OpSets: Sequential Specifications for Replicated Datatypes (Extended Version)

Martin Kleppmann

Computer Laboratory, University of Cambridge, UK
mk428@cl.cam.ac.uk
 https://orcid.org/0000-0001-7252-6958

Victor B. F. Gomes

Computer Laboratory, University of Cambridge, UK
vb358@cl.cam.ac.uk
 https://orcid.org/0000-0002-2954-4648

Dominic P. Mulligan

Security Research Group, Arm Research, Cambridge, UK
Dominic.Mulligan@arm.com
 https://orcid.org/0000-0003-4643-3541

Alastair R. Beresford

Computer Laboratory, University of Cambridge, UK
arb33@cl.cam.ac.uk
 https://orcid.org/0000-0003-0818-6535

—— **Abstract** ——

We introduce OpSets, an executable framework for specifying and reasoning about the semantics of replicated datatypes that provide eventual consistency in a distributed system, and for mechanically verifying algorithms that implement these datatypes. Our approach is simple but expressive, allowing us to succinctly specify a variety of abstract datatypes, including maps, sets, lists, text, graphs, trees, and registers. Our datatypes are also composable, enabling the construction of complex data structures. To demonstrate the utility of OpSets for analysing replication algorithms, we highlight an important correctness property for collaborative text editing that has traditionally been overlooked; algorithms that do not satisfy this property can exhibit awkward interleaving of text. We use OpSets to specify this correctness property and prove that although one existing replication algorithm satisfies this property, several other published algorithms do not. We also show how OpSets can be used to develop new replicated datatypes: we provide a simple specification of an atomic move operation for trees, an operation that had previously been thought to be impossible to implement without locking. We use the Isabelle/HOL proof assistant to formalise the OpSets approach and produce mechanised proofs of correctness of the main claims in this paper, thereby eliminating the ambiguity of previous informal approaches, and ruling out reasoning errors that could occur in handwritten proofs.

1 Introduction

A common requirement across many distributed systems is that several nodes may concurrently access and manipulate some shared data structure. Examples include everything from journalists using their laptops to work on a shared text document to a set of web servers manipulating a common database. In doing so, it is important that the shared data satisfies certain *consistency guarantees*. For example, *strong consistency models* such as serializability [31] or linearizability [26] make a system behave like a single sequentially executing node, even when it is in fact replicated and concurrent. An unavoidable downside of these models is that any operation or transaction must wait for network communication before it is allowed to complete [13, 19]. Thus, in a system with strong consistency, a node cannot make progress while it is offline or partitioned from other nodes.

On the other hand, *eventual consistency* [5, 9, 56, 60] allows each participant to modify a local copy (*replica*) of a shared data structure while offline, but its definition is very weak: *"if no new updates are made to the shared state, all nodes will eventually have the same data."* The premise *"if no new updates are made"* may never be true if the shared state is continually modified (i.e. the system is never quiescent). Moreover, nothing in the definition of eventual consistency specifies which final states are legal.

Conflict-free Replicated Data Types, or CRDTs [51, 52], are abstractions for replicated state that have received significant attention in recent years (see § 6). The primary correctness property for CRDTs is *convergence* [52, 21], defined as: *"whenever any two replicas have applied the same set of updates, they are in the same state"*, even if each replica applies the updates in a different order. Convergence is a stronger property than eventual consistency, but it also fails to define what exactly the converged state should be.

In this work we introduce *Operation Sets* (or *OpSets* for short), a novel approach for specifying the semantics of replicated datatypes, and for reasoning about algorithms for concurrent data access and manipulation. We go beyond merely ensuring replica convergence: the OpSets approach is an executable specification that precisely defines the permitted states of a replica after some set of updates has been applied. Our contributions in this paper are as follows:

- In § 2 we introduce the OpSet, which provides a simple abstraction for specifying and reasoning about the consistency properties of concurrently editable data structures.
- On top of this abstraction, in § 3 and § 5, we specify a variety of composable abstract datatypes (maps, sets, lists, text, graphs, trees, and registers), and we argue that our specifications are both simple and precise, making them a suitable tool for reasoning about replicated data.
- In § 4 we demonstrate how the OpSet abstraction can be used to reason about existing algorithms. We highlight an important correctness property for collaborative text editing that has been overlooked by prior work in this area. Our specification is, to our knowledge, the first that correctly captures this property. We then review a selection of text editing CRDTs from the literature, prove that one satisfies our specification, and identify several others that fail to satisfy our correctness property.
- In § 5 we show how the OpSet abstraction can be used to develop new replicated datatypes. In particular we describe, for the first time, how an atomic move operation can be defined for a tree CRDT. This operation can be used to move a subtree to a new position within the tree, or to rename a key in a map, or to reorder items in a list. The OpSets approach enables a simple definition of this operation that had previously been thought impossible to implement without locking [40, 41].

▬ Using the Isabelle/HOL proof assistant [64] we formalise the OpSets approach, producing mechanised proofs of correctness of the main claims in this paper. In particular, we prove that our list specification is strictly stronger than the recent specification of collaborative text editing by Attiya et al. [4]. By using mechanised proofs we eliminate the ambiguity of previous informal approaches, and rule out reasoning errors that could occur in handwritten proofs. Moreover, the proof framework we have developed is reusable and can be leveraged to verify other datatypes in the future.

The appendices contain an overview of our Isabelle/HOL mechanisation, and pseudocode for the replicated datatypes discussed in this paper. The full formal proof development is published in the Isabelle Archive of Formal Proofs [33].

2   The OpSets Approach

The OpSets approach is a simple abstraction for describing the consistency properties of a replicated data system. We outline the general approach in this section, before describing concrete data structures and specifications in § 3 and § 5.

2.1   System Model

We assume that the system consists of a set of *nodes* connected by a *network*. These nodes concurrently access some *shared data structure*, which may be a relational database (consisting of rows in tables), a text document (a sequence of characters), a vector-graphics document (a tree of records describing graphical objects), a filesystem (a tree of directories and files), or any other kind of data structure.

New nodes can be added at any time, and the set of nodes need not be known in advance. Nodes might be mobile devices, and hence we assume that nodes are sometimes *offline*, i.e. temporarily unable to communicate with other nodes. We require that nodes can access the shared data anytime, even while offline. Thus, each node has a local copy of the shared data structure, which it can read and modify without waiting for any communication or coordination with other nodes.

Whenever a node makes a modification to that structure, it records the change as an *operation*. For example, an operation may describe an insertion at a particular position in a text document. Each node locally maintains a set of operations, the *OpSet*. Whenever a node makes a change to the shared data, it adds the corresponding operation to its OpSet, and also sends *messages* containing the operation to other nodes. Whenever a node receives a message from another node, the operation in that message is added to the recipient's local OpSet. Operations remain immutable throughout this process.

We make no assumptions about the reliability of the network: messages may be lost, duplicated, or arbitrarily reordered. Reflecting the characteristics of real networks, we assume that lost messages are retransmitted when possible (e.g. using TCP), but messages may be permanently lost due to network or node failures. Since the OpSet at each node is a monotonically growing set of operations, any two communicating nodes can merge their OpSets using the standard set union operator ∪. Set union is commutative, associative, and idempotent, ensuring that communicating nodes converge towards the same OpSet contents.

We assume that each operation has a unique identifier (ID), that new IDs can be generated by any node without communication with other nodes, and that we have a total ordering on operation IDs. These requirements can easily be met by using Lamport timestamps [35] as IDs. A Lamport timestamp is a pair (*counter*, *nodeID*) that is constructed as follows:

- *counter* is an integer. To generate a new ID, find the maximum counter of any existing operation ID in the local OpSet, and increment that number.
- *nodeID* is a string that uniquely identifies the node generating the ID, e.g. a UUID [36].

Although different nodes may generate IDs with the same counter value, each node generates IDs with strictly monotonically increasing counter values, and thus IDs are globally unique. We define the total order on IDs as being the lexicographic order:

$$(ctr_1, node_1) < (ctr_2, node_2) \iff ctr_1 < ctr_2 \ \lor \ (ctr_1 = ctr_2 \land node_1 < node_2).$$

2.2   Interpreting an OpSet

Most definitions of operation-based CRDTs describe how a node's local state is manipulated by operations [51, 52]. We now depart from this convention and present an alternative formulation of replicated datatypes.

In the OpSets approach, we require that the shared data structure is never manipulated directly. Instead, we use an *interpretation function* $[\![-]\!]$ that takes an OpSet O and returns the current state $[\![O]\!]$ of the shared data structure described by the OpSet. The interpretation function is *pure*, i.e. deterministic, side-effect free, and its result depends only on O. All nodes in the system employ the same interpretation function.

Consequently, whenever any two nodes have the same OpSet O, their view of the shared data structure $[\![O]\!]$ must also be equal. This construction trivially ensures eventual consistency: as two nodes converge towards the same OpSet contents, any data structure that is deterministically derived from the OpSet must also converge.

In principle, any deterministic function can serve as interpretation function. However, in defining the semantics of CRDTs (see § 3 and § 5), we have found it useful to specialise $[\![-]\!]$ such that we can interpret one operation at a time.

Let the OpSet O be a set of pairs (id, op), where id is a unique operation identifier and op is an arbitrary description of the change that occurred. Assume that we have a total ordering $<$ on identifiers, as explained in § 2.1. Then observe that for any OpSet there exists a unique sequence of operations, containing all operations of the OpSet in ascending order of their identifier. We can specify the semantics of each operation — that is, the effect of the operation on the OpSet interpretation — when applied in this sequential order.

Formally, we can define the interpretation $[\![O]\!]$ of the OpSet O as follows:

$$[\![\emptyset]\!] = \mathsf{InitialState}$$
$$[\![O \ \cup \ \{(id, \ op)\}]\!] = \mathsf{interp}\big[[\![O]\!], \ (id, \ op)\big] \qquad \text{provided that } \forall \, (id', \ op') \in O. \ id' < id$$

where $\mathsf{interp}\big[S, \ (id, \ op)\big]$ is the interpretation of the operation $(id, \ op)$ in the state S, and $\mathsf{InitialState}$ is a fixed minimal element (e.g. the empty tree, or empty list) of the replicated type described. In other words, if S is the result of interpreting all operations with identifiers less than id, then $\mathsf{interp}\big[S, \ (id, \ op)\big]$ is the interpretation of the OpSet to which $(id, \ op)$ has been added. For example, if $id_1 < id_2 < id_3$, we have:

$$[\![\{(id_1, \ op_1), (id_2, \ op_2), (id_3, \ op_3)\}]\!] \ =$$
$$\mathsf{interp}\big[\mathsf{interp}\big[\mathsf{interp}\big[\mathsf{InitialState}, \ (id_1, \ op_1)\big], \ (id_2, \ op_2)\big], \ (id_3, \ op_3)\big]$$

Provided that the operation interpretation $\mathsf{interp}\big[S, \ (id, \ op)\big]$ is deterministic, the OpSet interpretation function $[\![-]\!]$ is also deterministic, due to the fact that the operation order in the OpSet is unique.

2.3   Receiving Messages Out-of-order

Many computing systems are based on the idea of putting operations in some total order, and executing them in that order. For example, serializable transactions [31] and state machine replication [50] follow this approach. However, it is important to understand that the OpSet interpretation of § 2.2 relies on a weaker notion of ordering than most systems.

With serializable transactions and state machine replication, once a transaction/operation has been executed in some state, its results are expected to be durable. Thus, before executing some transaction T_i, the system needs to ensure that there is no pending transaction with a lower ID than T_i (which would need to be executed before T_i), since otherwise the subsequent arrival of a transaction with lower ID would invalidate the state in which T_i was executed. However, ensuring this precondition is expensive: as we show in § 6.1, it requires communication with at least a quorum of nodes; if the IDs are Lamport timestamps, it even requires communication with every single node [35]. If too many nodes are offline, the system cannot execute any transactions.

By contrast, our system model of § 2.1 requires nodes to always be able to read and modify the shared data, even when all nodes are offline. Moreover, we do not assume any ordering guarantees from the network. Thus, whenever there is some operation $(id_1, op_1) \in O$ in the OpSet O of some node, it is possible that the node will subsequently receive a message containing (id_2, op_2), where $id_2 < id_1$; that is, the later-arriving operation needs to be applied before the existing operation (id_1, op_1) in the OpSet interpretation $[\![O]\!]$.

In the OpSet model, such out-of-order delivery of operations is no problem: the order in which operations are received has no effect on the OpSet O, and since we assume the interpretation function to be pure and side-effect free, the interpretation $[\![O]\!]$ can always be recomputed whenever new operations are added to O.

The interpretation function is an *executable specification* that defines the expected result of interpreting a set of operations. Presenting replicated datatypes in this manner has two significant advantages:

1. Unlike typical definitions of CRDT algorithms [51, 52], it is not necessary for the interpretation function $\mathsf{interp}\big[S, (id, op)\big]$ to commute with respect to other operations: any pure function can be used. This fact makes it much simpler to specify the interpretation of operations, as we shall see in § 3 and § 5.
2. We can guarantee the existence of an implementation of each described datatype: the specification itself. This is in contrast to axiomatic specifications, which may not be implementable, and require additional work to demonstrate than an implementation exists which satisfies the axiomatic description.

For practical implementations of replicated datatypes, a naive OpSet interpretation may exhibit poor performance, since nodes must potentially apply the same subset of operations repeatedly. More efficient (and, most likely, more complex) algorithms for CRDTs can therefore be developed and shown to satisfy the OpSet-based specification—we do this in § 4.

However, we have developed a practical JavaScript CRDT implementation around the OpSet model [34], and found it to have some advantages: for example, users can easily inspect the editing history of a document, since every past version of the document is the interpretation of a particular subset of operations. Moreover, using OpSets provides a straightforward mechanism for recovering from network partitions and failures, as missing operations may be retransmitted and added to the OpSets of previously partitioned nodes. The details of this implementation are beyond the scope of this paper.

3 Specifying a Graph of Lists, Maps, and Registers

We now make the OpSets approach concrete by defining example semantics for commonly-used data structures: maps (which associate values with user-specified keys) and lists (linear sequences of values). The map datatype can also represent a set (by using keys as members of the set, and ignoring values). The list datatype can also represent text (by mapping each character to a list element). In both lists and maps the values may be primitives (such as numbers or strings), or references to other map or list objects. Using these references we can construct arbitrary object graphs, including cycles of object references, like in object-oriented programming languages. In § 5 we will show how to restrict this object graph so that it conforms to a tree structure.

We treat each key of a map, and each element of a list, as a multi-value register. That is, if there are several concurrent assignments to the same map key or list element, our datatype preserves all concurrently written values. Thus, reading a map key or list element may return multiple values, which may be merged explicitly by the user. Assigning a new value to a map key or list element overwrites all causally preceding values. Different register behaviour, such as last-writer-wins (arbitrarily picking one of the concurrently written values as winner), can easily be defined, as we show later.

3.1 Generating Operations

An OpSet for these datatypes may contain six types of operation:

- $(id, \mathsf{MakeMap})$ creates a new, empty map object that is identified by id.
- $(id, \mathsf{MakeList})$ creates a new, empty list object that is identified by id.
- $(id, \mathsf{MakeVal}(val))$ associates the ID id with the primitive value val (e.g. a number, string, or boolean). This operation is used to "wrap" any primitive value, allowing Assign operations (see below) to always use IDs as values, regardless of whether the value is a primitive value, or a reference to a map or list object.
- $(id, \mathsf{InsertAfter}(ref))$ creates a new list element with ID id, and inserts it into a list. If ref is the ID of a prior $\mathsf{MakeList}$ operation, then the new element is inserted at the head of that list. Otherwise ref must be the ID of an existing list element (i.e. a prior $\mathsf{InsertAfter}$ operation), in which case the new list element is inserted immediately after the referenced list element. Note that the $\mathsf{InsertAfter}$ operation does not associate a value with the new list element; that is done by a subsequent Assign operation.
- $(id, \mathsf{Assign}(obj, key, val, prev))$ assigns a new value to a key within a map (if obj is the ID of a prior $\mathsf{MakeMap}$ operation), or to a list element (if obj is the ID of a prior $\mathsf{MakeList}$ operation). In the case of map assignment, key is the user-specified key to be updated, which may be any primitive value such as a string or integer. In the case of a list, key is the ID of the list element to be updated (i.e. the ID of a prior $\mathsf{InsertAfter}$ operation). val is the ID of the value being assigned, which may identify a $\mathsf{MakeMap}$, $\mathsf{MakeList}$, or $\mathsf{MakeVal}$ operation. $prev$ is the set of IDs of prior Assign operations to the same key in the same object, which are overwritten by the present operation.
- $(id, \mathsf{Remove}(obj, key, prev))$ removes a key-value pair from a map, or an element from a list. As with Assign, obj is the ID of the prior $\mathsf{MakeMap}$ or $\mathsf{MakeList}$ operation that created the object being updated, and key identifies the key or list element being removed. $prev$ is the set of IDs of prior Assign operations to the same key in the same object, which are removed by the present operation.

Pseudocode for generating these operations is given in Appendix A.

3.2 Interpreting Operations

We use the sequential OpSet interpretation given in § 2.2. To encode the current state of map and list data structures we use a pair of relations (E, L):

The element relation $E \subseteq (\mathbf{ID} \times \mathbf{ID} \times (\mathbf{ID} \cup \mathbf{Key}) \times \mathbf{ID})$ is a set of 4-tuples containing the values currently assigned to map keys and list elements. If $(id, obj, key, val) \in E$, then an Assign operation with ID id updated the object with ID obj, assigning the value with ID val to the map key or list element key. If obj references a list object, key is the ID of an element in the list relation L (see below). If obj references a map object, any primitive value such as string or integer may be used as key.

The list relation $L \subseteq (\mathbf{ID} \times (\mathbf{ID} \cup \{\bot\}))$ is a set of pairs that indicates the order of list elements. If $(prev, next) \in L$, that means the list element with ID $prev$ is immediately followed by the list element with ID $next$. We use $(last, \bot) \in L$ to indicate that list element $last$ has no successor. To indicate that $head$ is the first element in the list obj (i.e. obj is the ID of the MakeList operation that created the list) we have $(obj, head) \in L$.
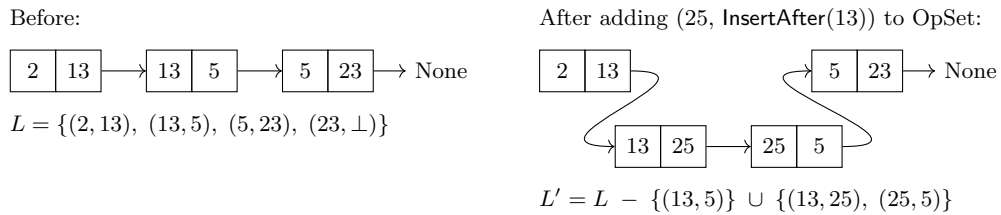
Initially, both relations are empty; that is, we have $\llbracket \emptyset \rrbracket = \mathsf{InitialState} = (\emptyset, \emptyset)$. We can then define the interpretation of the six operation types as follows:

$$\mathsf{interp}\big[(E, L), (id, \mathsf{Assign}(obj, key, val, prev))\big] =$$
$$\Big(\{(id', obj', key', val') \in E \mid id' \notin prev\} \cup \{(id, obj, key, val)\}, \ L\Big)$$

$$\mathsf{interp}\big[(E, L), (id, \mathsf{Remove}(obj, key, prev))\big] =$$
$$\Big(\{(id', obj', key', val') \in E \mid id' \notin prev\}, \ L\Big)$$

$$\mathsf{interp}\big[(E, L), (id, \mathsf{InsertAfter}(ref))\big] =$$
$$\begin{cases} (E, L) & \text{if } \nexists n. \, (ref, n) \in L \\ \Big(E, \, \{(p, n) \in L \mid p \neq ref\} \cup \{(ref, id)\} \cup \{(id, n) \mid (ref, n) \in L\}\Big) \\ \qquad \text{if } \exists n. \, (ref, n) \in L \end{cases}$$

$$\mathsf{interp}\big[(E, L), (id, \mathsf{MakeList})\big] = \Big(E, \, L \cup \{(id, \bot)\}\Big)$$

$$\mathsf{interp}\big[(E, L), (id, \mathsf{MakeMap})\big] = (E, L)$$

$$\mathsf{interp}\big[(E, L), (id, \mathsf{MakeVal}(val))\big] = (E, L)$$

The interpretation of Assign and Remove updates only E and leaves L unchanged; conversely, the interpretation of InsertAfter and MakeList updates only L. Both the Assign and Remove interpretations remove any tuples from causally prior assignments (those whose IDs appear in $prev$), but leave any tuples from concurrent assignments unchanged. This is the behaviour of a multi-value register; if a last-writer-wins register is required, the condition $id' \notin prev$ can be changed to $obj' \neq obj \ \vee \ key' \neq key$, which removes any existing tuples with the same object ID and key.

The interpretation of InsertAfter resembles the insertion into a linked list, as illustrated in Figure 1. For example, to interpret $(id, \mathsf{InsertAfter}(ref))$, if we have $(ref, next) \in L$, we remove the pair $(ref, next)$ from L, and add the pairs (ref, id) and $(id, next)$ to L. Thus, the new list element id is inserted between ref and $next$.

Note that L never shrinks, it only ever grows through interpreting InsertAfter operations. When a list element is removed by a Remove operation, the effect is that all values are

Before:

After adding $(25, \mathsf{InsertAfter}(13))$ to OpSet:

| 2 | 13 | $\rightarrow$ | 13 | 5 | $\rightarrow$ | 5 | 23 | $\rightarrow$ None |

$L = \{(2, 13),\ (13, 5),\ (5, 23),\ (23, \bot)\}$

$L' = L\ -\ \{(13, 5)\}\ \cup\ \{(13, 25),\ (25, 5)\}$

Figure 1 Illustration of the interpretation of an $\mathsf{InsertAfter}$ operation.

Hello!

Insert " Alice" between "o" and "!"

Insert " Charlie" between "o" and "!"

Hello Alice!

Hello Charlie!

Merge concurrent edits

Hello Al Ciharcliee!
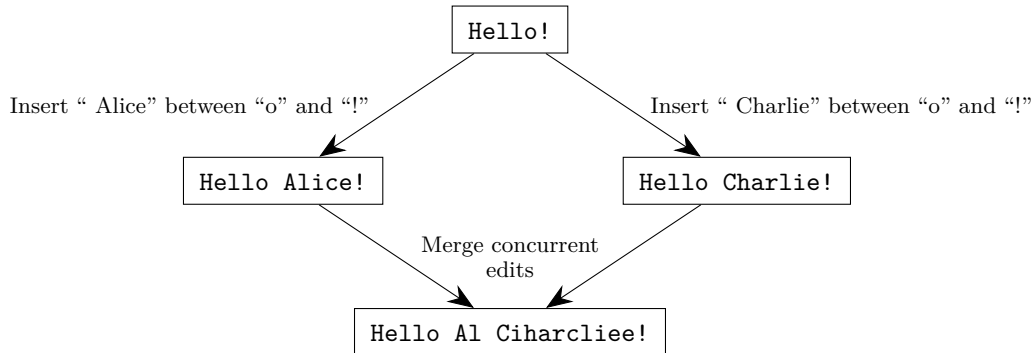
Figure 2 Two concurrent insertions at the same position are interleaved.

removed from the list element in the element relation E, but the list element remains in L as a *tombstone*, so that any concurrent $\mathsf{InsertAfter}$ operations can still locate the referenced list position. Thus, from a user's point of view a list element only exists if it has at least one associated value in the E relation; any list elements without an associated value should be ignored.

4   Discussion: Merging Text Edits

The datatypes we have specified in § 3 can support a wide range of applications. For example, the list datatype can be used to implement a collaborative text editor: by treating the text as a list of individual characters, every edit can be expressed as a sequence of insertion or deletion operations on the list.

The problem of collaborative text editing has been studied extensively, using two main approaches: Operational Transformation and CRDTs. We discuss this prior work in § 6. We will now highlight a scenario that, to our knowledge, has not been considered by any previous work on collaborative text editing.

Consider the execution illustrated in Figure 2. In this example, two users are concurrently editing a text document that initially reads "Hello!". The user on the left changes it to read "Hello Alice!", while concurrently the user on the right changes the document to read "Hello Charlie!". When the concurrent edits are merged, the algorithm randomly interleaves the two insertions of " Alice" and " Charlie" character by character, resulting in an unreadable jumble of characters.

The problem is even worse if the concurrent insertions are not just a single word, but an entire paragraph or section. In these cases, interleaving the users' insertions would most likely result in an entirely incomprehensible text that would have to be deleted and rewritten. Even though the merge in Figure 2 is so obviously undesirable, there is to our knowledge

no formal specification of collaborative text editing that rules out such an interleaving of insertions.

▶ **Theorem 1.** *The $\mathcal{A}_{strong}$ specification of collaborative text editing by Attiya et al. [4] allows the outcome in Figure 2; that is, an algorithm that interleaves concurrent insertions at the same position may nevertheless satisfy the $\mathcal{A}_{strong}$ specification. Moreover, the text editing CRDT algorithms Logoot [62, 63] and LSEQ [42, 43] also allow this outcome.*

Proof. Follows directly from the respective definitions, which are all based on the idea of assigning each character a position in a totally ordered identifier space, such that the order of identifiers corresponds to the order of characters in the document. When a new character is inserted, it is assigned an identifier that lies somewhere between the identifiers of its predecessor and successor. However, when concurrent insertions with the same predecessor and successor are performed, those insertions are ordered arbitrarily. Repeated insertions within the same predecessor-successor interval may thus be interleaved arbitrarily.

We also performed tests with open source implementations of Logoot [1, 3] and LSEQ [12, 42], and observed this interleaving anomaly occurring in practice. ◄

Rather than interleaving characters, a better approach to merging is to keep all insertions by a particular user together as a continuous sequence. With this constraint, there are two acceptable merged results in the example of Figure 2: either "Hello Alice Charlie!" or "Hello Charlie Alice!". The choice between these two outcomes is arbitrary, as there is no *a priori* requirement for one user's insertions to come before the other's.

▶ **Theorem 2.** *The list specification from § 3 does not allow interleaving of concurrent insertions. That is, if one user inserts a character sequence $\langle x_1, x_2, \ldots, x_n \rangle$ and another user concurrently inserts a character sequence $\langle y_1, y_2, \ldots, y_m \rangle$ at the same position, the merged document contains either the character sequence $\langle x_1, x_2, \ldots, x_n, y_1, y_2, \ldots, y_m \rangle$ or the character sequence $\langle y_1, y_2, \ldots, y_m, x_1, x_2, \ldots, x_n \rangle$ at the specified position.*

Proof. We formalise the list specification and Theorem 2 using the Isabelle/HOL proof assistant [64]. The formal proof development is summarised in Appendix C.3. ◄

For an informal argument why interleaving is ruled out, see Figure 3, which shows an editing scenario similar to Figure 2, but with the insertions of " Alice" and " Charlie" shortened to "Al" and "Ch" respectively. The example contains four insertion operations ("A", "l", "C", and "h"), which can be ordered in six possible ways. However, among the six possible operation orderings there are only two possible results: `ChAl` or `AlCh`. Interleavings such as `CAhl` or `AChl` never occur.

In fact, the end result depends only on the relative ordering of the operations that insert "A" and "C", respectively. All other operations can be reordered without affecting the outcome. Thus, even if the inserted strings are longer than two characters, their relative ordering only depends on the IDs of their first character. The remaining characters follow their initial character without interleaving.

Note that there are only six possible orderings of the four operations, and not $4! = 24$, because the Lamport timestamp ordering on identifiers (as given in § 2.1) is a linear extension of the causal order. In this example we assume that text is typed from left to right (that is, "A" is always inserted before "l", and "C" is inserted before "h"). This implies that the ID of the operation inserting "l" must be greater than that of the insertion of "A", and likewise the "h" insertion must be greater than the "C" insertion.

id_1, InsertAfter(id_0), "A" → A	id_1, InsertAfter(id_0), "A" → A	id_1, InsertAfter(id_0), "A" → A
id_2, InsertAfter(id_1), "l" → Al	id_2, InsertAfter(id_0), "C" → CA	id_2, InsertAfter(id_0), "C" → CA
id_3, InsertAfter(id_0), "C" → CAl	id_3, InsertAfter(id_1), "l" → CAl	id_3, InsertAfter(id_2), "h" → ChA
id_4, InsertAfter(id_3), "h" → ChAl	id_4, InsertAfter(id_2), "h" → ChAl	id_4, InsertAfter(id_1), "l" → ChAl
id_1, InsertAfter(id_0), "C" → C	id_1, InsertAfter(id_0), "C" → C	id_1, InsertAfter(id_0), "C" → C
id_2, InsertAfter(id_0), "A" → AC	id_2, InsertAfter(id_0), "A" → AC	id_2, InsertAfter(id_1), "h" → Ch
id_3, InsertAfter(id_2), "l" → AlC	id_3, InsertAfter(id_1), "h" → ACh	id_3, InsertAfter(id_0), "A" → ACh
id_4, InsertAfter(id_1), "h" → AlCh	id_4, InsertAfter(id_2), "l" → AlCh	id_4, InsertAfter(id_3), "l" → AlCh

Figure 3 All possible operation orderings when the strings "Al" (for "Alice") and "Ch" (for "Charlie") are concurrently inserted at the same position. The operation IDs are arbitrary; we only require that $id_0 < id_1 < id_2 < id_3 < id_4$.

▶ **Theorem 3.** *The OpSet list specification from § 3 is strictly stronger than the $\mathcal{A}_{\mathsf{strong}}$ specification of Attiya et al [4]. That is, any algorithm that satisfies the list specification given in § 3 also satisfies $\mathcal{A}_{\mathsf{strong}}$, but the converse is not true.*

Proof. We formalise the $\mathcal{A}_{\mathsf{strong}}$ specification with Isabelle/HOL, and produce a mechanically verified proof that every possible execution of the list specification from § 3 satisfies all conditions of $\mathcal{A}_{\mathsf{strong}}$. The formal proof development is summarised in Appendix C.5. The fact that our specification is *strictly* stronger follows from Theorems 1 and 2.   ◀

▶ **Theorem 4.** *The RGA algorithm [49] satisfies the OpSet list specification introduced in this paper, while Logoot [62, 63] and LSEQ [42, 43] do not.*

Proof. We use Isabelle/HOL to prove that RGA satisfies our specification, as described in Appendix C.4. Our Isabelle/HOL implementation of RGA is based on the formalisation that we developed in previous work [20, 21]. The fact that Logoot and LSEQ do not satisfy our specification follows directly from Theorems 1 and 2.   ◀

5   A Replicated Tree Datatype

In § 3 we gave an OpSet specification of a replicated object graph datatype. In this model, every map or list object has a unique ID (namely, the ID of the MakeMap or MakeList operation that created it), and objects can reference each other using these IDs.

We now build upon this model, showing how to restrict the object graph so that it is always a tree. A tree is a graph in which every vertex has exactly one parent (except for the root, which has no parent), and in which the parent relation has no cycles. Tree data structures are useful in many applications: for example, file systems (consisting of directories and files) and XML or JSON documents are trees. Branch nodes in this tree may be either maps or lists, and leaf nodes are primitive values (wrapped in a MakeVal operation).

5.1   The Difficulty of a Move Operation

In applications that use tree-structured data, a frequently required operation is to *move* a subtree to a new location within the tree. For example:

- In a filesystem, renaming a directory can be expressed as moving the directory node from the old name to the new name. Similarly, a directory may be moved to a new path.
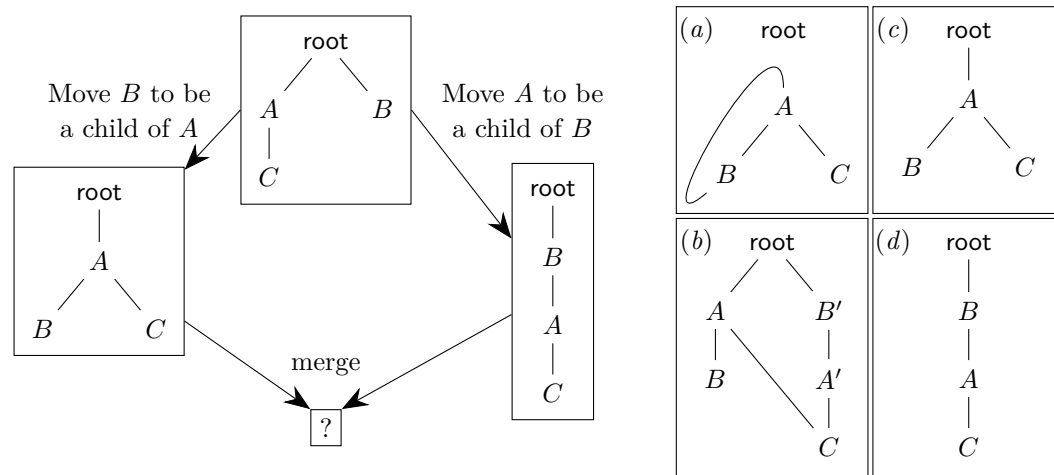
█ **Figure 4** Initially, A and B are siblings. B is moved to be a child of A, while concurrently A is moved to be a child of B. Boxes (a) to (d) show possible outcomes of the merge.

- In vector graphics applications, several graphical objects may be grouped together as a logical unit. This operation can be expressed by creating a new branch node to represent the group, and then moving the individual objects to be children of that group node.

- In a to-do list application, users may use the order of items in the list to denote a priority order, and they may drag and drop items to change their relative order. Reordering items is equivalent to moving items to new locations within the list.

A move operation can be naively emulated by deleting the subtree from its old location and recreating it at the new location. However, if two users perform this process concurrently, the resulting tree will contain two copies of the moved subtree, which would be undesirable in all of the application examples given above. Thus, we require an *atomic move* operation that does not create duplicate objects in case of concurrent moves.

A more subtle kind of conflict is illustrated in Figure 4. Here, B is moved to be a child of A, while concurrently A is moved to be a child of B. If the CRDT does not take care to detect this situation, it may introduce a cycle in the merged result, as shown in Figure 4(a); this result is no longer a tree. Handling such conflicting move operations is a challenging problem, and to our knowledge no existing implementation of a tree CRDT has found an adequate solution to this problem.

Several CRDT tree datatypes for XML [37, 45] and JSON data [32, 34, 58] have been developed, but to our knowledge, none of them define a move operation. Tao et al. [55] implemented a CRDT-based replicated filesystem, resolving concurrent moves with an approach illustrated in Figure 4(b): conflicting branch nodes (directories) are duplicated, and leaf nodes (files) may be referenced from multiple branch nodes. Thus, Tao et al.'s data structure is strictly a DAG, not a tree.

Najafzadeh [40, 41] also implemented a CRDT-based replicated filesystem, but chose a different approach: move operations must acquire a global lock before they can proceed, which ensures that conflicting concurrent move operations cannot occur in the first place. This conservative approach rules out move conflicts, but the resulting datatype is not strictly a CRDT, since some operations require strongly consistent synchronisation.

5.2 Specifying a Tree with Atomic Moves

We now demonstrate the power of the OpSets approach by using it to define a tree CRDT with an anomaly-free atomic move operation. Our specification rules out violations of the tree structure such as those in Figure 4(*a*,*b*), and concurrent moves do not duplicate tree nodes. Moreover, our CRDT does not require any locks or global synchronisation.

When the OpSet contains conflicting move operations, our specification chooses one of them as the one that takes effect, and simply ignores the other conflicting operations. Thus, in the example of Figure 4, the merged outcome of the two conflicting move operations is either (*c*) or (*d*). If two users concurrently move the same item to different locations, the move operation with the greater ID determines the item's final location. However, in non-conflict situations, all concurrent move operations take effect.

We define a tree to be a restricted form of the object graph specified in § 3. First, we require that there is a designated root object: assume that we have an operation ID root that is less than all other operation IDs (according to the total order on identifiers, introduced in § 2.1). Further assume that for any OpSet O specifying a tree, we have either (root, MakeList) $\in O$ or (root, MakeMap) $\in O$, depending on whether the root node is a list or a map. We define an object x to be the *parent* of an object y if one of the values in x is a reference to y. The *ancestor* relation is the transitive closure of the parent relation, defined using the element relation E:

$$\mathsf{parent}(E,\, i) = \begin{cases} \big\{(obj,\, val) \mid \exists\, id, key.\ (id,\, obj,\, key,\, val) \in E\big\} & \text{if } i = 1 \\ \big\{(x,\, z) \mid (x,\, y) \in \mathsf{parent}(E,\, i-1)\, \wedge\, (y,\, z) \in \mathsf{parent}(E,\, 1)\big\} & \text{if } i > 1 \end{cases}$$

$$\mathsf{ancestor}(E) = \bigcup_{i\, \geq\, 1} \mathsf{parent}(E,\, i)$$

An object graph is a tree if the root has no parent, every non-root node has exactly one parent, and if the ancestor relation has no cycles. We can redefine the operation interpretations from § 3.2 to preserve this tree invariant. In fact, it is sufficient to redefine only the interpretation of Assign, and to leave the interpretation of the other five operation types unchanged:

$$\mathsf{interp}\big[(E,\, L),\, (id,\, \mathsf{Assign}(obj, key, val, prev))\big] \;=$$
$$\begin{cases} (E,\, L) & \text{if } (val,\, obj) \in \mathsf{ancestor}(E) \\ \Big(\big\{(id', obj', key', val') \in E \mid id' \notin prev \wedge val' \neq val\big\}\, \cup\, \big\{(id, obj, key, val)\big\},\, L\Big) \\ \qquad\qquad \text{if } (val,\, obj) \notin \mathsf{ancestor}(E) \end{cases}$$

This definition differs in two ways from that in § 3.2. Firstly, the operation has no effect if *val* is already an ancestor of the proposed parent *obj*, since the operation would otherwise introduce a cycle. Secondly, any existing tuple in E that references the same value *val* is removed, preserving the invariant that every non-root node must have exactly one parent.

This interpretation of Assign performs an atomic move whenever *val* is the ID of an existing object in the tree; in that case, it is moved from its existing position to the key *key* in the object *obj*. If *val* does not currently exist in the tree (e.g. because it has just been created), the operation behaves like conventional assignment.

6   Related Work

6.1   Interpretation of Operation Sequences

The general idea of establishing a total order of operations, and executing them in that order, appears in many areas of computing: for example, in the state machine approach to replication [50], the event sourcing approach to data modelling [59], write-ahead logs for crash recovery [38], serializable transactions [13], and scalable multicore data structures [8]. However, beneath the superficial similarity of these approaches there are important differences that need to be distinguished.

As discussed in § 2.3, many of these systems rely on the property that after some operation is executed, all subsequent operations will appear *after* it in the total order. In other words, the operation sequence is an append-only log, and new operations never need to be inserted ahead of an existing operation in the total order. This is a very strong property: in the context of a distributed system, it requires an atomic broadcast (or total order broadcast) protocol [15], which is equivalent to solving distributed consensus [11]. This class of protocols requires communication with a quorum of nodes in order to make progress [27], and it cannot guarantee progress in a fully asynchronous setting [17].

By contrast, the sequential OpSet interpretation of § 2.2 does not require atomic broadcast because it allows operations to be added to the OpSet in any order, and it assigns operation IDs without any coordination. Few systems use this approach; the most closely related prior work are the Bayou system [57], which executes tentative transactions deterministically in timestamp order, and Burckhardt's *standard conflict resolution* [9, § 4.3.3]. Both of these share the OpSet approach's characteristic that operations with a higher ID need to be undone and re-applied when a new operation with a lower ID is received.

Our contribution in this paper is to formulate the OpSet approach more generally as a tool for specifying and reasoning about complex replicated data structures, such as lists and trees. Our work is the first to use this approach in mechanised proofs, in which we show that a non-OpSet list CRDT (RGA) satisfies an OpSet-based specification, and prove the absence of the interleaving anomaly in Figure 2.

Baquero et al. [6] and Grishchenko [23] have proposed representing CRDTs in terms of a partially-ordered log of operations, where the partial order captures the causal relationships between operations. The OpSet approach can be seen as a variant of this idea, in which we define the total order on identifiers to be a linear extension of the partial order.

6.2   Specification and Verification of Replicated Datatypes

Algorithms for collaboratively editing a shared data structure have been the topic of active research for approximately 30 years, under the headings of Operational Transformation [16, 48, 53, 47] and CRDTs [51, 52]. However, throughout this time, the exact consistency properties provided by the algorithms have been somewhat unclear. For example, Sun et al. [54] identified three desirable properties that they articulated informally: *convergence*, *causality preservation*, and *intention preservation*. While the definition of the first two properties is fairly unambiguous, the definition of "intention preservation" leaves much more room for interpretation. Efforts to formally specify and verify the semantics of replicated datatypes have replaced such informal statements with precise consistency properties.

Burckhardt et al. [10] provide a wide-ranging formal account of CRDTs, covering their specification, verification, and optimality, with the semantics of an operation on a replicated datatype given as a function of the operation, *o*, and a *operation context*—the set of operations

visible to a node at the time that *o* was received. Our OpSets can be seen as an explicitly executable variation on this idea: nodes record all operations that they have ever received in a monotonically growing set, and the interpretation function builds the result "bottom up" in a fold-like operation. In contrast to Burckhardt et al., who focus on applying their techniques to set and counter datatypes, we apply our approach to the specification of lists, maps, and trees, using our OpSets as a tool for designing new replicated datatypes—including those previously thought impossible, such as our replicated tree with atomic move. Gotsman et al. [22] extend Burckhardt et al.'s formalism to reason about hybrid consistency models, providing a modular proof rule inspired by permissions-based logics to enforce an integrity invariant for a given consistency model.

Bieniusa et al. [7] articulate a *principle of permutation equivalence* that partially specifies the expected semantics of replicated datatypes, but which leaves some combinations of operations unspecified. Zeller et al. [65] formalise counters, registers, and sets using Isabelle/HOL and provide mechanised proofs of their correctness. Attiya et al. [4] give two specifications of collaborative text editing ($\mathcal{A}_{\mathsf{strong}}$ and $\mathcal{A}_{\mathsf{weak}}$), prove that the RGA CRDT [49] satisfies $\mathcal{A}_{\mathsf{strong}}$, and conjecture that the Operational Transformation algorithm Jupiter [44] satisfies $\mathcal{A}_{\mathsf{weak}}$. Wei et al. [61] complete the proof that Jupiter satisfies $\mathcal{A}_{\mathsf{weak}}$.

In our prior work [21] we establish a formal verification framework for CRDTs in Isabelle/HOL, and verify the strong eventual consistency properties (in particular, convergence) of a list, set, and counter datatype. The Isabelle implementation of RGA we use in § 3 is based on this work [20]. However, this work does not specify the datatype semantics beyond the convergence property.

Gaducci et al. [18] develop a semantics for replicated datatypes, placing a focus on compositionality, where a replicated datatype is modelled as a function from labelled directed acyclic graphs of events to sets of values, with each value in this set potentially observable at a node under different ordering of events observed at that node. A notion of behavioural *refinement* for replicated datatypes induced by set inclusion is also defined, along with a generalisation of their relational semantics to a categorical one.

Mukund et al. [39] use traces to provide bounded declarative specifications of CRDTs and show how Counter Example Guided Abstract Refinement (CEGAR) can be used to automatically verify a reference CRDT implementation against its bounded specification.

6.3 Collaborative Tree Datatypes

For collaborative editing of tree data structures, several CRDTs [37, 32] and Operational Transformation algorithms [29, 28, 14] have been proposed. However, most of them only consider insertion and deletion of tree nodes, but do not support a move operation.

As explained in § 5, supporting an operation that can move a subtree to a new location within a tree introduces new conflicts that need to be handled. Ahmed-Nacer et al. [2] survey approaches to handling these conflicts without providing concrete algorithms. Tao et al. [55] propose handling conflicting move operations by allowing the same object to appear in more than one location; thus, their datatype is strictly a DAG, not a tree.

Najafzadeh [40, 41] asserts that concurrent move operations on a tree cannot safely be implemented in a CRDT, since the precondition of a move operation is not stable. Najafzadeh suggests the use of locks to globally synchronise move operations, preventing a scenario such as that in Figure 4 from ever occurring. However, the resulting datatype is not strictly a CRDT, since some operations require strongly consistent synchronisation.

To our knowledge, our move semantics specified in § 5 is the first definition of such an operation on a fully asynchronous tree CRDT. We avoid the apparent contradiction with

Najafzadeh's assertion by evaluating the precondition $(val, obj) \notin \mathsf{ancestor}(E)$ at the same time as applying the operation, rather than at the time when the operation is generated, and by applying all operations in the OpSet in a deterministic order.

7 Conclusion

In this work we have introduced Operation Sets (OpSets), a simple but powerful approach for specifying the semantics of replicated datatypes. We specified a variety of common, composable replicated datatypes in the OpSets model, and used Isabelle/HOL to formally reason about their properties. We have used this specification to highlight an interleaving anomaly that affects some existing collaborative text editing algorithms, and proved that the RGA algorithm satisfies our list specification. Finally, we demonstrated how the OpSet model to can be used to develop new replication algorithms, and we introduced a specification for an atomic move operation in a tree CRDT.

The OpSets approach is an executable specification that precisely defines the permitted states of a replica after some set of updates have been applied. In this paper we have used a sequential OpSet interpretation: operations are applied in strict ascending order of ID. This property is very useful as it trivially ensures convergence, and it simplifies reasoning about specifications and invariants for CRDTs. In contrast, the traditional approach to defining CRDTs requires operations to be commutative, increasing their complexity. In proving that RGA satisfies our list specification, we demonstrate a correspondence between sequential specification and commutative implementation; for future work it will be interesting to further explore this correspondence for other datatypes. In particular, we hypothesise that it is possible to derive a tree CRDT with a commutative move operation from the specification in § 5, which could then be used to implement a distributed peer-to-peer file system.

Although we focussed on sequential OpSet interpretations in this paper, note that any deterministic function can be used as interpretation function $[\![-]\!]$. In particular, one can view the OpSet as a *database of facts*, containing all changes ever made to the shared data, and the interpretation function as a *query* over this database. The resulting datatype is then a *materialized view* in database terminology. When new operations are added to an OpSet O, computing the corresponding change to $[\![O]\!]$ is a materialized view maintenance problem, for which optimised algorithms have been developed [24]. We hypothesise that these techniques can be applied to replicated datatypes, allowing efficient CRDT implementations to be derived from an OpSet-based specification.

Acknowledgements

───── **References** ─────

1 Mehdi Ahmed-Nacer, Claudia-Lavinia Ignat, Gérald Oster, Hyun-Gul Roh, and Pascal Urso. Evaluating CRDTs for real-time document editing. In *11th ACM Symposium on Document Engineering (DocEng)*, pages 103–112, September 2011. `doi:10.1145/2034691.2034717`.

2   Mehdi Ahmed-Nacer, Stéphane Martin, and Pascal Urso. File system on CRDT. Technical Report RR-8027, INRIA, July 2012. URL: `https://hal.inria.fr/hal-00720681/`.

3   Mehdi Ahmed-Nacer, Gérald Oster, and Pascal Urso. Java benchmarker of optimistic replication algorithms. URL: `https://github.com/PascalUrso/ReplicationBenchmark`.

4   Hagit Attiya, Sebastian Burckhardt, Alexey Gotsman, Adam Morrison, Hongseok Yang, and Marek Zawirski. Specification and complexity of collaborative text editing. In *ACM Symposium on Principles of Distributed Computing (PODC)*, pages 259–268, July 2016. `doi:10.1145/2933057.2933090`.

5   Peter Bailis and Ali Ghodsi. Eventual consistency today: Limitations, extensions, and beyond. *ACM Queue*, 11(3), March 2013. `doi:10.1145/2460276.2462076`.

6   Carlos Baquero, Paulo Sérgio Almeida, and Ali Shoker. Making operation-based CRDTs operation-based. In *14th IFIP International Conference on Distributed Applications and Interoperable Systems (DAIS)*, pages 126–140, June 2014. `doi:10.1007/978-3-662-43352-2_11`.

7   Annette Bieniusa, Marek Zawirski, Nuno Preguiça, Marc Shapiro, Carlos Baquero, Valter Balegas, and Sérgio Duarte. Brief announcement: Semantics of eventually consistent replicated sets. In *26th International Symposium on Distributed Computing (DISC)*, pages 441–442, October 2012. `doi:10.1007/978-3-642-33651-5_48`.

8   Silas Boyd-Wickizer, M Frans Kaashoek, Robert Morris, and Nickolai Zeldovich. OpLog: a library for scaling update-heavy data structures. Technical Report MIT-CSAIL-TR-2014-019, MIT CSAIL, September 2014. URL: `http://hdl.handle.net/1721.1/89653`.

9   Sebastian Burckhardt. Principles of eventual consistency. *Foundations and Trends in Programming Languages*, 1(1-2):1–150, October 2014. `doi:10.1561/2500000011`.

10   Sebastian Burckhardt, Alexey Gotsman, Hongseok Yang, and Marek Zawirski. Replicated data types: Specification, verification, optimality. In *41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 271–284, January 2014. `doi:10.1145/2535838.2535848`.

11   Tushar Deepak Chandra and Sam Toueg. Unreliable failure detectors for reliable distributed systems. *Journal of the ACM*, 43(2):225–267, March 1996. `doi:10.1145/226643.226647`.

12   Chat-Wane. LSEQTree. URL: `https://github.com/Chat-Wane/LSEQTree`.

13   Susan B Davidson, Hector Garcia-Molina, and Dale Skeen. Consistency in partitioned networks. *ACM Computing Surveys*, 17(3):341–370, September 1985. `doi:10.1145/5505.5508`.

14   Aguido Horatio Davis, Chengzheng Sun, and Junwei Lu. Generalizing operational transformation to the Standard General Markup Language. In *ACM Conference on Computer Supported Cooperative Work (CSCW)*, pages 58–67, November 2002. `doi:10.1145/587078.587088`.

15   Xavier Défago, André Schiper, and Péter Urbán. Total order broadcast and multicast algorithms: Taxonomy and survey. *ACM Computing Surveys*, 36(4):372–421, December 2004. `doi:10.1145/1041680.1041682`.

16   Clarence Ellis and S J Gibbs. Concurrency control in groupware systems. In *ACM International Conference on Management of Data (SIGMOD)*, pages 399–407, May 1989. `doi:10.1145/67544.66963`.

17   Michael J Fischer, Nancy A Lynch, and Michael S Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, April 1985. `doi:10.1145/3149.214121`.

18   Fabio Gadducci, Hernán C. Melgratti, and Christian Roldán. A denotational view of replicated data types. In *19th International Conference on Coordination Models and Languages (COORDINATION)*, pages 138–156, June 2017. `doi:10.1007/978-3-319-59746-1_8`.

19 Seth Gilbert and Nancy A Lynch. Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services. *ACM SIGACT News*, 33(2):51–59, 2002. `doi: 10.1145/564585.564601`.

20 Victor B F Gomes, Martin Kleppmann, Dominic P Mulligan, and Alastair R Beresford. A framework for establishing strong eventual consistency for conflict-free replicated data types. *Archive of Formal Proofs*, July 2017. URL: `http://isa-afp.org/entries/CRDT.html`.

21 Victor B F Gomes, Martin Kleppmann, Dominic P Mulligan, and Alastair R Beresford. Verifying strong eventual consistency in distributed systems. *Proceedings of the ACM on Programming Languages (PACMPL)*, 1(OOPSLA), October 2017. `doi:10.1145/3133933`.

22 Alexey Gotsman, Hongseok Yang, Carla Ferreira, Mahsa Najafzadeh, and Marc Shapiro. 'Cause I'm strong enough: reasoning about consistency choices in distributed systems. In *43rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 371–384, January 2016. `doi:10.1145/2837614.2837625`.

23 Victor Grishchenko. Citrea and Swarm: Partially ordered op logs in the browser. In *1st Workshop on Principles and Practice of Eventual Consistency (PaPEC)*, April 2014. `doi:10.1145/2596631.2596641`.

24 Ashish Gupta and Inderpal Singh Mumick. *Materialized Views: Techniques, Implementations, and Applications.* MIT Press, May 1999.

25 Florian Haftmann and Makarius Wenzel. Local theory specifications in Isabelle/Isar. In *International Workshop on Types for Proofs and Programs (TYPES)*, pages 153–168, 2008. `doi:10.1007/978-3-642-02444-3_10`.

26 Maurice P Herlihy and Jeannette M Wing. Linearizability: A correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 12(3):463–492, July 1990. `doi:10.1145/78969.78972`.

27 Heidi Howard, Dahlia Malkhi, and Alexander Spiegelman. Flexible Paxos: Quorum intersection revisited. In *20th International Conference on Principles of Distributed Systems (OPODIS)*, December 2016. `doi:10.4230/LIPIcs.OPODIS.2016.25`.

28 Claudia-Lavinia Ignat and Moira C Norrie. Customizable collaborative editor relying on treeOPT algorithm. In *8th European Conference on Computer-Supported Cooperative Work (ECSCW)*, pages 315–334, September 2003. `doi:10.1007/978-94-010-0068-0_17`.

29 Tim Jungnickel and Tobias Herb. Simultaneous editing of JSON objects via operational transformation. In *31st Annual ACM Symposium on Applied Computing (SAC)*, pages 812–815, April 2016. `doi:10.1145/2851613.2852003`.

30 Florian Kammüller, Markus Wenzel, and Lawrence C. Paulson. Locales - A sectioning concept for Isabelle. In *12th International Conference on Theorem Proving in Higher Order Logics (TPHOLs)*, pages 149–166, 1999. `doi:10.1007/3-540-48256-3_11`.

31 Martin Kleppmann. *Designing Data-Intensive Applications.* O'Reilly Media, April 2017.

32 Martin Kleppmann and Alastair R Beresford. A conflict-free replicated JSON datatype. *IEEE Transactions on Parallel and Distributed Systems*, 28(10):2733–2746, April 2017. `doi: 10.1109/TPDS.2017.2697382`.

33 Martin Kleppmann, Victor B. F. Gomes, Dominic P. Mulligan, and Alastair R. Beresford. OpSets: Sequential specifications for replicated datatypes (proof document), May 2018. URL: `https://www.isa-afp.org/entries/OpSets.html`.

34 Martin Kleppmann, Peter van Hardenberg, and Orion Henry. Automerge. URL: `https://github.com/automerge/automerge`.

35 Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, July 1978. `doi:10.1145/359545.359563`.

36 Paul J Leach, Michael Mealling, and Rich Salz. A Universally Unique IDentifier (UUID) URN namespace. IETF Standards Track, RFC 4122, July 2005. `doi:10.17487/rfc4122`.

37 Stéphane Martin, Pascal Urso, and Stéphane Weiss. Scalable XML collaborative editing with undo. In *On the Move to Meaningful Internet Systems*, pages 507–514, October 2010. `doi:10.1007/978-3-642-16934-2_37`.

38 C Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, and Peter Schwarz. ARIES: A transaction recovery method supporting fine-granularity locking and partial rollbacks using write-ahead logging. *ACM Transactions on Database Systems (TODS)*, 17(1):94–162, March 1992. `doi:10.1145/128765.128770`.

39 Madhavan Mukund, Gautham Shenoy R., and S. P. Suresh. Effective verification of replicated data types using later appearance records (LAR). In *13th International Symposium on Automated Technology for Verification and Analysis (ATVA)*, pages 293–308, October 2015. `doi:10.1007/978-3-319-24953-7_23`.

40 Mahsa Najafzadeh. *The Analysis and Co-design of Weakly-Consistent Applications*. PhD thesis, Université Pierre et Marie Curie, August 2016. URL: `https://tel.archives-ouvertes.fr/tel-01351187v1`.

41 Mahsa Najafzadeh, Marc Shapiro, and Patrick Eugster. Co-design and verification of an available file system. In *19th International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI)*, pages 358–381, January 2018. `doi:10.1007/978-3-319-73721-8_17`.

42 Brice Nédelec, Pascal Molli, and Achour Mostefaoui. CRATE: Writing stories together with our browsers. In *25th International World Wide Web Conference (WWW)*, pages 231–234, April 2016. `doi:10.1145/2872518.2890539`.

43 Brice Nédelec, Pascal Molli, Achour Mostefaoui, and Emmanuel Desmontils. LSEQ: an adaptive structure for sequences in distributed collaborative editing. In *13th ACM Symposium on Document Engineering (DocEng)*, pages 37–46, September 2013. `doi:10.1145/2494266.2494278`.

44 David A Nichols, Pavel Curtis, Michael Dixon, and John Lamping. High-latency, low-bandwidth windowing in the Jupiter collaboration system. In *8th Annual ACM Symposium on User Interface Software and Technology (UIST)*, pages 111–120, November 1995. `doi:10.1145/215585.215706`.

45 Petru Nicolaescu, Kevin Jahns, Michael Derntl, and Ralf Klamma. Yjs: A framework for near real-time P2P shared editing on arbitrary data types. In *15th International Conference on Web Engineering (ICWE)*, June 2015. `doi:10.1007/978-3-319-19890-3_55`.

46 Tobias Nipkow and Gerwin Klein. *Concrete Semantics - With Isabelle/HOL*. Springer, 2014. `doi:10.1007/978-3-319-10542-0`.

47 Gérald Oster, Pascal Molli, Pascal Urso, and Abdessamad Imine. Tombstone transformation functions for ensuring consistency in collaborative editing systems. In *2nd International Conference on Collaborative Computing (CollaborateCom)*, 2006. `doi:10.1109/COLCOM.2006.361867`.

48 Matthias Ressel, Doris Nitsche-Ruhland, and Rul Gunzenhäuer. An integrating, transformation-oriented approach to concurrency control and undo in group editors. In *ACM Conference on Computer Supported Cooperative Work (CSCW)*, pages 288–297, November 1996. `doi:10.1145/240080.240305`.

49 Hyun-Gul Roh, Myeongjae Jeon, Jin-Soo Kim, and Joonwon Lee. Replicated abstract data types: Building blocks for collaborative applications. *Journal of Parallel and Distributed Computing*, 71(3):354–368, 2011. `doi:10.1016/j.jpdc.2010.12.006`.

50 Fred B Schneider. Implementing fault-tolerant services using the state machine approach: A tutorial. *ACM Computing Surveys*, 22(4):299–319, December 1990. `doi:10.1145/98163.98167`.

51 Marc Shapiro, Nuno Preguiça, Carlos Baquero, and Marek Zawirski. A comprehensive study of convergent and commutative replicated data types. Technical Report 7506, INRIA, 2011. URL: `http://hal.inria.fr/inria-00555588/`.

52 Marc Shapiro, Nuno Preguiça, Carlos Baquero, and Marek Zawirski. Conflict-free replicated data types. In *13th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)*, pages 386–400, October 2011. `doi:10.1007/978-3-642-24550-3_29`.

53 Chengzheng Sun and Clarence Ellis. Operational transformation in real-time group editors: Issues, algorithms, and achievements. In *ACM Conference on Computer Supported Cooperative Work (CSCW)*, pages 59–68, November 1998. `doi:10.1145/289444.289469`.

54 Chengzheng Sun, Xiaohua Jia, Yanchun Zhang, Yun Yang, and David Chen. Achieving convergence, causality preservation, and intention preservation in real-time cooperative editing systems. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 5(1):63–108, 1998. `doi:10.1145/274444.274447`.

55 Vinh Tao, Marc Shapiro, and Vianney Rancurel. Merging semantics for conflict updates in geo-distributed file systems. In *8th ACM International Systems and Storage Conference (SYSTOR)*, May 2015. `doi:10.1145/2757667.2757683`.

56 Douglas B Terry, Alan J Demers, Karin Petersen, Mike J Spreitzer, Marvin M Theimer, and Brent B Welch. Session guarantees for weakly consistent replicated data. In *3rd International Conference on Parallel and Distributed Information Systems (PDIS)*, pages 140–149, September 1994. `doi:10.1109/PDIS.1994.331722`.

57 Douglas B Terry, Marvin M Theimer, Karin Petersen, Alan J Demers, Mike J Spreitzer, and Carl H Hauser. Managing update conflicts in Bayou, a weakly connected replicated storage system. In *15th ACM Symposium on Operating Systems Principles (SOSP)*, pages 172–182, December 1995. `doi:10.1145/224056.224070`.

58 Frank S Thomas. crjdt: A conflict-free replicated JSON datatype in Scala. URL: `https://github.com/fthomas/crjdt`.

59 Vaughn Vernon. *Implementing Domain-Driven Design*. Addison-Wesley Professional, February 2013.

60 Werner Vogels. Eventually consistent. *Communications of the ACM*, 52(1):40–44, January 2009. `doi:10.1145/1435417.1435432`.

61 Hengfeng Wei, Yu Huang, and Jian Lu. Specification and implementation of replicated list: The Jupiter protocol revisited. *arxiv.org*, August 2017. URL: `https://arxiv.org/abs/1708.04754`.

62 Stéphane Weiss, Pascal Urso, and Pascal Molli. Logoot: A scalable optimistic replication algorithm for collaborative editing on P2P networks. In *29th IEEE International Conference on Distributed Computing Systems (ICDCS)*, pages 404–412, June 2009. `doi:10.1109/ICDCS.2009.75`.

63 Stéphane Weiss, Pascal Urso, and Pascal Molli. Logoot-Undo: Distributed collaborative editing system on P2P networks. *IEEE Transactions on Parallel and Distributed Systems*, 21(8):1162–1174, January 2010. `doi:10.1109/TPDS.2009.173`.

64 Makarius Wenzel, Lawrence C. Paulson, and Tobias Nipkow. The Isabelle framework. In *21st International Conference on Theorem Proving in Higher Order Logics (TPHOLs)*, pages 33–38, August 2008. `doi:10.1007/978-3-540-71067-7_7`.

65 Peter Zeller, Annette Bieniusa, and Arnd Poetzsch-Heffter. Formal specification and verification of CRDTs. In *34th IFIP International Conference on Formal Techniques for Distributed Objects, Components and Systems (FORTE)*, June 2014. `doi:10.1007/978-3-662-43613-4_3`.

A   Generating Operations

Listing 1 gives pseudocode for functions that generate these operations. Intuitively, these functions form an API through which nodes can modify OpSets to indirectly describe replicated list and map datatypes. The first parameter O of each function is the OpSet that defines the current state of the node, and the functions return an updated OpSet containing new operations. The interpretation $[\![O]\!]$ returns a pair (E, L) as defined in § 3.2. The function newID(O) returns a unique ID (e.g. a Lamport timestamp [35]) that is greater than any existing ID in the OpSet O. Note that we elide explicit network broadcasts reflecting changes in a node's OpSet.

The function SETMAPKEY can be called by a user to update a map object with ID *map*, setting a key *key* to a value *val*. If *val* references an existing object, id_1 is set to the ID of that existing object; otherwise, the function VALUEID generates a new Make$\cdots$ operation for the value, and the new operation is added to O'. A new ID id_2 is generated for the Assign operation, in which the key *key* is set to id_1 (which identifies the value *val*). Finally, the function returns the OpSet with the Assign operation included. The other definitions follow a similar pattern.

The list manipulation functions SETLISTINDEX, INSLISTINDEX and REMOVELISTINDEX take a numeric index as argument to identify the position in the list being edited. The numeric index is translated into the ID of a list element using the function idxKey$_{E, L}$():

$$
\text{idxKey}_{E, L}(obj, key, i) \;=\; \begin{cases} \text{idxKey}_{E, L}(obj, n, i - 1) \\ \quad \text{if } i > 0 \land (key, n) \in L \; \land \; \exists\, id, val.\, (id, obj, key, val) \in E \\ \text{idxKey}_{E, L}(obj, n, i) \\ \quad\quad\quad \text{if } (key, n) \in L \; \land \; \nexists\, id, val.\, (id, obj, key, val) \in E \\ key \quad\quad \text{if } i = 0 \quad\quad \land\; \exists\, id, val.\, (id, obj, key, val) \in E \end{cases}
$$

key is initially the ID of the MakeList operation that created the list. The function recursively moves along the linked list structure in L, decrementing the index for every list element that has an associated value, and not counting any list elements without associated value (which are treated as deleted). Eventually, it returns the ID of the list element with the desired index.

B   Introduction to Isabelle/HOL

To help any readers who are not familiar with Isabelle/HOL, this appendix provides a brief introduction to the key concepts and syntax, taken from our previous work [21]. A more detailed introduction can be found in the standard tutorial material [46].

B.1   Syntax of expressions.

Isabelle/HOL is a logic with a strict, polymorphic, inferred type system. *Function types* are written $\tau_1 \Rightarrow \tau_2$, and are inhabited by *total* functions, mapping elements of τ_1 to elements of τ_2. We write $\tau_1 \times \tau_2$ for the *product type* of τ_1 and τ_2, inhabited by pairs of elements of type τ_1 and τ_2, respectively. *Type operators* are applied to arguments in reverse order: τ *list* denotes the type of lists of elements of type τ, and τ *set* denotes the type of mathematical (i.e., potentially infinite) sets of type τ, for instance. Type variables are written in lowercase, and preceded with a prime: $'a \Rightarrow 'a$ denotes the type of a polymorphic identity function, for

Listing 1 Generating new operations for modifying maps and lists.

function SETMAPKEY(O, map, key, val)
  $(E, L) = [\![O]\!]$
  $(id_1, op_1) = $ VALUEID(O, val)
  if $op_1 \neq \bot$ **then**
    $O := O \cup \{(id_1, op_1)\})$
  end if
  $id_2 = $ newID(O)
  $prev = \{id \mid \exists v.\ (id, map, key, v) \in E\}$
  return $O \cup$
    $\{(id_2, \mathsf{Assign}(map, key, id_1, prev))\}$
end function

function REMOVEMAPKEY(O, map, key)
  $(E, L) = [\![O]\!]$
  $id_1 = $ newID(O)
  $prev = \{id \mid \exists v.\ (id, map, key, v) \in E\}$
  return $O \cup$
    $\{(id_1, \mathsf{Remove}(map, key, prev))\}$
end function

function VALUEID(O, val)
  if val is a primitive type **then**
    return (newID(O), $\mathsf{MakeVal}(val)$)
  else if $val = $ [] (empty list literal) **then**
    return (newID(O), $\mathsf{MakeList}$)
  else if $val = $ {} (empty map literal) **then**
    return (newID(O), $\mathsf{MakeMap}$)
  else (val is an existing object)
    return (objID(val), $\bot$)
  end if
end function

function SETLISTINDEX(O, $list$, $index$, val)
  $(E, L) = [\![O]\!]$
  $(id_1, op_1) = $ VALUEID(O, val)
  if $op_1 \neq \bot$ **then**
    $O := O \cup \{(id_1, op_1)\})$
  end if
  $id_2 = $ newID(O)
  $key = $ idxKey$_{E,\, L}(list, list, index)$
  $prev = \{id \mid \exists v.\ (id, list, key, v) \in E\}$
  return $O \cup$
    $\{(id_2, \mathsf{Assign}(list, key, id_1, prev))\}$
end function

function INSLISTINDEX(O, $list$, $index$, val)
  $(E, L) = [\![O]\!]$
  $(id_1, op_1) = $ VALUEID(O, val)
  if $op_1 \neq \bot$ **then**
    $O := O \cup \{(id_1, op_1)\})$
  end if
  $id_2 = $ newID(O)
  if $index = 0$ **then**
    $ref = list$
  else
    $ref = $ idxKey$_{E,\, L}(list, list, index - 1)$
  end if
  $O := O \cup \{(id_2, \mathsf{InsertAfter}(ref))\}$
  $id_3 = $ newID(O)
  return $O \cup$
    $\{(id_3, \mathsf{Assign}(list, id_2, id_1, \emptyset))\}$
end function

function REMOVELISTINDEX(O, $list$, $index$)
  $(E, L) = [\![O]\!]$
  $id_1 = $ newID(O)
  $key = $ idxKey$_{E,\, L}(list, list, index)$
  $prev = \{id \mid \exists v.\ (id, list, key, v) \in E\}$
  return $O \cup$
    $\{(id_1, \mathsf{Remove}(list, key, prev))\}$
end function

example. *Tagged union* types are introduced with the **datatype** keyword, with *constructors* of these types usually written with an initial upper case letter.

In Isabelle/HOL's term language we write $t :: \tau$ for a *type ascription*, constraining the type of the term t to the type τ. We write $\lambda x.\, t$ for an anonymous function mapping an argument x to $t(x)$, and write the application of term t with function type to an argument u as $t\ u$. Terms of list type are introduced using one of two constructors: the empty list $[\,]$ or 'nil', and the infix 'cons' operator $\#$, which prepends an element to an existing list. We use $[t_1, \ldots, t_n]$ as syntactic sugar for a list literal, and $xs\,@\,ys$ to express the concatenation (appending) of two lists xs and ys. We write $\{\,\}$ for the empty set, and use usual mathematical notation for set union, disjunction, membership tests, and so on: $t \cup u$, $t \cap u$, and $x \in t$. We write $t \longrightarrow s$ for logical implication between formulae (terms of type *bool*). Strictly speaking Isabelle is a logical framework, providing a weak meta-logic within which object logics are embedded, including the Isabelle/HOL object logic that we use in this work. Accordingly, the implication arrow of Isabelle's meta-logic, $t \Longrightarrow u$, is required in certain contexts over the object-logic implication arrow, $t \longrightarrow s$, already introduced. However, for purposes of an intuitive understanding, the two forms of implication can be regarded as equivalent by the reader, with the requirement to use one over the other merely being an implementation detail of Isabelle itself. We will sometimes use the shorthand $[\![H_1;\ \ldots;\ H_n]\!] \Longrightarrow C$ instead of iterated meta-logic implications, i.e., $H_1 \Longrightarrow \ldots \Longrightarrow H_n \Longrightarrow C$.

B.2   Definitions and theorems.

New non-recursive definitions are entered into Isabelle's global context using the **definition** keyword. Recursive functions are defined using the **fun** keyword, and support *pattern matching* on their arguments. All functions are total, and therefore every recursive function must be provably terminating. All termination proofs in this work are generated automatically by Isabelle itself.

Inductive relations are defined with the **inductive** keyword. For example, the definition

> **inductive** *only-fives* :: *nat list* $\Rightarrow$ *bool* **where**
>   *only-fives* $[\,]$ $\mid$
> $[\![$ *only-fives xs* $]\!] \Longrightarrow$ *only-fives* $(5\#xs)$

introduces a new constant *only-fives* of type *nat list* $\Rightarrow$ *bool*. The two clauses in the body of the definition enumerate the conditions under which *only-fives xs* is true, for arbitrary *xs*: firstly, *only-fives* is true for the empty list; and secondly, if you know that *only-fives xs* is true for some *xs*, then you can deduce that *only-fives* $(5\#xs)$ (i.e., *xs* prefixed with the number 5) is also true. Moreover, *only-fives xs* is true in no other circumstances—it is the *smallest* relation closed under the rules defining it. In short, the clauses above state that *only-fives xs* holds exactly in the case where *xs* is a (potentially empty) list containing only repeated copies of the natural number 5.

Lemmas, theorems, and corollaries can be asserted using the **lemma**, **theorem**, and **corollary** keywords, respectively. There is no semantic difference between these keywords in Isabelle, and they serve only to mark certain results as especially important (or unimportant) for human readers. For example,

> **theorem** *only-fives-concat*:
>   **assumes** *only-fives xs* **and** *only-fives ys*
>   **shows** *only-fives (xs @ys)*

conjectures that if *xs* and *ys* are both lists of fives, then their concatenation *xs @ ys* is also a list of fives. Isabelle then requires that this claim be proved by using one of its proof methods, for example by induction. Some proofs can be automated, whilst others require the user to provide explicit reasoning steps. The theorem is assigned a name, here *only-fives-concat*, so that it may be referenced in later proofs.

C   Statements of Mechanised Proofs

In this appendix we provide a copy of the Isabelle/HOL definitions and proof statements that support the central claims in the paper. For space reasons, the actual proofs are omitted; the full formal proof development can be found in the Isabelle Archive of Formal Proofs [33]. The source code is available at `https://github.com/trvedata/opsets`.

C.1   Abstract OpSet

In this section, we define a general-purpose OpSet abstraction that is not specific to any one particular datatype. An OpSet is a set of (ID, operation) pairs with an associated total order on IDs (represented here with the *linorder* typeclass), and satisfying the following properties:

1. The ID is unique (that is, if any two pairs in the set have the same ID, then their operation is also the same).
2. If the operation references the IDs of any other operations, those referenced IDs are less than that of the operation itself, according to the total order on IDs. To avoid assuming anything about the structure of operations here, we use a function *deps* that returns the set of dependent IDs for a given operation. This requirement is a weak expression of causality: an operation can only depend on causally prior operations, and by making the total order on IDs a linear extension of the causal order, we can easily ensure that any referenced IDs are less than that of the operation itself.
3. The OpSet is finite (but we do not assume any particular maximum size).

We define it as follows in Isabelle:[1]

> **locale** *opset* =
>   **fixes** *opset* :: (*'oid*::{*linorder*} × *'oper*) *set*
>   **and** *deps* :: *'oper* ⇒ *'oid set*
>   **assumes** *unique-oid*: (*oid*, *op1*) ∈ *opset* ⟹ (*oid*, *op2*) ∈ *opset* ⟹ *op1* = *op2*
>   **and** *ref-older*: (*oid*, *oper*) ∈ *opset* ⟹ *ref* ∈ *deps oper* ⟹ *ref* < *oid*
>   **and** *finite-opset*: *finite opset*

We prove that any subset of an OpSet is also a valid OpSet. This is the case because, although an operation can depend on causally prior operations, the OpSet does not require those prior operations to actually exist. This weak assumption makes the OpSet model more general and simplifies reasoning about OpSets.

----

[1] In programming terms, a *locale* (or 'local theory') may be thought of as an interface with associated laws that implementations must obey. When showing that an implementation matches this interface, one must also show that the implementation satisfies all assumed laws of the locale. Moreover, locales can be extended with new assumed facts and fixed constants to form a hierarchy, and definitions and theorems may be defined and declared within a locale and made available to all of its implementations. See the standard Isabelle tutorial material, as well as [30] and [25] for a more detailed explanation of locales.

> **lemma** *opset-subset*:
>   **assumes** *opset Y deps*
>   **and** $X \subseteq Y$
>   **shows** *opset X deps*

C.1.1   The *spec-ops* **predicate**

The *spec-ops* predicate describes a list of (ID, operation) pairs that corresponds to the linearisation of an OpSet, and which we use for sequentially interpreting the OpSet. A list satisfies *spec-ops* iff it is sorted in ascending order of IDs, if the IDs are unique, and if every operation's dependencies have lower IDs than the operation itself. A list is implicitly finite in Isabelle/HOL.

> **definition** *spec-ops* :: $('oid::\{linorder\} \times 'oper)$ *list* $\Rightarrow$ $('oper \Rightarrow 'oid\ set)$ $\Rightarrow$ *bool*
> **where**
>   *spec-ops ops deps* $\equiv$ (*sorted* (*map fst ops*) $\wedge$ *distinct* (*map fst ops*) $\wedge$
>   ($\forall$ *oid oper ref*. (*oid, oper*) $\in$ *set ops* $\wedge$ *ref* $\in$ *deps oper* $\longrightarrow$ *ref* $<$ *oid*))

We prove that for any given OpSet, a *spec-ops* linearisation exists:

> **lemma** *spec-ops-exists*:
>   **assumes** *opset ops deps*
>   **shows** $\exists$ *op-list. set op-list = ops* $\wedge$ *spec-ops op-list deps*

Conversely, for any given *spec-ops* list, the set of pairs in the list is an OpSet:

> **lemma** *spec-ops-is-opset*:
>   **assumes** *spec-ops op-list deps*
>   **shows** *opset* (*set op-list*) *deps*

C.1.2   The *crdt-ops* **predicate**

Like *spec-ops*, the *crdt-ops* predicate describes the linearisation of an OpSet into a list. Like *spec-ops*, it requires IDs to be unique. However, its other properties are different: *crdt-ops* does not require operations to appear in sorted order, but instead, whenever any operation references the ID of a prior operation, that prior operation must appear previously in the *crdt-ops* list. Thus, the order of operations is partially constrained: operations must appear in causal order, but concurrent operations can be ordered arbitrarily.

This list describes the operation sequence in the order it is typically applied to an operation-based CRDT. Applying operations in the order they appear in *crdt-ops* requires that concurrent operations commute. For any *crdt-ops* operation sequence, there is a permutation that satisfies the *spec-ops* predicate. Thus, to check whether a CRDT satisfies its sequential specification, we can prove that interpreting any *crdt-ops* operation sequence with the commutative operation interpretation results in the same end result as interpreting the *spec-ops* permutation of that operation sequence with the sequential operation interpretation.

inductive *crdt-ops* :: $('oid::\{linorder\} \times 'oper)$ *list* $\Rightarrow$ $('oper \Rightarrow 'oid \ set) \Rightarrow bool$
where
  crdt-ops $[]$ *deps* |
  $[\![crdt\text{-}ops \ xs \ deps;$
    $oid \notin set \ (map \ fst \ xs);$
    $\forall \ ref \in deps \ oper. \ ref \in set \ (map \ fst \ xs) \wedge ref < oid$
  $]\!] \implies crdt\text{-}ops \ (xs \ @ \ [(oid, \ oper)]) \ deps$

C.2 Specifying List Insertion

In this section we consider only list insertion. We model an insertion operation as a pair (*ID, ref*), where *ref* is either *None* (signifying an insertion at the head of the list) or *Some r* (an insertion immediately after a reference element with ID *r*). If the reference element does not exist, the operation does nothing.

    We provide two different definitions of the interpretation function for list insertion: *insert-spec* and *insert-alt*. The *insert-alt* definition matches the paper, while *insert-spec* uses the Isabelle/HOL list datatype, making it more suitable for formal reasoning. In section C.2.2 we prove that the two definitions are in fact equivalent.

fun *insert-spec* :: $'oid \ list \Rightarrow ('oid \times 'oid \ option) \Rightarrow 'oid \ list$
where
  insert-spec xs   (oid, *None*)   = $oid\#xs$ |
  insert-spec $[]$ (oid, -)   = $[]$ |
  insert-spec ($x\#xs$) (oid, *Some ref*) =
    (*if* $x = ref$ *then* $x \ \# \ oid \ \# \ xs$
             else $x \ \# \ (insert\text{-}spec \ xs \ (oid, \ Some \ ref)))$

fun *insert-alt* :: $('oid \times 'oid \ option) \ set \Rightarrow ('oid \times 'oid) \Rightarrow ('oid \times 'oid \ option) \ set$
where
  insert-alt list-rel (oid, *ref*) = (
    if $\exists \ n. \ (ref, \ n) \in list\text{-}rel$
    then $\{(p, \ n) \in list\text{-}rel. \ p \neq ref\} \cup \{(ref, \ Some \ oid)\} \cup$
        $\{(i, \ n). \ i = oid \wedge (ref, \ n) \in list\text{-}rel\}$
    else list-rel)

interp-ins is the sequential interpretation of a set of insertion operations. It starts with an empty list as initial state, and then applies the operations from left to right.

definition *interp-ins* :: $('oid \times 'oid \ option) \ list \Rightarrow 'oid \ list$ **where**
  interp-ins ops $\equiv$ *foldl insert-spec* $[]$ *ops*

C.2.1 The *insert-ops* predicate

We now specialise the definitions from section C.1 for list insertion. *insert-opset* is an opset consisting only of insertion operations, and *insert-ops* is the specialisation of the *spec-ops* predicate for insertion operations.

> **locale** *insert-opset = opset opset set-option*
>   **for** *opset* :: (*'oid*::{*linorder*} × *'oid option*) *set*
>
> **definition** *insert-ops* :: (*'oid*::{*linorder*} × *'oid option*) *list* ⇒ *bool* **where**
>   *insert-ops list ≡ spec-ops list set-option*

C.2.2   Equivalence of the two definitions of insertion

We now prove that the two definitions of insertion, *insert-spec* and *insert-alt*, are equivalent.
First we define how to derive the successor relation from an Isabelle list. This relation contains
(*id, None*) if *id* is the last element of the list, and (*id1, id2*) if *id1* is immediately followed by
id2 in the list.

> **fun** *succ-rel* :: *'oid list* ⇒ (*'oid* × *'oid option*) *set*
> **where**
>   *succ-rel* [] = {} |
>   *succ-rel* [*head*] = {(*head, None*)} |
>   *succ-rel* (*head#x#xs*) = {(*head, Some x*)} ∪ *succ-rel* (*x#xs*)

interp-alt is the equivalent of interp-ins, but using *insert-alt* instead of *insert-spec*. To match
the paper, it uses a distinct head element to refer to the beginning of the list.

> **definition** *interp-alt* :: *'oid* ⇒ (*'oid* × *'oid option*) *list* ⇒ (*'oid* × *'oid option*) *set*
> **where**
>   *interp-alt head ops ≡ foldl insert-alt* {(*head, None*)}
>   (*map* (λ*x. case x of*
>   (*oid, None*)   ⇒ (*oid, head*) |
>   (*oid, Some ref*) ⇒ (*oid, ref*))
>   *ops*)

We can now prove that *insert-spec* and *insert-alt* are equivalent:

> **theorem** *insert-alt-equivalent*:
>   **assumes** *insert-ops ops*
>   **and** *head ∉ fst ' set ops*
>   **and** ⋀*r. Some r ∈ snd ' set ops ⟹ r ≠ head*
>   **shows** *succ-rel* (*head # interp-ins ops*) = *interp-alt head ops*

C.3   No Interleaving

The predicate *insert-seq start ops* is true iff *ops* is a list of insertion operations that begins
by inserting after *start*, and then continues by placing each subsequent insertion directly
after its predecessor. This definition models the sequential insertion of text at a particular
place in a text document.

> **inductive** *insert-seq* :: *'oid option* ⇒ (*'oid* × *'oid option*) *list* ⇒ *bool* **where**
>   *insert-seq start* [(*oid, start*)] |
>   ⟦*insert-seq start* (*list* @ [(*prev, ref*)])⟧
>   ⟹ *insert-seq start* (*list* @ [(*prev, ref*), (*oid, Some prev*)])

Consider an execution that contains two distinct insertion sequences, *xs* and *ys*, that both begin at the same initial position *start*. We prove that, provided the starting element exists, the two insertion sequences are not interleaved. That is, in the final list order, either all insertions by *xs* appear before all insertions by *ys*, or vice versa.

> **theorem** *no-interleaving*:
>   **assumes** *insert-ops ops*
>   **and** *insert-seq start xs* **and** *insert-ops xs*
>   **and** *insert-seq start ys* **and** *insert-ops ys*
>   **and** *set xs* ⊆ *set ops* **and** *set ys* ⊆ *set ops*
>   **and** *distinct* (*map fst xs* @ *map fst ys*)
>   **and** *start* = *None* ∨ (∃ *r*. *start* = *Some r* ∧ *r* ∈ *set* (*interp-ins ops*))
>   **shows** (∀ *x* ∈ *set* (*map fst xs*). ∀ *y* ∈ *set* (*map fst ys*). *list-order ops x y*) ∨
>   (∀ *x* ∈ *set* (*map fst xs*). ∀ *y* ∈ *set* (*map fst ys*). *list-order ops y x*)

For completeness, we also prove what happens if there are two insertion sequences, *xs* and *ys*, but their reference element *start* does not exist. In this failure case, none of the insertions in *xs* or *ys* take effect.

> **theorem** *missing-start-no-insertion*:
>   **assumes** *insert-ops ops*
>   **and** *insert-seq* (*Some start*) *xs* **and** *insert-ops xs*
>   **and** *insert-seq* (*Some start*) *ys* **and** *insert-ops ys*
>   **and** *set xs* ⊆ *set ops* **and** *set ys* ⊆ *set ops*
>   **and** *start* ∉ *set* (*interp-ins ops*)
>   **shows** ∀ *x* ∈ *set* (*map fst xs*) ∪ *set* (*map fst ys*). *x* ∉ *set* (*interp-ins ops*)

C.4 The Replicated Growable Array (RGA)

The RGA algorithm [49] is a replicated list (or collaborative text-editing) algorithm. In this section we prove that RGA satisfies our list specification. The Isabelle/HOL definition of RGA in this section is based on our prior work on formally verifying CRDTs [21, 20].

fun *insert-body* :: *'oid*::{*linorder*} *list* $\Rightarrow$ *'oid* $\Rightarrow$ *'oid list* **where**
  insert-body []   *e* = [*e*] |
  insert-body (*x* # *xs*) *e* =
    (*if x* < *e then e* # *x* # *xs*
          else x # *insert-body xs e*)

fun *insert-rga* :: *'oid*::{*linorder*} *list* $\Rightarrow$ (*'oid* $\times$ *'oid option*) $\Rightarrow$ *'oid list* **where**
  insert-rga xs   (*e, None*)   = *insert-body xs e* |
  insert-rga []   (*e, Some i*) = [] |
  insert-rga (*x* # *xs*) (*e, Some i*) =
    (*if x* = *i then*
      x # *insert-body xs e*
    else
      x # *insert-rga xs* (*e, Some i*))

definition *interp-rga* :: (*'oid*::{*linorder*} $\times$ *'oid option*) *list* $\Rightarrow$ *'oid list* **where**
  interp-rga ops $\equiv$ *foldl insert-rga* [] *ops*

definition *rga-ops* :: (*'oid*::{*linorder*} $\times$ *'oid option*) *list* $\Rightarrow$ *bool* **where**
  rga-ops list $\equiv$ *crdt-ops list set-option*

We can then prove that RGA satisfies our list specification:

theorem *rga-meets-spec*:
  assumes *rga-ops xs*
  shows $\exists$ *ys. set ys* = *set xs* $\wedge$ *insert-ops ys* $\wedge$ *interp-ins ys* = *interp-rga xs*

C.5   Relationship to Strong List Specification

In this section we show that our list specification is stronger than the $\mathcal{A}_{\mathsf{strong}}$ specification of collaborative text editing by Attiya et al. [4]. We do this by showing that the OpSet interpretation of any set of insertion and deletion operations satisfies all of the consistency criteria that constitute the $\mathcal{A}_{\mathsf{strong}}$ specification.

Attiya et al.'s specification is as follows [4]:

An abstract execution $A = (H, \mathsf{vis})$ belongs to the *strong list specification* $\mathcal{A}_{\mathsf{strong}}$ if and only if there is a relation $\mathsf{lo} \subseteq \mathsf{elems}(A) \times \mathsf{elems}(A)$, called the *list order*, such that:

1. Each event $e = do(op, w) \in H$ returns a sequence of elements $w = a_0 \ldots a_{n-1}$, where $a_i \in \mathsf{elems}(A)$, such that

  a. w contains exactly the elements visible to e that have been inserted, but not deleted:

$$\forall a.\, a \in w \iff (do(\mathsf{ins}(a, \_), \_) \leq_{\mathsf{vis}} e) \wedge \neg(do(\mathsf{del}(a), \_) \leq_{\mathsf{vis}} e).$$

  b. The order of the elements is consistent with the list order:

$$\forall i, j.\, (i < j) \implies (a_i, a_j) \in \mathsf{lo}.$$

      c. Elements are inserted at the specified position: if $op = \mathsf{ins}(a, k)$, then $a = a_{\min\{k,\, n-1\}}$.

  2. The list order lo is transitive, irreflexive and total, and thus determines the order of all insert operations in the execution.

This specification considers only insertion and deletion operations, but no assignment. Moreover, it considers only a single list object, not a graph of composable objects like in our paper. Thus, we prove the relationship to $\mathcal{A}_{\mathsf{strong}}$ using a simplified interpretation function that defines only insertion and deletion on a single list.

We first define a datatype for list operations, with two constructors: *Insert ref val*, and *Delete ref*. For insertion, the *ref* argument is the ID of the existing element after which we want to insert, or *None* to insert at the head of the list. The *val* argument is an arbitrary value to associate with the list element. For deletion, the *ref* argument is the ID of the existing list element to delete.

> **datatype** (′*oid*, ′*val*) *list-op* =
>   *Insert* ′*oid option* ′*val* |
>   *Delete* ′*oid*

When interpreting operations, the result is a pair (*list*, *vals*). The *list* contains the IDs of list elements in the correct order (equivalent to the list relation in the paper), and *vals* is a mapping from list element IDs to values (equivalent to the element relation in the paper).

Insertion delegates to the previously defined *insert-spec* interpretation function. Deleting a list element removes it from *vals*.

> **fun** *interp-op* :: (′*oid list* × (′*oid* ⇀ ′*val*)) ⇒ (′*oid* × (′*oid*, ′*val*) *list-op*)
>   ⇒ (′*oid list* × (′*oid* ⇀ ′*val*)) **where**
>   *interp-op* (*list*, *vals*) (*oid*, *Insert ref val*) = (*insert-spec list* (*oid*, *ref*), *vals*(*oid* ↦ *val*)) |
>   *interp-op* (*list*, *vals*) (*oid*, *Delete ref*   ) = (*list*, *vals*(*ref* := *None*))

> **definition** *interp-ops* :: (′*oid* × (′*oid*, ′*val*) *list-op*) *list* ⇒ (′*oid list* × (′*oid* ⇀ ′*val*))
> **where**
>   *interp-ops ops* ≡ *foldl interp-op* ([], *Map.empty*) *ops*

list-order ops x y holds iff, after interpreting the list of operations *ops*, the list element with ID x appears before the list element with ID y in the resulting list.

> **definition** *list-order* :: (′*oid* × (′*oid*, ′*val*) *list-op*) *list* ⇒ ′*oid* ⇒ ′*oid* ⇒ *bool* **where**
>   *list-order ops x y* ≡ ∃ *xs ys zs*. *fst* (*interp-ops ops*) = *xs* @ [*x*] @ *ys* @ [*y*] @ *zs*

The *make-insert* function generates a new operation for insertion into a given index in a given list. The exclamation mark is Isabelle's list subscript operator.

> **fun** *make-insert* :: ′*oid list* ⇒ ′*val* ⇒ *nat* ⇒ (′*oid*, ′*val*) *list-op* **where**
>   *make-insert list val 0*   = *Insert None val* |
>   *make-insert* []   *val k*   = *Insert None val* |
>   *make-insert list val* (*Suc k*) = *Insert* (*Some* (*list* ! (*min k* (*length list* − *1*)))) *val*

The *list-ops* predicate is a specialisation of *spec-ops* to the *list-op* datatype: it describes a list of (ID, operation) pairs that is sorted by ID, and can thus be used for the sequential interpretation of the OpSet.

fun *list-op-deps* :: (*'oid*, *'val*) *list-op* ⇒ *'oid set* **where**
   list-op-deps (*Insert* (*Some ref*) -) = {*ref*} |
   list-op-deps (*Insert   None   * -) = {} |
   list-op-deps (*Delete   ref   *) = {*ref*}

locale *list-opset* = *opset opset list-op-deps*
   for *opset* :: (*'oid*::{*linorder*} × (*'oid*, *'val*) *list-op*) *set*

definition *list-ops* :: (*'oid*::{*linorder*} × (*'oid*, *'val*) *list-op*) *list* ⇒ *bool* **where**
   list-ops ops ≡ *spec-ops ops list-op-deps*

C.5.1   Satisfying all conditions of $\mathcal{A}_{\mathsf{strong}}$

Part 1(a) of Attiya et al.'s specification states that whenever the list is observed, the elements of the list are exactly those that have been inserted but not deleted. $\mathcal{A}_{\mathsf{strong}}$ uses the visibility relation $\leq_{\mathsf{vis}}$ to capture the operations known to a node at some arbitrary point in the execution; in the OpSet model, we can simply prove the theorem for an arbitrary OpSet, since the contents of the OpSet at a particular time on a particular node correspond exactly to the set of operations known to that node at that time.

theorem *inserted-but-not-deleted*:
   assumes *list-ops ops*
      and *interp-ops ops* = (*list*, *vals*)
   shows *a* ∈ *dom* (*vals*) ⟷ (∃ *ref val*. (*a*, *Insert ref val*) ∈ *set ops*) ∧
                                  (∄ *i*. (*i*, *Delete a*) ∈ *set ops*)

Part 1(b) states that whenever the list is observed, the order of list elements is consistent with the global list order. We can define the global list order simply as the list order that arises from interpreting the OpSet containing all operations in the entire execution. Then, at any point in the execution, the OpSet is some subset of the set of all operations.

We can then rephrase condition 1(b) as follows: whenever list element x appears before list element y in the interpretation of *some-ops*, then for any OpSet *all-ops* that is a superset of *some-ops*, x must also appear before y in the interpretation of *all-ops*. In other words, adding more operations to the OpSet does not change the relative order of any existing list elements.

theorem *list-order-consistent*:
   assumes *list-ops some-ops* **and** *list-ops all-ops*
      and *set some-ops* ⊆ *set all-ops*
      and *list-order some-ops x y*
   shows *list-order all-ops x y*

Part 1(c) states that inserted elements appear at the specified position: that is, immediately after an insertion of *oid* at index k, the list index k does indeed contain *oid* (provided that k is less than the length of the list). We prove this property below.

theorem *correct-position-insert*:
   assumes *list-ops* (*ops* @ [(*oid*, *ins*)])
      and *ins* = *make-insert* (*fst* (*interp-ops ops*)) *val k*
      and *list* = *fst* (*interp-ops* (*ops* @ [(*oid*, *ins*)]))
   shows *list* ! (*min k* (*length list* − *1*)) = *oid*

Part 2 states that the list order relation must be transitive, irreflexive, and total. These three properties are straightforward to prove, using our definition of the *list-order* predicate.

> **theorem** *list-order-trans*:
>   **assumes** *list-ops ops*
>   **and** *list-order ops x y*
>   **and** *list-order ops y z*
>   **shows** *list-order ops x z*

> **theorem** *list-order-irrefl*:
>   **assumes** *list-ops ops*
>   **shows** ¬ *list-order ops x x*

> **theorem** *list-order-total*:
>   **assumes** *list-ops ops*
>   **and** $x \in set\ (fst\ (interp\text{-}ops\ ops))$
>   **and** $y \in set\ (fst\ (interp\text{-}ops\ ops))$
>   **and** $x \neq y$
>   **shows** *list-order ops x y* ∨ *list-order ops y x*