a parent is always placed at the leftmost position (cf. Figure 2). One way to impose a total order on the partial order represented through a version tree is to perform a preorder traversal and aggregate the traversed nodes in a list known as the *version list*. In fact, we do not need to maintain the version tree explicitly to create the version list. If we simply insert a new version immediately after its parent in the list, it maintains the preorder traversal, which can be proved using mathematical induction [4]. This insertion scheme also implies that the version list has a special property: for any version  $D$  the descendants of  $D$  in the version tree occur consecutively in the version list, starting with  $D$ .

### 3 Version Trees as Replica States

The versioning scheme described in Section 2 can be used to model the divergence (using version trees) and the subsequent convergence (using version lists) among replicas in a distributed system. We consider a distributed system with  $n$  replicas (nodes). Each replica is assigned a replica ID from 1 to  $n$ . Each replica maintains a log of operations (cf. Section 1) where the elements in set  $L$  can be sorted according to some criteria, e.g., lexicographically. As we will see in Section 3.1, it is crucial to have this sorting ability for achieving consistency.

We assume that each replica  $R_i$  maintains a log of operations  $L_i$  (cf. Section 1) implicitly



**Figure 2.** Example of two version trees. In the linear tree, a new version is obtained by always updating the latest version. Tree 2 illustrates a scenario where a new version can be obtained by updating some previous version. Although version 3 in 2 was obtained by updating the then latest version 2, version 5 was also created by updating version 2. Note that as 5 is greater (younger) than 3, 5 is placed to the left of 3. In general, the children of a node in a version tree are arranged in descending order of their version stamps. A version list of 2 is  $[1, 2, 5, 6, 3, 4]$ . records a linear version tree  $(- \beta)$ . Note that even if the underlying data structure is not versioned (i.e., it is mutable), we can still map the divergence and convergence using version tree and version list respectively. In essence, we are using the *causal* or *happens-before* relationship embedded within the version tree/list to achieve consistency. This relationship will be present among the operations of any data structure, irrespective of the data structure being ephemeral or persistent.

We further assume that to achieve consistency, each replica reads the OpLog of another replica at a regular interval and performs a *merge* step. A merge step is always between two replicas, one known as the *source* and the other known as the *reader*. During a merge step, the reader incorporates all the operations unknown to it but known to the source into the reader's OpLog. Note that a merge step is unidirectional, i.e., operations known to the reader but unknown to the source are incorporated into the OpLog of the source during some other merge step where the current source is the reader. Once all replicas observe the full set of operations and apply a consistent ordering scheme, the system achieves consistency. As described in Section 3.3, the ordering scheme involves implicitly creating a version tree from the OpLogs of the source and the reader, followed by creating a linear tree from a modified preorder traversal of the version tree, which we call *InterleavedPreorder* (cf. Section 3.2).

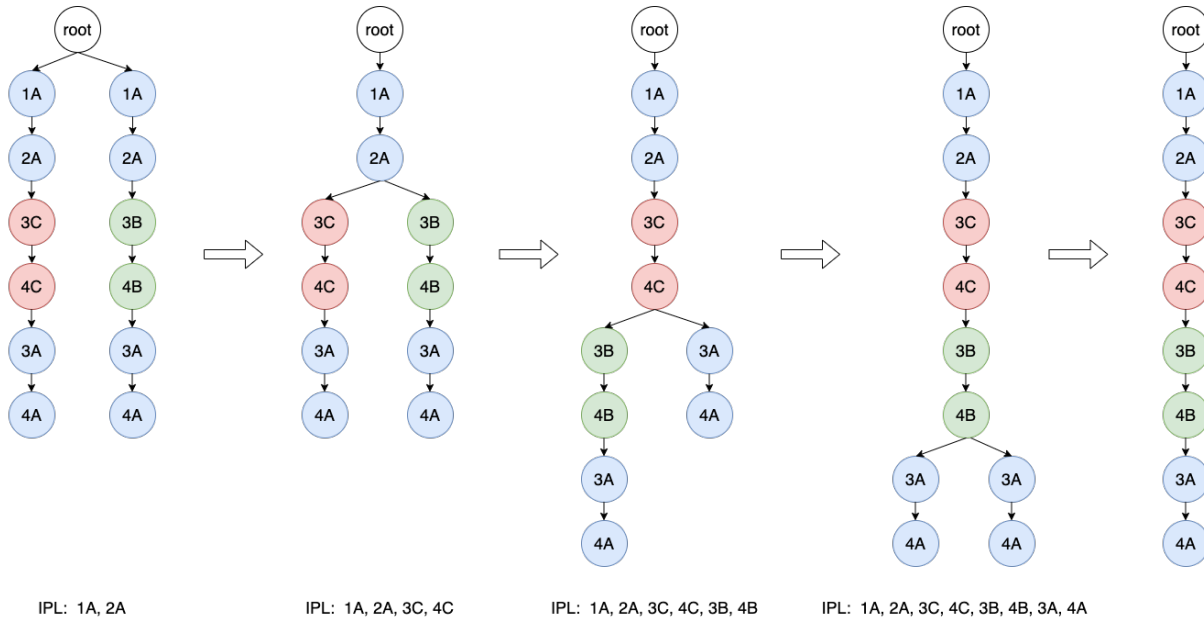


Figure 3. InterleavedPreorder traversal.

### 3.1 Version Stamps

The version stamps described in Section 2 are monotonically increasing. This is relatively simple to maintain in a single-machine system. However, in a distributed system with multiple machines, it requires either a special sequencer node or a complex coordination to ensure that the next higher version stamp is allotted to only one replica. Hence, we propose using a concatenation of a monotonically increasing counter value (local to each replica) and the replica ID to represent the version stamp. We represent the counter and the replica ID of a version stamp by  $EB$  and  $D$  respectively. We say version stamp  $EB$  is less than version stamp  $EB'$  if (i) the counter of  $EB$  is less than that of  $EB'$  or (ii) both the counters are the same but the replica ID of  $EB$  is less than that of  $EB'$ .

When replica  $B$  executes a new operation in response to a client (i.e., a process that can send update/access request to any replica) request,  $B$  tags the operation with version stamp  $EB$  which is greater than all other version stamps it has observed so far (operations that happened before). Thus if operation  $O_0$  happens before  $O_1$ ,  $EB_0 < EB_1$  where  $EB_0$  and  $EB_1$  are the version stamps of operations  $O_0$  and  $O_1$ , respectively. We further assume that any version stamp is greater than the null version stamp, i.e., a version stamp having invalid/null values for counter and replica ID fields. This form of version stamp is essentially the Lamport timestamp [9]. This allows us to generate version stamps based only on local information without any complex coordination, all the while capturing the causal or happens-before relationship.

### 3.2 InterleavedPreorder Traversal

In this section, we define the interleavedPreorder traversal for trees with at most two branches. Trees with more branches can be generalized from this construction, by “cutting” branches in pairs (cf. Section 3.4). Each node in a version tree contains three fields: (i) a version stamp, (ii) a left pointer, and (iii) a right pointer. We represent a node  $G$  as  $(G.V, G.L, G.R)$ . If a node has only one child, the left pointer points to the child whereas the right pointer points to the null node  $q$ .

To generate the InterleavedPreorder traversal list  $\mathcal{L}$ , we first add all nodes up to the branching point in  $\mathcal{L}$ . We maintain two pointers, one for the top of the left branch and another for the top of the right branch, advancing them when the corresponding nodes are incorporated in  $\mathcal{L}$ . Next, we start adding nodes from the left branch until we reach a node  $G$  such that the version stamp of  $G$  is smaller than or equal to that of the topmost node of the right branch. In case the version stamps of the top of the branches are equal, we advance both the pointers. Once the top of the right branch has a greater version stamp than that of  $G$ , we create an intermediate tree where each node in  $\mathcal{L}$  appears one after another in a linear structure. The branch pointed to by the right pointer becomes the new left branch and the portion of the previous left branch starting from  $G$  becomes the new right branch. We then perform InterleavedPreorder traversal on this intermediate tree. This continues recursively until we end up with a tree having only a single branch. A linear traversal of this tree gives us the final value of  $\mathcal{L}$ . Algorithm 1 presents the recursive procedure to populate  $\mathcal{L}$ . The initial call to the procedure is made with a dummy



root  $\ell$  connecting two linear trees and an empty list where the IPL is populated. Figure 3 illustrates InterleavedPreorder traversal on a tree with two branches.

**Algorithm 1** InterleavedPreorder Traversal

```

Require: VT node  $\ell$ , list
Ensure: IPL is populated in
1: procedure  $\diamond\diamond\diamond\diamond\diamond\diamond\diamond\diamond$  IPL( $\ell$ )
2:   if  $\ell = q$  then
3:     return
4:   end if
5:   while  $\ell \neq q \wedge [A86\#C]q$  do  $\ell$  traverse linear part
6:     [  $\ell$ ;  $\ell$ ;  $\ell$  C
7:      $\ell$ ?? $\ell$ = $\ell$ [  $\ell$ .EB
8:   end while
9:   if  $\ell \neq q \wedge [A86\#C]q$  then  $\ell$  all nodes traversed
10:    return
11:  end if
12:  [  $\ell$ ;  $\ell$ ;  $\ell$  C
13:  [A [A86#C
14:  [A86#Cq
15:  while  $\ell < q \wedge [A < q \wedge [EB [AEB]$  do  $\ell$  add from the
    left branch until a version stamp is observed that is less than
    that of the version stamp at the top of the right branch
16:    [  $\ell$ ;  $\ell$ ;  $\ell$  C;
17:    [ = [  $\ell$ ;  $\ell$  C
18:     $\ell$ ?? $\ell$ = $\ell$ [  $\ell$ .EB
19:    if [  $\ell$ .EB [AEB] then
20:      [A = [A;  $\ell$  C
21:    end if
22:    [  $\ell$ ; = [  $\ell$ ; ;
23:  end while
24:  [  $\ell$ ;  $\ell$ ;  $\ell$  C  $\ell$  swap left and right branch for the next
    recursive call
25:  [A86#C];
26:   $\diamond\diamond\diamond\diamond\diamond\diamond\diamond\diamond$  IPL( $\ell$ )
27: end procedure
    
```

**3.3 Mapping Divergence/Convergence between Replicas using Version Trees**

An OpLog implicitly records a linear version tree. During a merge step, the reader scans the OpLogs of both the sources and the reader from the top and skips over all the common elements until it finds a mismatch. Note that although this can be optimized so that the reader does not have to scan the logs from the top, we leave it out from our discussion as it is not the primary focus of our current exposition. Once it finds a mismatch, this indicates the occurrence of concurrent operations, i.e., two operations that cannot be ordered according to a happens-before relationship or causality. We can consider the paths after this point of mismatch as two branches of the version tree. In accordance with how version trees are created, we put the branch containing the greater version stamp at the top to the left of the branch containing

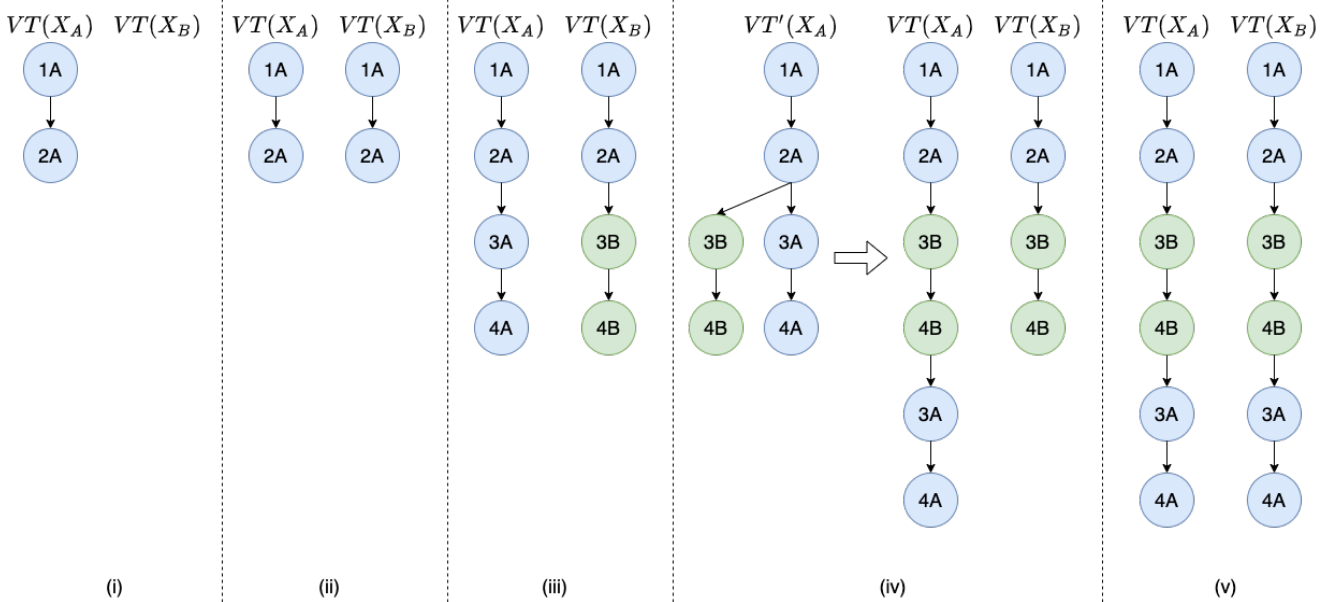
the smaller version stamp. After that, we perform an InterleavedPreorder traversal of this intermediate tree which results in the final merged version tree.

Figure 4 illustrates how we can model the divergence between two replicas as a version tree. In column (i) executes two operations 1 and 2. In column (ii),  $\ell$  (reader) performs a merge step with  $\ell$  (source). In this case, as the version tree of the reader, i.e.,  $\ell$  is empty, it trivially merges to that of the source. In column (iii), both  $\ell$  and  $\ell$  independently execute two operations. Note how the version stamps for these operations can be derived from local information alone. In column (iv),  $\ell$  (reader) performs a merge step with  $\ell$  (source).  $\ell$  skips over the first two elements in  $\ell$  ( $\ell$ ) and  $\ell$  ( $\ell$ ) as the corresponding elements are the same. However, after that  $\ell$  ( $\ell$ ) has 3 and  $\ell$  ( $\ell$ ) has 3. As 3  $\ell$  3, the branch containing 3 is placed to the left of the branch containing 3 to create the intermediate tree  $\ell$  ( $\ell$ ). Finally, an InterleavedPreorder “flattening” of  $\ell$  ( $\ell$ ) results in  $\ell$  ( $\ell$ ). In column (v),  $\ell$  (reader) performs a merge step with  $\ell$  (source). All the current elements in  $\ell$  ( $\ell$ ) are present in the same order from the top in  $\ell$  ( $\ell$ ). Therefore,  $\ell$  skips over all these elements and simply adds the rest of the elements from  $\ell$  ( $\ell$ ) to the tail of  $\ell$  ( $\ell$ ). At this point, both  $\ell$  and  $\ell$  have observed the same set of operations (order of observance was different) and have incorporated all these operations in the same order (the final order in OpLogs is the same), i.e., the system is in a consistent state. Note that even if the order of merge steps was altered in columns (iv) and (v), we arrive at the same final state as shown in column (v).

Although at first sight it seems we could have used a simple preorder traversal instead of the complex Interleaved-Preorder traversal, this is not the general case. As an exception, we consider the version tree in Figure 3. A possible sequence of actions that lead to the initial tree of Figure 3 is: (i)  $\ell$  executes 1 2. (ii)  $\ell$  and  $\ell$  (both separately as the reader) merges with  $\ell$  (source). (iii)  $\ell$ ,  $\ell$  and  $\ell$  execute 3 4 ; 3  $\ell$   $\ell$  and 3...4... respectively. (iv)  $\ell$  (reader) merges with  $\ell$  (source). (v)  $\ell$  (reader) merges with  $\ell$  (source). (vi) Finally,  $\ell$  (reader) attempts to merge with  $\ell$  (source). If we perform a simple preorder traversal, this results in duplicated operations for 3 and 4. Hence, we need InterleavedPreorder traversal.

**3.4 Convergence among More than Two Replicas**

In a distributed system with more than two replicas, it might appear that we can end up with a version tree with more than two branches. However, from our discussion in Section 3, we know that a merge step is performed among two replicas at a time. Hence, while we can represent the state of multiple replicas using a single version tree with two or more branches, we can also represent incrementally updated states as a sequence of version trees having at most two



**Figure 4** Merging version trees (i) - executes operations concurrently but independently. (ii) - merges with. (iii) - and - executes operations concurrently but independently. (iv) - merges with. (v) - merges with.

branches. In this case, the reader merges with replicas one-unknown operations have to be placed before the existing by-one following some strategy (random, round-robin, etc.) ones (cf. Fig 1). This means that the underlying data structure until it observes all the operations. At the start of each merge also needs to rollback its states which were obtained through step, the reader creates a version tree with two branches, on the application of the rolled back operations. Depending on branch being its own OpLog and the other being the remote the exact data structure, rolling back state can be complex. (i.e. source) OpLog.

Note that the convergence does not depend on the order inrolled back by performing a corresponding decrement operation which two replicas execute the merge steps. To understand how InterleavedPreorder traversal results in a consistent state of the underlying data structure once every operation has been observed by all the replicas, we note a couple of points. First, upon branching the version stamps in a version tree are ordered arbitrarily but deterministically. This follows from the construction of the version tree, where the branch containing the greater version stamp at the top is placed to the left of the branch containing the smaller version stamp at the top. Second, InterleavedPreorder traversal does not change any existing relative order of two version stamps in a version tree (and thus the OpLog), only imposes a new order among version stamps that were not previously ordered (i.e., version stamps in two different branches). This holds as the IPL includes all the nodes in the version tree in order up to the branching point, upon which it includes unique version stamps deterministically. Hence once all the replicas have observed the same set of operations, the system achieves consistency.

### 3.5 Supporting Generic Data Structures

Although InterleavedPreorder traversal facilitates the consistent ordering of operations, we must rollback OpLog if

For example, an increment operation on a counter can be rolled back by performing a corresponding decrement operation and vice versa. However, a binary search tree insertion cannot always be rolled back by performing a corresponding deletion, because the resulting tree structure might differ from the one that would have existed had the insert not taken place at all.

One way to ensure the underlying data structure states can also be rolled back is to represent the states using append-only logs as well. While representing data types such as registers and counters using logs can be relatively simple, other data structures such as linked lists and binary search trees require complex algorithms. Once a data structure is represented using logs, we must maintain a map from the operations to the range of sequence numbers of the log(s) used to represent the state of the data structure. Then during rollback of the OpLog, we can use this map to perform rollback on other log(s) as well.

As an example, let us assume that apart from appending to the  $log(-)$  for each operation, the underlying data structure needs to perform one additional append to an auxiliary log,  $DG(-)$ . Then to rollback the last operations, we must rollback the last entries from both  $log(-)$  and  $DG(-)$ . In practice, an operation might append to more than one log and different operations might append

to different sets of logs. Hence this information must be recorded in a map, where the key is the unique version stamp of the operation and the value is a list of tuples. Each tuple contains the name of an auxiliary log and the range of sequence numbers corresponding to the entries in that log appended during the execution of operation.

### 4 Replicated Growable Arrays (RGA)

As an example of a CRDT used to maintain order in a list or sequence, we review Replicated Growable Array (RGA). [RGA uses the same algorithm of Causal Trees (CT) as shown in 2]. We also explore the similarity between the version tree approach and RGA (and thus other approaches including CT and any method based on RGA). This reveals that the principles used in these algorithms are the same ones used in version trees, which predates these approaches by more than 20 years.

RGA implements a sequence as a linked list. It supports operations to add element immediately after element. Elements can be uniquely identified using timestamps which are ordered consistently with causality. If a client invokes addRight operations twice at the same place one after another, followed by the latter insert occurs to the left of the former and has a higher timestamp.

This is the exact same construction of from version trees during a merge step as described in Section 3. To see this, let us consider the OpLog of the reader as a linked list. The OpLog of the source then implicitly provides a list of addRight operations that need to be incorporated into the OpLog of the reader. Specifically, if element immediately follows in the OpLog of the reader, this translates to the operation. If and are already present in the OpLog of the reader one after another, nothing is changed. This is equivalent to the case in IPL formation where the two branches have the same elements (i.e., version stamps), in which case values from only one branch is added to IPL but a pointer for the other branch is advanced to avoid duplicates. If is the last element of OpLog of the source, can be simply appended, e.g., transition from column (iv) to (v) in Figure 4. However, if an element already exists after ( $D < E$ ) in the OpLog of the reader, we need to order based on the value of the version stamps. As we place the greater version stamp to the left of the smaller version stamp at the branching point, the InterleavedPreorder traversal places the greater version stamp to the left of the smaller one in IPL; giving us the exact same order as in RGA. Hence, the underlying algorithm of both RGA and IPL are the same.

### 5 Conclusion and Future Work

In this work, we showed how to order operations for generic replicated data types using version trees. The key to supporting generic data types is to represent them using append-only

logs, which facilitates rollback of operations required while achieving a consistent order. We also explore the similarities between our approach and an existing approach to maintain order in sequence or list CRDT. This implies that we can potentially use versioning schemes while ordering operations, which can provide multiple benefits to the system including facilitating debugging, efficiently executing temporal query, and supporting data lineage. In the future, we plan to extend this work by creating a portfolio of versioned data types that can be represented using append-only logs.

Although we can guarantee convergence of an arbitrary data structure by ordering operations, convergence itself does not mean the user expectations on the semantics are satisfied. As an example, the consistency of a conventional set can be achieved by consistent ordering of operations. However, the consistency for the more constrained OR-set following an add-wins strategy might violate the expected semantics as a delete can supersede a concurrent add if the former is ordered after the latter following the strategy described in Section 3.2 in the future, we plan to devise mechanisms to respect the expectations on the semantics of the underlying data structure as well.

### Acknowledgments

This work has been supported in part by NSF awards CNS-2107101, CNS-1703560, ACI-1541215.

### References

- [1] Paulo Sérgio Almeida, Ali Shoker, and Carlos Baquero. 2020. Data state replicated data types. *Parallel and Distrib. Comput.* 111 (2018), 162–173.
- [2] Hagit Attiya, Sebastian Burckhardt, Alexey Gotsman, Adam Morrison, Hongseok Yang, and Marek Zawirski. 2016. Specification and complexity of collaborative text editing. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*. 259–268.
- [3] Kevin De Porre, Florian Myter, Christophe De Troyer, Christophe Scholliers, Wolfgang De Meuter, and Elisa Gonzalez Boix. 2019. Putting order in strong eventual consistency. In *IFIP International Conference on Distributed Applications and Interoperable Systems*. Springer, 36–56.
- [4] J. Driscoll, N. Sarnak, D. Sleator, and R. Tarjan. 1989. Making data structures persistent. *J. Comput. Syst. Sci.* 38, 1 (1989).
- [5] Victor Grishchenko. [n. d.]. Causal trees: towards real-time read-write hypertext.
- [6] Victor Grishchenko. 2010. Deep hypertext with embedded revision control implemented in regular expressions. In *Proceedings of the 6th International Symposium on Wikis and Open Collaboration*. 1–10.
- [7] Robert T Hood and Robert C Melville. 1988. Real time queue operations in pure Lisp. Technical Report. Cornell University.
- [8] Martin Kleppmann, Dominic P Mulligan, Victor BF Gomes, and Alastair R Beresford. 2020. A highly-available move operation for replicated trees and distributed systems.
- [9] Leslie Lamport. 2019. Time, clocks, and the ordering of events in a distributed system. In *Concurrency: the Works of Leslie Lamport*. 179–196.
- [10] Eugene W Myers. 1983. An applicative random-access stack transformation processing letters 17, 5 (1983), 241–248.

- [11] Brice Nédelec, Pascal Molli, Achour Mostefaoui, and Emmanuel Desmontils. 2013. LSEQ: an adaptive structure for sequences in distributed collaborative editing. In *Proceedings of the 2013 ACM symposium on Document engineering*. 37–46.
- [12] Diego Ongaro and John Ousterhout. 2014. In search of an understandable consensus algorithm. In *2014 USENIX Annual Technical Conference (USENIX ATC 14)*. 305–319.
- [13] Gérard Oster, Pascal Urso, Pascal Molli, and Abdessamad Imine. 2006. Data consistency for P2P collaborative editing. In *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*. 259–268.
- [14] Nuno Preguiça, Carlos Baquero, and Marc Shapiro. 2009. *Conflict-Free Replicated Data Types CRDTs*. Springer International Publishing, Cham. 491–500. [https://doi.org/10.1007/978-3-319-77525-8\\_185](https://doi.org/10.1007/978-3-319-77525-8_185)
- [15] Nuno Preguiça, Joan Manuel Marques, Marc Shapiro, and Mihai Letia. 2009. A commutative replicated data type for cooperative editing. In *2009 29th IEEE International Conference on Distributed Computing Systems*. IEEE, 395–403.
- [16] Thomas Reps, Tim Teitelbaum, and Alan Demers. 1983. Segmental context-dependent analysis for language-based editors. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 5, 3 (1983), 449–477.
- [17] Hyun-Gul Roh, Myeongjae Jeon, Jin-Soo Kim, and Joonwon Lee. 2011. Replicated abstract data types: Building blocks for collaborative applications. *J. Parallel and Distrib. Comput.* 71, 3 (2011), 354–368.
- [18] Nazmus Saquib, Chandra Krintz, and Rich Wolski. 2021. RDALS: Persisting Versioned Data Structures. In *2021 IEEE International Conference on Cloud Engineering (IC2E)*. IEEE, 179–190.
- [19] Neil Sarnak and Robert E Tarjan. 1983. Planar point location using persistent search trees. *Commun. ACM* 29, 7 (1986), 669–679.
- [20] Marc Shapiro, Nuno Preguiça, Carlos Baquero, and Marek Zawirski. 2011. *A comprehensive study of convergent and commutative replicated data types*. Technical Report RR-7506. Inria.
- [21] Marc Shapiro, Nuno Preguiça, Carlos Baquero, and Marek Zawirski. 2011. Conflict-free replicated data types. In *Symposium on Self-Stabilizing Systems*. Springer, 386–400.
- [22] Garret Frederick Swart. 1986. EFFICIENT ALGORITHMS FOR COMPUTING GEOMETRIC INTERSECTIONS (DECISION TREE, HIDDEN LINE REMOVAL, GRAPHICS, COMPLEXITY). (1986).