



For-Each Operations in Collaborative Apps

Matthew Weidner
maweidne@andrew.cmu.edu
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Benito Geordie
bg31@rice.edu
Rice University
Houston, Texas, USA

Ria Pradeep
rpradeep@alumni.cmu.edu
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Heather Miller
heather.miller@cs.cmu.edu
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Abstract

Conflict-free Replicated Data Types (CRDTs) allow collaborative access to an app's data. We describe a novel CRDT operation, for-each on the list of CRDTs, and demonstrate its use in collaborative apps. Our for-each operation applies a given mutation to each element of a list, including elements inserted concurrently. This often preserves user intention in a way that would otherwise require custom CRDT algorithms. We give example applications of our for-each operation to collaborative rich-text, recipe, and slideshow editors.

CCS Concepts: • **Theory of computation** → **Distributed algorithms**; • **Human-centered computing** → **Collaborative and social computing systems and tools**.

Keywords: collaboration, CRDTs, concurrency

ACM Reference Format:

Matthew Weidner, Ria Pradeep, Benito Geordie, and Heather Miller. 2023. For-Each Operations in Collaborative Apps. In *10th Workshop on Principles and Practice of Consistency for Distributed Data (PaPoC '23)*, May 8, 2023, Rome, Italy. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3578358.3591323>

1 Introduction

Lists of mutable values are common in collaborative apps. Examples include the list of slides in a slideshow editor, or the list of rich characters (characters plus formatting attributes) in a rich-text editor.

To allow collaborative access to this data, we would like to use a Conflict-free Replicated Data Type (CRDT) [11, 14]. Let us assume that we already have a CRDT C representing the list's mutable value type. Then one can construct a *list of Cs* CRDT representing the entire list.



This work is licensed under a Creative Commons Attribution International 4.0 License.

PaPoC '23, May 8, 2023, Rome, Italy

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0086-6/23/05.

<https://doi.org/10.1145/3578358.3591323>

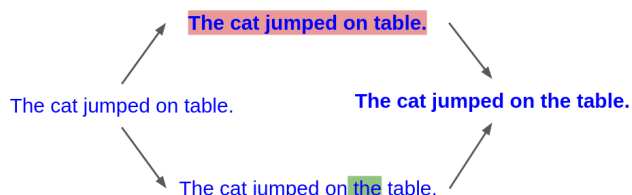


Figure 1. Typical user intention for a bold-range operation (top) concurrent to text insertion (bottom). When using a forEachPrior operation, "the" would not be bolded.

The basic operations on a list of C s allow users to insert, delete, and apply C operations to individual list elements. However, many user operations instead take the form of a "for-each" loop: for each element of the list meeting some condition, apply a C operation to that element. For example:

- In a slideshow editor, a user selects slides 3–7 and changes the background color to blue. This does: for each slide, if its index is in the range $[3, 7]$, then set its "background color" property to "blue".
- In a rich-text editor, a user selects all and clicks the "bold" formatting button. This does: for each rich character, set its "bold" formatting attribute to "true".

The easy way to implement such for-each operations is using a literal for-each loop on the initiating user's replica. That is, the user's device loops through its replica of the list and performs CRDT operations on individual elements. We call this a forEachPrior operation, since it acts on elements that were inserted (causally) prior to the for-each operation.

However, forEachPrior operations do not always capture user intention. Figure 1 shows a classic example: in a rich text editor, if one user bolds a range of text, while concurrently, another user types in the middle of the range, then the latter text should also be bolded. A literal for-each loop on the first user's replica will not do so because it is not aware of concurrently-inserted characters.

Traditionally, collaborative rich-text editors accomplish Figure 1's user intention using a specialized tree structure (e.g. Iagnet et al. [5]) or formatting markers at both ends of the range (e.g. Peritext [9]). However, both techniques require a




Figure 2. Light cone diagram for a forEach operation.

careful analysis of operations' interactions, and they do not generalize beyond rich-text editing.

1.1 Contributions

In this paper, we propose a novel for-each operation on lists of CRDTs. Unlike forEachPrior, it applies a given mutation to every list element that is inserted prior *or concurrently* to the for-each operation. We call this operation *for-each* (without qualification), to distinguish it from forEachPrior.

Figure 2 illustrates how an element's insert operation may relate to a for-each operation: causally prior, concurrent, or causally future. Our for-each operation affects the "prior" and "concurrent" categories, while forEachPrior only affects the "prior" category.

Using our for-each operation, we easily implement Figure 1's intended behavior: issue a for-each operation with the mutation "if the character is in the range, set its "bold" formatting attribute to "true". See Section 4.1 for details.

We hope that the intended semantics (i.e., user-visible behavior) of the list of C s and our for-each operation are already clear. However, some technicalities arise, especially when applying for-each to concurrently-inserted elements. Sections 2 and 3 discuss these technical details, including algorithms for all of our constructions.

A hurried reader may skip directly to Section 4, which applies for-each operations to example collaborative apps.

1.2 Background

We assume familiarity with the causal order on CRDT operations [8]. The terms "(causally) prior", "concurrent", and "(causally) future" reference this order. Our algorithms use vector clocks [4, 10] to query the causal order relationship between CRDT operations.

Throughout the paper, we use the language of operation-based CRDTs [14], although our constructions can easily be reformulated as state-based CRDTs. Each CRDT operation is described in terms of a *generator* and an *effector*. The generator is called to handle user input on the user's local replica, and it returns a *message* to be broadcast to other replicas. Each replica, including the sender, applies the operation by passing this message to the corresponding effector; the sender does so atomically with the generator call. We assume that messages are received exactly once on each replica, and in causal order.

2 List of CRDTs

We begin with a formal description of the *list of CRDTs*. It is modeled on Yjs's $Y.Array$ shared type [6].

First, a *list CRDT* is a classic CRDT type whose external interface is a list (ordered sequence) of immutable values, e.g., the characters in a text document [1, 14]. Since the same value may appear multiple times in a list, we use *element* to refer to a unique instance of a value. A list CRDT has operations to insert and delete elements.

A *list of CRDTs* is a more general CRDT in which the list values are themselves mutable CRDTs. Specifically, let C be an operation-based CRDT. The external interface of a *list of Cs* is a list of mutable values of type C . The operations on the list are:

- $\text{insert}(i, \sigma)$: Inserts a new element with initial value σ —a state of C —into the list at index i , between the existing elements at indices $i - 1$ and i . All later elements (index $\geq i$) shift to an incremented index.
- $\text{delete}(i)$: Deletes the element at index i . All later elements (index $\geq i + 1$) shift to a decremented index.
- $\text{apply}(i, o)$: Applies a C operation o to the element at index i . All replicas update their copy of the element's value (a state of C) in the usual way for C operations. A concurrent delete operation may cause a replica to receive the apply message after deleting the element; in this case, it ignores the apply message (the delete "wins").

Example 2.1. As a running example, consider a collaborative rich-text document, such as a Google Doc. We can represent one character using a *rich character CRDT*. Its state is a pair $(\text{char}, \text{attrs})$, where *char* is an immutable character and *attrs* is a map CRDT [14] for formatting attributes. Then a *list of rich character CRDTs* models the entire rich-text

document's state. E.g., the text “ab” is represented as

```
[{char : "a", attrs : {}},
 {char : "b", attrs : {"bold", true}}]
```

We construct the list of C s using C and an ordinary list CRDT \mathcal{L} . See Algorithm 1 for pseudocode.

Specifically, we assume that \mathcal{L} produces *positions* that are unique, immutable, and drawn from a dense total order $<$, e.g., Logoot's “position identifiers” [18].¹ Then the list of C s is implemented as:

State A list of elements (p, σ) , where p is a position from \mathcal{L} and σ is a state of C , sorted by p . An application using the list usually only looks at the values σ , but it may also use the positions, e.g., for cursor locations.

Insert, delete Similar to \mathcal{L} .

Apply Similar to C , except that the message sent to remote replicas is tagged with the element's position p . In the pseudocode, we use $C.\text{gen}(o, \sigma)$ to represent C 's generator for an operation o , and we use $C.\text{eff}(m, \sigma)$ to represent C 's effector for a message m .

We sketch a proof of strong eventual consistency in the extended version [17, Thm A.1].

3 For-Each Operation

We now define our new CRDT operation, *for-each*, on the list of C s.

Let O denote the set of all C operations. For technical reasons (described in Section 3.1 below), we restrict for-each to pure operations, where an operation is *pure* if its generated message is just the operation itself. Formally:

Definition 3.1. Let $C.\text{gen}(o, \sigma)$ denote C 's generator. An operation $o \in O$ is *pure* if $C.\text{gen}(o, \sigma) = o$ for all states σ .

Baquero et al. [2] show that many classic CRDTs' operations are pure, at least with the relaxations discussed in Section 3.1.

Let $O_P \subset O$ denote the subset of pure C operations. Let

$$f : (p, \text{prior}) \rightarrow O_P \cup \{\text{del}, \text{null}\}$$

be a function that takes as input a list element's position p and a boolean *prior* described below, and returns one of:

- $o \in O_P$: a pure C operation to apply to the element.
- **del**: an instruction to delete the element.
- **null**: an instruction to do nothing.

Then the operation `forEach(f)` loops over $elts$, applies f to each element, then performs the operation specified by f . Specifically, it loops over all elements that are inserted causally prior or concurrently to the for-each operation itself, but not causally future elements. It also computes the argument *prior* for f , which indicates whether each element is causally prior (true) or concurrent (false).

¹If \mathcal{L} uses extra state (e.g., tombstones) or messages to manage its positions, then those are implicitly added to the state or messages for the list of C s.

1 per-replica CRDT state:

2 $elts$: A list of elements (p, σ) , where p is a position from \mathcal{L} and σ is a state of C , sorted by p

3 **query** elements()

4 **return** $elts$

5 **update** insert

6 **generator** (i, σ)

7 $p \leftarrow$ new \mathcal{L} position between the positions at indices $i - 1$ and i in $elts$

8 **return** (insert, p, σ)

9 **effector** (insert, p, σ)

10 Insert (p, σ) into $elts$

11 **update** delete

12 **generator** (i)

13 $elt \leftarrow$ i -th element in $elts$

14 **return** (delete, $elt.p$)

15 **effector** (delete, p)

16 $elt \leftarrow$ unique element of $elts$ s.t. $elt.p = p$, or null if none exists (already deleted)

17 **if** $elt \neq \text{null}$ **then**

18 Delete elt from $elts$

19 **update** apply

20 **generator** (i, o)

21 $elt \leftarrow$ i -th element in $elts$

22 $m \leftarrow C.\text{gen}(o, elt.\sigma)$

23 **return** (apply, $elt.p, m$)

24 **effector** (apply, p, m)

25 $elt \leftarrow$ unique element of $elts$ s.t. $elt.p = p$, or null if none exists (already deleted)

26 **if** $elt \neq \text{null}$ **then**

27 $elt.\sigma \leftarrow C.\text{eff}(m, elt.\sigma)$

Algorithm 1: List of C s as an operation-based CRDT.

Example 3.2. In a rich text document, a user bolds a range of text. Let *start* and *end* be the positions of the first and last-plus-1 characters in the range, so that the range is $[start, end)$. Define:

function $f(p, \text{prior})$

if $start \leq p < end$ **then**

return ($rich \mapsto rich.attrs.set(\text{"bold"}, \text{true})$)

else return null

Then `forEach(f)` implements the intended behavior in Figure 1: all characters in the range are bolded, including those inserted concurrently.

Example 3.3. Again in a rich-text document, a user deletes a range of text $[start, end)$. To delete only existing characters, we consult *prior*:




Figure 3. Typical user intention for a delete-range operation (top) concurrent to text insertion (bottom): the concurrent text is not deleted, to avoid data loss.

```

function  $f(p, prior)$ 
  if  $prior$  and  $start \leq p < end$  then
    return del
  else return null

```

Then $forEach(f)$ implements the behavior shown in Figure 3. Note that we could instead use a $forEachPrior$ operation, i.e., an ordinary loop on the initiating replica. However, $forEach(f)$ generates less network traffic: a single $forEach$ message for the entire range, instead of a separate delete message per deleted character.

Algorithm 2 gives a pseudocode implementation of $forEach$, which we now describe.

We first modify the list of C s to track each element’s logical insertion time t —namely, its sender’s vector clock entry.² We also add a list $buffer$ to the internal state.

When a user calls $forEach(f)$, their replica broadcasts f together with the operation’s vector clock w .

Upon receiving this message, a replica first loops over its current elements. For each element elt , the replica computes $f(elt.p, prior)$ and does as instructed, but only locally; it does not broadcast any new messages. Here $prior$ indicates whether elt was inserted causally prior to $forEach(f)$. Note that we do not apply f to elements that were already deleted on this replica, including by concurrent delete operations (the delete “wins”).

Next, the receiving replica stores the message in its $buffer$. In the future, whenever the replica receives an insert message, it checks whether the insert operation is concurrent to $forEach(f)$. If so, the replica computes $f(elt.p, false)$ and does as instructed, again only locally. Note that the message stays in the buffer forever, although in principle it could be discarded once all concurrent operations are received (i.e., it is causally stable [2]).

3.1 On Pure Operations

Our restriction to pure operations is not arbitrary: we need to know what message to pass to $C.eff$, even for concurrent elements. Such elements did not yet exist on the initiating replica, hence the replica could not pass their states to $C.gen$.

²Some authors call this a *causal dot*.

1 per-replica CRDT state:

- 2 $elts$: A list of elements (p, σ, t) , where p is a position from \mathcal{L} , σ is a state of C , and $t = (senderID, clock)$ is a vector clock entry; sorted by p
- 3 $buffer$: A list of pairs (f, u) , where f is the $forEach$ argument and u is a vector clock entry
- 4 vc : the local vector clock, in the form of a function from replica IDs to \mathbb{N} ; initially the all-0 function
- 5 $replicaID$: the unique ID of this replica

6 **function** $execute(f, elt, prior)$

- 7 $op \leftarrow f(elt.p, prior)$
- 8 **if** $op \in O$ **then**
- 9 $elt.\sigma \leftarrow C.eff(op, elt.\sigma)$
- 10 **else if** $op = del$ **then**
- 11 Delete elt from $elts$

12 **update** insert

- 13 **generator** (i, σ)
- 14 $p \leftarrow$ new \mathcal{L} position between the positions at indices $i - 1$ and i in $elts$
- 15 $v \leftarrow$ copy of vc
- 16 $v[replicaID] \leftarrow v[replicaID] + 1$
- 17 **return** (insert, $p, \sigma, v, replicaID$)

18 **effector** (insert, $p, \sigma, v, senderID$)

- 19 $vc[senderID] = v[senderID]$
- 20 $t \leftarrow (senderID, v[senderID])$
- 21 Insert (p, σ, t) into $elts$
- 22 // Loop over concurrent for-each operations.
- 23 **for** (f, u) **in** $buffer$ **do**
- 24 $concurrent \leftarrow (v[u.senderID] < u.clock)$
- 25 **if** $concurrent$ **then** $execute(f, elt, false)$

26 **update** $forEach$

- 27 **generator** (f)
- 28 $w \leftarrow$ copy of vc
- 29 $w[replicaID] \leftarrow w[replicaID] + 1$
- 30 **return** (forEach, $f, w, replicaID$)
- 31 **effector** (forEach, $f, w, senderID$)
- 32 $vc[senderID] = w[senderID]$
- 33 **for** elt **in** $elts$ **do**
- 34 $prior \leftarrow (w[elt.t.senderID] \geq elt.t.clock)$
- 35 $execute(f, elt, prior)$
- 36 $u \leftarrow (senderID, w[senderID])$
- 37 Append (f, u) to $buffer$

Algorithm 2: List of C s with our for-each operation. Blocks not shown here are the same as in Algorithm 1 (elements, delete, and apply).

With pure operations, we know that the generated message is just o itself, as used on line 9.

In practice, we can relax the pure restriction by passing additional metadata to $C.\text{eff}$. In particular, we may pass in f 's vector clock: line 9 of `execute` becomes

$$\text{elt}.\sigma \leftarrow C.\text{eff}((op, w), \text{elt}.\sigma).$$

This does not threaten strong eventual consistency because w is consistent across replicas.

C can use the provided vector clocks to query the causal order on operations. That is sufficient to implement most CRDTs using only pure operations [2]. List CRDTs' insert operations are a notable exception.

3.2 Correctness

Informally, we claim that Algorithm 2 matches the semantics described in the introduction. That is, a for-each operation's f is applied to exactly the causally prior and concurrent elements, minus deleted elements, regardless of message order.

We give a precise correctness claim and proof sketch in the extended version [17, Thm A.2].

3.3 Other Data Structures

For-each works equally well if we ignore the list order but still assign a unique ID p to each element, like in a U-set CRDT [14]. That is, we can define a for-each operation on a *U-set of CRDTs* in which each added element is assigned a unique ID.

Likewise, one can define for-each on a CRDT-valued map in which each key-value pair is assigned a unique ID when set, like Yjs's `Y.Map` shared type [6].

However, our construction does not work with a Riak-style map [3] in which a key's value CRDT is created on first use instead of explicitly set: two users may create the same key's value CRDT concurrently, complicating the choice of which for-each operations to apply [15, §4]. We expect similar issues for the list of CRDTs in Kleppmann and Beresford's JSON CRDT [7], in which an element may reappear after deletion.

We leave full descriptions to future work.

4 Examples

We now describe example uses of our for-each operation in collaborative apps, at a high level. As in Section 3, we write a for-each operation as `forEach(f)`, where

$$f : (p, \text{prior}) \rightarrow O_P \cup \{\text{del}, \text{null}\}$$

is a function that takes as input a list element's position p and whether it is causally prior (else concurrent), and outputs an instruction for that element: apply a (pure) operation $o \in O_P$, delete the element, or do nothing.

4.1 Rich-Text Editor

Let us begin with a collaborative rich-text editor, as described in the introduction. To recap Examples 2.1 and 3.2, we can

represent a rich-text document as a list of *rich character CRDTs* ($\text{char}, \text{attrs}$), where char is an immutable character and attrs is a map CRDT for formatting attributes. Given list positions start and end , define:

```
function  $f(p, \text{prior})$ 
  if  $\text{start} \leq p < \text{end}$  then
    return ( $\text{rich} \mapsto \text{rich}.\text{attrs}.\text{set}(\text{"bold"}, \text{true})$ )
  else return null
```

Then `forEach(f)` bolds the range $[\text{start}, \text{end})$ with the intended behavior in Figure 1: all characters in the range are bolded, including concurrently-inserted ones.

It is possible to use a closed interval $[\text{start}, \text{end}']$ instead of the half-open interval $[\text{start}, \text{end})$. Here end' is the position of the last character in the original range, while end is the last-plus-one position. The difference is that $[\text{start}, \text{end})$ will also format concurrently-inserted characters at the end of the range, while $[\text{start}, \text{end}']$ will not. The latter behavior is typical for hyperlink formatting [9].

Other formatting attributes are similar. However, for deletions, one typically deletes only causally prior characters, as in Example 3.3. This is safer because deletions are monotonic (permanent), making unintended deletions harder to undo.

Note that a literal list of rich character CRDTs is memory-inefficient, since it stores a map CRDT per character. However, one can use this theoretical model as a guide, then implement an equivalent but more efficient CRDT. For example, one can store attrs 's state explicitly only when it differs from the previous character, like in Peritext [9].

4.2 Recipe Editor

A collaborative recipe editor allows multiple users to view and edit a recipe for a meal. Let us consider in particular the list of ingredients. We can model it as a list of *ingredient CRDTs*, where each ingredient CRDT has sub-CRDTs for its name and amount.

Suppose we add a "scale recipe" button that multiplies every amount by a given value. If one user scales the recipe, while concurrently, another user inserts a new ingredient, then it is important that the new ingredient's amount is also scaled. Otherwise, it will be out of proportion with the other ingredients.

To implement such a "scale recipe" operation, let s be the scaling amount. Define f by:

```
function  $f(p, \text{prior})$ 
  return ( $\text{ingredient} \rightarrow \text{ingredient}.\text{amount}.\text{mult}(s)$ )
```

Then `forEach(f)` scales every ingredient's amount, including ingredients inserted concurrently.³

³Here we assume an operation `mult(s)` on the "ingredient amount" CRDT. This is nontrivial if you also allow `set(value)` operations, but it can be implemented using another list with for-each operations; we omit the details.




Figure 4. Typical user intention for a translation operation (top) concurrent to a group rotation (bottom): the rectangle is translated within the rotated group.

4.3 Slideshow Editor

A slideshow editor is another collaborative app that can use for-each operations. A single slide might contain multiple images, shapes, or text boxes. These objects can be edited individually or together. For example, a user might translate (shift) a single object while another simultaneously rotates all objects on the slide.

To implement these translations and rotations, each translation on an object can be represented as a translation vector. Then the object's position is represented by a list CRDT \mathcal{L} containing all translations made so far; the actual position is the sum of all the vectors in that list.

The entire slide can be represented as a list CRDT \mathcal{L}' , where each element is a list CRDT \mathcal{L}_p of an object p 's translation vectors.⁴

A user shifts an object p by appending a translation vector to its list \mathcal{L}_p , and rotates an object p by multiplying corresponding translation vectors in \mathcal{L}_p by a rotation matrix. When a group of objects are edited, the same operation should be applied to each object's list.

For example, say a user rotates a group of objects 30 degrees clockwise, like the bottom operation of Figure 4. We want to rotate each object in the group. To keep objects aligned within the group, we also want to rotate any concurrent translation of those objects, such as the top translation of Figure 4.

To implement this, let *updatedObjects* be a set of the positions p of the selected objects. Define g and f as:

```

function  $g(q, \text{prior})$ 
  return  $\left( \text{vector} \mapsto \text{vector}.\text{mult} \left( \begin{bmatrix} \cos 30 & \sin 30 \\ -\sin 30 & \cos 30 \end{bmatrix} \right) \right)$ 

function  $f(p, \text{prior})$ 
  if  $p \in \text{updatedObjects}$  then
    return  $(\text{object} \mapsto \text{object}.\text{forEach}(g))$ 
  else return null
```

Then $\text{forEach}(f)$ rotates all *updatedObjects* 30 degrees clockwise. The degree of rotation can also be stored to render the rotated object correctly.

⁴Since the objects on a slide are unordered, we use their positions p merely as IDs, ignoring their total order. See Section 3.3.

5 Related Work

Dataflow programming and stream processing both perform operations “for each” element of a stream. Unlike this work, they typically apply a for-each operation to all regions in Figure 2, including the causal future. In particular, FlowPools [12] allow issuing a for-each operation after a FlowPool (stream) begins; they apply the operation to all existing elements immediately, then store it as a callback for concurrent or future elements, similar to our algorithm (Algorithm 2).

Operational Transformation [13] allows every operation on a collaborative app to perform a transformation for each concurrent operation. In contrast, we allow for-each operations to transform list elements (equivalently, insert operations) but not each other. Thus we do not need complicated algebraic rules to ensure eventual consistency.

The semidirect product of CRDTs [15] combines the operations of two CRDTs, C_1 and C_2 , in a single CRDT. It essentially implements the rule: to apply a C_2 operation, “act on” each prior and concurrent C_1 operation in some way, then reduce over those C_1 operations to get the current state. However, instead of storing the literal list of C_1 operations, it only stores their reduced form (the actual state). Thus one can view the semidirect product as an optimized but less intuitive version of our list with for-each operations.

6 Conclusions

We formalized the list of CRDTs and described a novel for-each operation on this list. The resulting CRDT models a list of mutable values in a collaborative app, equipped with the operation: for each element of the list, including ones inserted concurrently, apply some operation to that element. We gave several examples in which our for-each operation matches user intention better than a literal for-each loop.

For future work, we plan to implement our for-each operation in the Collabs CRDT library [16].

Acknowledgments

We thank James Riely for insightful questions about a previous paper [15] that inspired this work. We also thank the anonymous PaPoC reviewers for helpful feedback. Matthew Weidner was supported by an NDSEG Fellowship sponsored by the US Office of Naval Research. Benito Geordie was supported by an REU sponsored by the US National Science Foundation.

References

- [1] 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* (Chicago, Illinois, USA) (PODC '16). Association for Computing Machinery, New York, NY, USA, 259–268. <https://doi.org/10.1145/2933057.2933090>
- [2] Carlos Baquero, Paulo Sérgio Almeida, and Ali Shoker. 2014. Making Operation-Based CRDTs Operation-Based. In *Distributed Applications*

- and Interoperable Systems (Lecture Notes in Computer Science), Kostas Magoutis and Peter Pietzuch (Eds.). Springer, Berlin, Heidelberg, 126–140. https://doi.org/10.1007/978-3-662-43352-2_11
- [3] Basho. 2015. Riak datatypes. <http://github.com/basho>.
- [4] Colin J. Fidge. 1988. Timestamps in message-passing systems that preserve the partial ordering. *Proceedings of the 11th Australian Computer Science Conference* 10, 1 (1988), 56–66.
- [5] Claudia-Lavinia Ignat, Luc André, and Gérald Oster. 2021. Enhancing rich content wikis with real-time collaboration. *Concurrency and Computation: Practice and Experience* 33, 8 (2021), e4110. <https://doi.org/10.1002/cpe.4110> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/cpe.4110> e4110 cpe.4110.
- [6] Kevin Jahns. 2022. Yjs. GitHub repository. <https://github.com/yjs/yjs>
- [7] Martin Kleppmann and Alastair R Beresford. 2017. A Conflict-Free Replicated JSON Datatype. *IEEE Transactions on Parallel and Distributed Systems* 28, 10 (April 2017), 2733–2746. <https://doi.org/10.1109/TPDS.2017.2697382> arXiv:1608.03960
- [8] Leslie Lamport. 1978. Time, Clocks, and the Ordering of Events in a Distributed System. *Commun. ACM* 21, 7 (jul 1978), 558–565. <https://doi.org/10.1145/359545.359563>
- [9] Geoffrey Litt, Sarah Lim, Martin Kleppmann, and Peter van Hardenberg. 2022. Peritext: A CRDT for Collaborative Rich Text Editing. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 531 (nov 2022), 36 pages. <https://doi.org/10.1145/3555644>
- [10] Friedemann Mattern. 1989. Virtual Time and Global States of Distributed Systems. In *Parallel and Distributed Algorithms*. North-Holland, 215–226.
- [11] Nuno Preguiça, Carlos Baquero, and Marc Shapiro. 2018. *Conflict-Free Replicated Data Types CRDTs*. Springer International Publishing, Cham, 1–10. https://doi.org/10.1007/978-3-319-63962-8_185-1
- [12] Aleksandar Prokopec, Heather Miller, Tobias Schlatter, Philipp Haller, and Martin Odersky. 2013. FlowPools: A Lock-Free Deterministic Concurrent Dataflow Abstraction. In *Languages and Compilers for Parallel Computing*, Hironori Kasahara and Keiji Kimura (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 158–173.
- [13] Matthias Ressel, Doris Nitsche-Ruhland, and Rul Gunzenhäuser. 1996. An Integrating, Transformation-Oriented Approach to Concurrency Control and Undo in Group Editors. In *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work (Boston, Massachusetts, USA) (CSCW '96)*. Association for Computing Machinery, New York, NY, USA, 288–297. <https://doi.org/10.1145/240080.240305>
- [14] Marc Shapiro, Nuno Preguiça, Carlos Baquero, and Marek Zawirski. 2011. Conflict-Free Replicated Data Types. In *13th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS 2011)*. 386–400. https://doi.org/10.1007/978-3-642-24550-3_29
- [15] Matthew Weidner, Heather Miller, and Christopher Meiklejohn. 2020. Composing and Decomposing Op-Based CRDTs with Semidirect Products. *Proc. ACM Program. Lang.* 4, ICFP, Article 94 (aug 2020), 27 pages. <https://doi.org/10.1145/3408976>
- [16] Matthew Weidner, Heather Miller, Huairui Qi, Maxime Kjaer, Ria Pradeep, Ignacio Maronna, Benito Geordie, and Yicheng Zhang. 2022. Collabs. GitHub repository. <https://github.com/composablesys/collabs>
- [17] Matthew Weidner, Ria Pradeep, Benito Geordie, and Heather Miller. 2023. For-Each Operations in Collaborative Apps. arXiv:2304.03141 [cs.DC]
- [18] Stephane Weiss, Pascal Urso, and Pascal Molli. 2009. Logoot: A Scalable Optimistic Replication Algorithm for Collaborative Editing on P2P Networks. In *2009 29th IEEE International Conference on Distributed Computing Systems*. 404–412. <https://doi.org/10.1109/ICDCS.2009.75>