Strong Eventual Consistency of the Collaborative Editing Framework WOOT

Emin Karayel and Edgar Gonzàlez

Google, Mountain View

December 14, 2021

Abstract

Commutative Replicated Data Types (CRDTs) are a promising new class of data structures for large-scale shared mutable content in applications that only require eventual consistency. The WithOut Operational Transforms (WOOT) framework is a CRDT for collaborative text editing introduced by Oster et al. (CSCW 2006) for which the eventual consistency property was verified only for a bounded model to date. We contribute a formal proof for WOOTs strong eventual consistency.

Contents

1 Introduction **2**

2 Related Work **3**

3 Preliminary Notes **5**
- 3.1 Algorithms in Isabelle 5

4 The WOOT Framework **6**
- 4.1 Symbol Identifiers 7
  - 4.1.1 Extended Identifiers 8
- 4.2 Messages 8
- 4.3 States 9
- 4.4 Basic Algorithms 9
- 4.5 Edit Operations 10
- 4.6 Integration algorithm 12
- 4.7 Network Model 14

5 Formalized Proof **18**
5.1 Definition of Ψ . . . 19
5.2 Sorting . . . 22
5.3 Consistency of sets of WOOT Messages . . . 23
5.4 Create Consistent . . . 25
5.5 Termination Proof for *integrate-insert* . . . 27
5.6 Integrate Commutes . . . 28
5.7 Strong Convergence . . . 32

6 Strong Eventual Consistency **34**

7 Code generation **35**

8 Proof Outline **35**
8.1 Sort Keys . . . 37
8.2 Induction . . . 39

9 Example **40**

1 Introduction

A *Replicated (Abstract) Data Type (RDT)* consists of "*multiple copies of a shared Abstract Data Type (ADT) replicated over distributed sites, [which] provides a set of primitive operation types corresponding to that of normal ADTs, concealing details for consistency maintenance*" [22]. RDTs can be classified as *state-based* or *operation-based* depending on whether full states (e.g., a document's text) or only the operations performed on them (e.g., character insertions and deletions) are exchanged among replicas. Operation-based RDTs are *commutative* when the integration of any two concurrent operations on any reachable replica state commutes [24].

Commutative (Operation-Based) Replicated Data Types (CRDTs[1] from now on) enable sharing mutable content with optimistic replication—ensuring high-availability, responsive interaction, and eventual consistency without consensus-based concurrency control [13]. They are used in highly scalable robust distributed applications [26, 3].

An RDT is *eventually consistent* when, if after some point in time no further updates are made at any replica, all replicas eventually converge to equivalent states. It is *strongly eventually consistent* when it is eventually

[1]Note that other authors like Shapiro et al. [24] use CmRDT to refer to Commutative RDTs, with CRDT standing for *Conflict-free RDTs*.

consistent and, whenever any two peers have seen the same set of updates (in possibly different order), they reach equivalent states immediately [24].

The WithOut Operational Transforms (WOOT) Framework [19] was the first proposed CRDT for collaborative text editing [2]. It has been implemented as part of several OSS projects [4, 6, 8, 16]. However, the eventual consistency of WOOT has only been verified for a bounded model [19, 18]. A formal proof of WOOTs consistency can rigorously establish that there is no complex counter-example not identified by model checking.

The contribution of this work is one such proof that the WOOT Framework is strongly eventually consistent. Its central idea is the association of a value from a dense totally ordered space to each inserted (and potentially deleted) character, using a recursive definition with respect to the acyclic graph induced by the predecessor and successor relation of the characters. We then show that the strings in each peer remain sorted with respect to that value, i.e., that the values form a sort key for W-characters.[2] This resolves the conjecture posed by Oster et al. [18, conjecture 1] and is also the key lemma to establish that the WOOT Framework has the strong eventual consistency property.

After reviewing related work in the following section, we formalize the WOOT Framework as a distributed application in Section 4. We follow with the complete eventual consistency proof in Section 5 and summarize the established results in Section 6. In Section 8 we given overview of the proof and follow up with a conrete formalized example in Section 9.

The presentation is structured such that all the definitions necessary to review the established results in Section 6 are part of Section 4. This means it is possible to skip Section 5 entirely.

2 Related Work

The first collaborative text editing tools were based on operational transformations (OT), and introduced by Ellis and Gibbs [5]. The basic idea behind OT-based frameworks is to adjust edit operations, based on the effects of previously executed concurrent operations. For instance, in Figure 1a, peer B can execute the message received from peer A without correction, but peer A needs to transform the one received from peer B to reach the same state.

Proving the correctness of OT-based frameworks is error-prone and requires complicated case coverage [14, 17]. Counter-examples have been found in most OT algorithms [22][7, section 8.2].

[2]Note that the values themselves do not have to be actually computed, during the execution of the framework. Their existence and compatibility with the integration algorithm forms a witness for the consistency proof we are presenting.

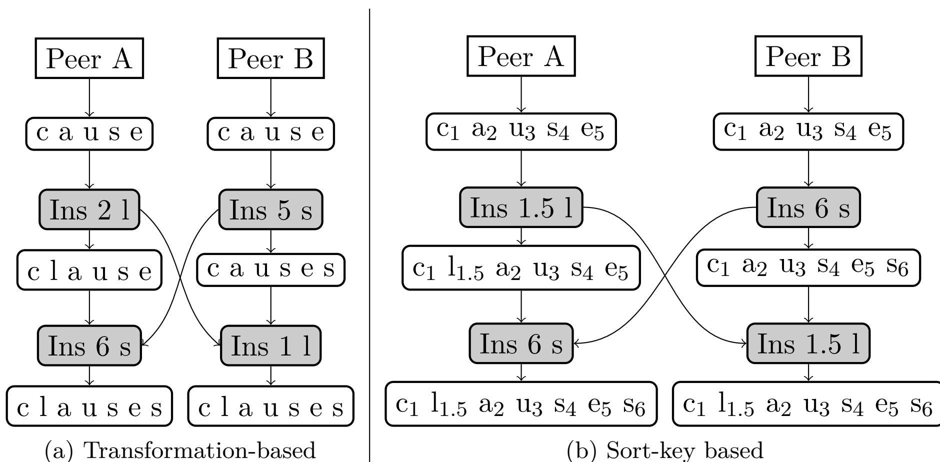

Figure 1: Collaborative text editing

LSEQ [15], LOGOOT [26] and TreeDoc [20] are CRDTs that create and send sort keys for symbols (e.g., 1.5 and 6 in Figure 1b). These keys can then be directly used to order them, without requiring any transformations, and are drawn from a dense totally ordered space. In the figure rational numbers were chosen for simplicity, but more commonly lexicographically ordered sequences are used.[3] The consistency property of these frameworks can be established easily. However, the space required per sort key potentially grows linearly with the count of edit operations. In LSEQ, a randomized allocation strategy for new identifiers is used to reduce the key growth, based on empirically determined edit patterns—but in the worst-case the size of the keys will still grow linearly with the count of insert operations. Preguica et al. [20] propose a solution for this problem using regular rebalancing operations. However, this can only be done using a consensus-based mechanism, which is only possible when the number of participating peers is small.

A benefit of LSEQ, LOGOOT, and TreeDoc is that deleted symbols can be garbage-collected (though delete messages may have to be kept in a buffer if the corresponding insertion message has not arrived at a peer), in contrast to the WOOT Framework, where deleted symbols (tombstones) cannot be removed.

Replicated Growable Arrays (RGAs) are another data structure for collaborative editing, introduced by Roh et al. [22]. Contrary to the previous approaches, the identifiers associated to the symbols are not sort keys, but are instead ordered consistently with the happened-before relation. A peer sends the identifier of the symbol immediately preceeding the new symbol

[3] In addition, peers draw sort keys from disjoint (but dense) subsets to avoid concurrently choosing the same sort key.

at the time it was created and the actual identifier associated to the new symbol. The integration algorithm starts by finding the preceeding symbol and skipping following symbols with a larger identifier before placing the new symbol. The authors provide a mathematical eventual consistency proof. Recently, Gomes et. al. [7] also formalized the eventual consistency property of RGAs using Isabelle/HOL.

In addition to the original design of WOOT by Oster et al. [19], a number of extensions have also been proposed. For instance, Weiss et al. [25] propose a line-based version WOOTO, and Ahmed-Nacer et al. [1] introduce a second extension WOOTH which improves performance by using hash tables. The latter compare their implementation in benchmarks against LOGOOT, RGA, and an OT algorithm.

To the best of our knowledge there are no publications that further expand on the correctness of the WOOT Framework. The fact that the general convergence proof is missing is also mentioned by Kumawat and Khunteta [11, Section 3.10].

3 Preliminary Notes

3.1 Algorithms in Isabelle

theory *ErrorMonad*
  imports
    Certification-Monads.Error-Monad
begin

Isabelle's functions are mathematical functions and not necessarily algorithms. For example, it is possible to define a non-constructible function:

fun *non-constructible-function* **where**
  $\textit{non-constructible-function}\ f = (\textit{if}\ (\exists n.\ f\ n = 0)\ \textit{then}\ 1\ \textit{else}\ 0)$

and even prove properties of them, like for example:

$$\textit{non-constructible-function}\ (\lambda x.\ \textit{Suc}\ 0) = 0$$

In addition to that, some native functions in Isabelle are under-defined, e.g., $[]\ !\ 1$. But it is still possible to show lemmas about these undefined values, e.g.: $[]\ !\ 1 = [a, b]\ !\ 3$. While it is possible to define a notion of algorithm in Isabelle [9], we think that this is not necessary for the purpose of this formalization, since the reader needs to verify that the formalized functions correctly model the algorithms described by Oster et al. [19] anyway. However, we show that Isabelle can generate code for the functions, indicating that non-constructible terms are not being used within the algorithms.

type-synonym *error* = *String.literal*

```
fun assert :: bool ⇒ error + unit
  where
    assert False = error (STR ''Assertion failed.'') |
    assert True = return ()

fun fromSingleton :: 'a list ⇒ error + 'a
  where
    fromSingleton [] = error (STR ''Expected list of length 1'') |
    fromSingleton (x # []) = return x |
    fromSingleton (x # y # ys) = error (STR ''Expected list of length 1'')
```

Moreover, we use the error monad—modelled using the *sum* type—and build wrappers around partially defined Isabelle functions such that the evaluation of undefined terms and violation of invariants expected by the algorithms result in error values.

We are able to show that all operations succeed without reaching unexpected states during the execution of the framework.

end

4 The WOOT Framework

```
theory Data
  imports Main Datatype-Order-Generator.Order-Generator
begin
```

Following the presentation by Oster et al. [19] we describe the WOOT framework as an operation-based CRDT [24].

In WOOT, the shared data type is a string over an alphabet $'\Sigma$. Each peer starts with a prescribed initial state representing the empty string. Users can perform two types of edit operations on the string at their peer:

- Insert a new character.
- Delete an existing character.

Whenever a user performs one of these operations, their peer will create an update message (see Section 4.5), integrate it immediately into its state, and send it to every other peer.

An update message created at a peer may depend on at most two of the previously integrated messages at that peer. A message cannot be delivered to a peer if its antecedents have not been delivered to it yet. In Section 4.7 we describe a few possible methods to implement this requirement, as there is a trade-off between causal consistency and scalability.

Once delivered to a remote peer, an update message will be integrated to the peers' state. The integration algorithm for an update message is the

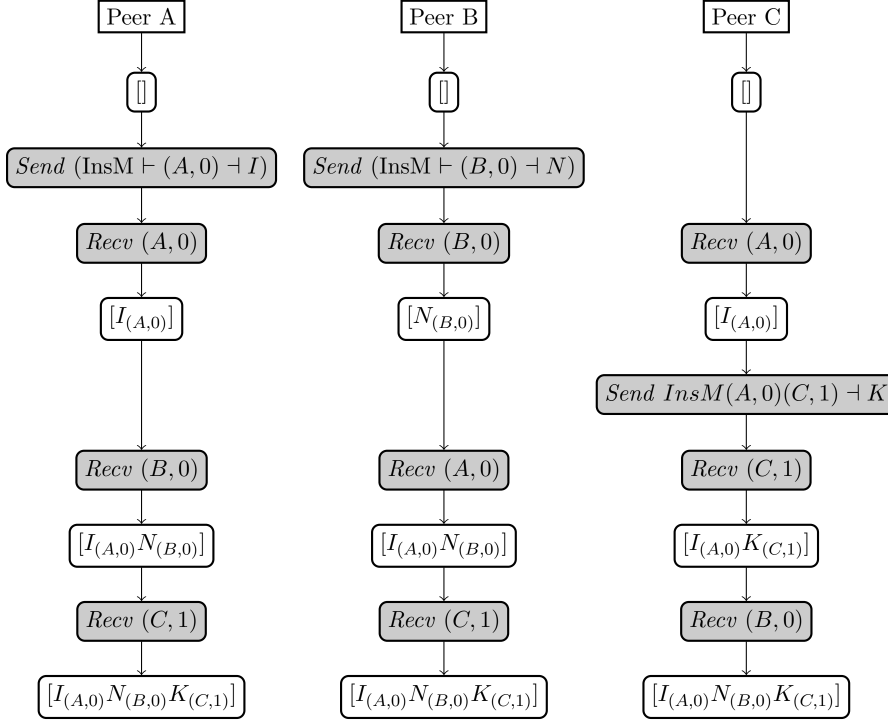

Figure 2: Example session with 3 peers. Each peer creates an update message and sends a copy of it to the other two peers. Each peer integrates the messages in a different order. The white rounded boxes represent states, for brevity we only show the W-character's symbol and identifier. Although a W-character's data structure stores the identifiers of its predecessor and successor from its original creation event. The gray round boxes represent events, we abbreviate the reception events, with the identifier of the W-character, although the peer receives the full insert message.

same whether the message originated at the same or at a different peer (see Section 4.6).

The interaction of the WOOT Framework can be visualized using a space-time diagram [10]. An example session between 3 peers is shown in Figure 2. Note that, each peer sees the edit operations in a different order.

4.1 Symbol Identifiers

The WOOT Framework requires a unique identifier for each insert operation, which it keeps associated with the inserted symbol. The identifier may not be used for another insertion operation on the same or any other peer. Moreover the set of identifiers must be endowed with a total linear order. We will denote the set of identifiers by $'\mathcal{I} :: \textit{linorder}$.

Note that the order on the identifiers is not directly used as a global order over the inserted symbols, in contrast to the sort-key based approaches: LSEQ, LOGOOT, or TreeDoc. In particular, this means we do not require the identifier space to be dense.

In the modelling in Section 4.7, we will use the pair consisting of a unique identifier for the peer and the count of messages integrated or sent by that peer, with the lexicographic order induced by the Cartesian product of the peer identifier and the counter.

It is however possible to use other methods to generate unique identifiers, as long as the above requirements are fulfilled.

4.1.1 Extended Identifiers

datatype $'\mathcal{I}$ *extended*
  = *Begin* ($\vdash$)
  | *InString* $'\mathcal{I}$ (($1[\![-]\!]$))
  | *End* ($\dashv$)
derive *linorder extended*

We embed the set of identifiers in an extension containing two additional elements denoting the smallest (resp. largest) element of the extension. The order of identifiers with respect to each other is preserved. The extended set is used in the corner cases, where a W-character is inserted at the beginning or end of the string - and there is no preceeding resp. succeeding W-character to reference. See also the following section.

4.2 Messages

datatype $('\mathcal{I}, '\Sigma)$ *insert-message* =
  InsertMessage (P:$'\mathcal{I}$ *extended*) (I:$'\mathcal{I}$) (S:$'\mathcal{I}$ *extended*) (Σ:$'\Sigma$)

datatype $'\mathcal{I}$ *delete-message* = *DeleteMessage* $'\mathcal{I}$

datatype $('\mathcal{I}, '\Sigma)$ *message* =
  Insert $('\mathcal{I}, '\Sigma)$ *insert-message* |
  Delete $'\mathcal{I}$ *delete-message*

Two kinds of update messages are exchanged in the WOOT Framework, indicating respectively an insertion or a deletion of a character. Thus the set of messages is a sum type *message*.

An insert message *Insert m* has the following four components:

- $P\ m$ and $S\ m$ denote the identifiers of the character immediately preceding (resp. succeeding) the character at the time of its insertion. The special value $\vdash$ (resp. $\dashv$) indicates that there was no such character, i.e., that it was inserted at the beginning (resp. end) of the string.

- $I\ m$ denotes the unique identifier associated to the character (as described in Subsection 4.1).

- $\Sigma\ m$ denotes the inserted character.

4.3 States

type-synonym $('\mathcal{I}, '\Sigma)$ *woot-character* $=$ $('\mathcal{I}, '\Sigma$ *option*$)$ *insert-message*

A W-character w is the representation of an inserted character in the state of a peer. It has the same semantics and notation for its components as an insert message, with the difference that $\Sigma\ w$ can be *Some* σ denoting an inserted character, or *None* if the character has already been deleted. Because of this overlap in semantics, we define the type of W-characters as a type synonym.

The state of a peer is then a string of W-characters $s :: ('\mathcal{I}, '\Sigma)$ *woot-character list*. The initial state is the empty string $[]$. The string the user sees is the sequence of symbols omitting *None*s, i.e., the sequence: $[\sigma.\ Some\ \sigma \leftarrow map\ \Sigma\ s]$.

fun *to-woot-char* $:: ('\mathcal{I}, '\Sigma)$ *insert-message* $\Rightarrow ('\mathcal{I}, '\Sigma)$ *woot-character*
 where
 to-woot-char $(InsertMessage\ p\ i\ s\ \sigma) = InsertMessage\ p\ i\ s\ (Some\ \sigma)$

An insert message can be converted into a W-character by applying *Some* to the symbol component.

end

4.4 Basic Algorithms

theory *BasicAlgorithms*
 imports *Data ErrorMonad*
begin

In this section, we introduce preliminary definitions and functions, required by the integration and edit algorithms in the following sections.

definition *ext-ids* $:: ('\mathcal{I}, '\Sigma)$ *woot-character list* $\Rightarrow '\mathcal{I}$ *extended list*
 where *ext-ids* $s = \vdash\#(map\ (\lambda x.\ [\![\ I\ x\ ]\!])\ s)@[\dashv]$

The function *ext-ids* returns the set of extended identifiers in a string s, including the beginning and end markers $\vdash$ and $\dashv$.

fun *idx* $:: ('\mathcal{I}, '\Sigma)$ *woot-character list* $\Rightarrow '\mathcal{I}$ *extended* $\Rightarrow$ *error* $+$ *nat*
 where
 $idx\ s\ i = fromSingleton\ (filter\ (\lambda j.\ (ext\text{-}ids\ s\ !\ j = i))\ [0..{<}(length\ (ext\text{-}ids\ s))])$

The function *idx* returns the index in w of a W-character with a given identifier i:

- If the identifier i occurs exactly once in the string then $idx\ s\ [\![i]\!] = Inr\ (j + 1)$ where $I\ (s\ !\ j) = i$, otherwise $idx\ s\ [\![i]\!]$ will be an error.
- $idx\ s \vdash = Inr\ 0$ and $idx\ s \dashv = Inr\ (length\ w + 1)$.

```
fun nth :: ('I, 'Σ) woot-character list ⇒ nat ⇒ error + ('I, 'Σ) woot-character
  where
    nth s 0 = error (STR ''Index has to be >= 1.'') |
    nth s (Suc k) = (
      if k < (length s) then
        return (s ! k)
      else
        error (STR ''Index has to be <= length s''))
```

The function *nth* returns the W-character at a given index in s. The first character has the index 1.

```
fun list-update ::
  ('I, 'Σ) woot-character list ⇒ nat ⇒ ('I, 'Σ) woot-character ⇒
  error + ('I, 'Σ) woot-character list
  where
    list-update s (Suc k) v = (
        if k < length s then
          return (List.list-update s k v)
        else
          error (STR ''Illegal arguments.'')) |
    list-update - 0 - = error (STR ''Illegal arguments.'')
```

The function *list-update* substitutes the W-character at the index k in s with the W-character v. As before, we use the convention of using the index 1 for the first character.

end

4.5 Edit Operations

```
theory CreateAlgorithms
  imports BasicAlgorithms
begin

fun is-visible :: ('I, 'Σ) woot-character ⇒ bool
  where is-visible (InsertMessage - - - s) = (s ≠ None)

fun nth-visible :: ('I, 'Σ) woot-character list ⇒ nat ⇒ error + 'I extended
  where
    nth-visible s k = (let v = ext-ids (filter is-visible s) in
      if k < length v then
        return (v ! k)
      else
        error (STR ''Argument k out of bounds.''))
```

Let l be the count of visible symbols in s. The function *nth-visible s n*:

- Returns the identifier of the n-th visible element in s if $1 \leq n \leq l$.
- Returns $\vdash$ if $n = 0$, and $\dashv$ if $n = l + 1$.
- Returns an error otherwise.

Note that, with this definition, the first visible character in the string has the index 1.

Algorithms *create-insert* and *create-delete* detail the process by which messages are created in response to a user action.

```
fun from-non-extended :: 'I extended ⇒ error + 'I
  where
    from-non-extended ⟦ i ⟧ = Inr i |
    from-non-extended - = error (STR ''Expected InString'')

fun create-insert ::
  ('I, 'Σ) woot-character list ⇒ nat ⇒ 'Σ ⇒ 'I ⇒ error + ('I, 'Σ) message
  where create-insert s n σ' i =
    do {
      p ← nth-visible s n;
      q ← nth-visible s (n + 1);
      return (Insert (InsertMessage p i q σ'))
    }
```

In particular, when a user inserts a character σ' between visible position n and its successor of the string of a peer with state s, *create-insert* starts by retrieving the identifiers p of the last visible character before n in w (or $\vdash$ if no such character exists) and q of the first visible one after n (or $\dashv$).

It then broadcasts the message *Insert* (*InsertMessage* p i q σ') with the new identifier i.

```
fun create-delete :: ('I, 'Σ) woot-character list ⇒ nat ⇒ error + ('I, 'Σ) message
  where create-delete s n =
    do {
      m ← nth-visible s n;
      i ← from-non-extended m;
      return (Delete (DeleteMessage i))
    }
```

When the user deletes the visible character at position n, *create-delete* retrieves the identifier i of the n'th visible character in s and broadcasts the message *Delete* (*DeleteMessage* i).

In both cases the message will be integrated into the peer's own state, with the same algorithm that integrates messages received from other peers.

end

4.6 Integration algorithm

In this section we describe the algorithm to integrate a received message into a peers' state.

```
theory IntegrateAlgorithm
  imports BasicAlgorithms Data
begin

fun fromSome :: 'a option ⇒ error + 'a
  where
    fromSome (Some x) = return x |
    fromSome None = error (STR ''Expected Some'')

lemma fromSome-ok-simp [simp]: (fromSome x = Inr y) = (x = Some y)
  ⟨proof⟩

fun substr :: 'a list ⇒ nat ⇒ nat ⇒ 'a list where
  substr s l u = take (u − (Suc l)) (drop l s)

fun concurrent ::
  ('a, 's) woot-character list
  ⇒ nat
  ⇒ nat
  ⇒ ('a, 's) woot-character
  ⇒ error + ('a extended list)
  where
    concurrent s l u w =
      do {
        p-pos ← idx s (P w);
        s-pos ← idx s (S w);
        return (if (p-pos ≤ l ∧ s-pos ≥ u) then [⟦I w⟧] else [])
      }

function integrate-insert
  where
    integrate-insert m w p s =
      do {
        l ← idx w p;
        u ← idx w s;
        assert (l < u);
        if Suc l = u then
          return ((take l w)@[to-woot-char m]@(drop l w))
        else do {
          d ← mapM (concurrent w l u) (substr w l u);
          assert (concat d ≠ []);
          (p', s') ← fromSome (find ((λx.⟦I m⟧ < x ∨ x = s) ∘ snd)
                       (zip (p#concat d) (concat d@[s])));
          integrate-insert m w p' s'
        }
```

}
⟨*proof*⟩

fun *integrate-delete* ::
('*a* :: *linorder*) *delete-message*
⇒ ('*a*, '*s*) *woot-character list*
⇒ *error* + ('*a*, '*s*) *woot-character list*
where
integrate-delete (*DeleteMessage i*) *s* =
do {
$k \leftarrow$ *idx s* ⟦*i*⟧;
$w \leftarrow$ *nth s k*;
list-update s k
(*case w of* (*InsertMessage p i u -*) ⇒ *InsertMessage p i u None*)
}

fun *integrate* ::
('*a*, '*s*) *woot-character list*
⇒ ('*a* :: *linorder*, '*s*) *message*
⇒ *error* + ('*a*, '*s*) *woot-character list*
where
integrate s (*Insert m*) = *integrate-insert m s* (*P m*) (*S m*) |
integrate s (*Delete m*) = *integrate-delete m s*

Algorithm *integrate* describes the main function that is called when a new message m has to be integrated into the state s of a peer. It is called both when m was generated locally or received from another peer. Note that we require that the antecedant messages have already been integrated. See also Section 4.7 for the delivery assumptions that ensure this requirement.

Algorithm *integrate-delete* describes the procedure to integrate a delete message: *DeleteMessage i*. The algorithm just replaces the symbol of the W-character with identifier i with the value *None*. It is not possible to entirely remove a W-character if it is deleted, since there might be unreceived insertion messages that depend on its position.

Algorithm *integrate-insert* describes the procedure to integrate an insert message: m = *InsertMessage p i s* σ. Since insertion operations can happen concurrently and the order of message delivery is not fixed, it can happen that a remote peer receiving m finds multiple possible insertion points between the predecessor p and successor s that were recorded when the message was generated. An example of this situation is the conflict between *InsertMessage* ⊢ (*A*, *0*) ⊣ *CHR ''I''* and *InsertMessage* ⊢ (*B*, *0*) ⊣ *CHR ''N''* in Figure 2.

A first attempt to resolve this would be to insert the W-characters by choosing an insertion point using the order induced by their identifiers to achieve a consistent ordering. But this method fails in some cases: a counter-example was found by Oster et al. [19, section 2].

The solution introduced by the authors of WOOT is to restrict the identifier comparison to the set of W-characters in the range *substr l u s* whose predecessor and successor are outside of the possible range, i.e. $idx\ s\ (P\ w) \leq l$ and $idx\ s\ (S\ w) \geq u$.

New narrowed bounds are selected by finding the first W-character within that restricted set with an identifier strictly larger than the identifier of the new W-character.

This leads to a narrowed range where the found character forms an upper bound and its immediately preceeding character the lower bound. The method is applied recursively until the insertion point is uniquely determined.

Note that the fact that this strategy leads to a consistent ordering has only been verified for a bounded model. One of the contributions of this paper is to provide a complete proof for it.

end

4.7 Network Model

In the past subsections, we described the algorithms each peer uses to integrate received messages and broadcast new messages when an edit operation has been made on that peer.

In this section, we model the WOOT Framework as a distributed application and set the basis for the consistency properties, we want to establish.

We assume a finite set of peers starting with the same initial state of an empty W-string, each peer reaches a finite set of subsequent states, caused by the integration of received (or locally generated messages). A message is always generated from a concrete state of a peer, using the algorithms described in Section 4.5. Moreover, we assume that the same message will only be delivered once to a peer. Finally, we assume that the happened before relation, formed by

- Subsequent states of the same peer
- States following the reception and states that were the generation sites

do not contain loops. (Equivalently that the transitive closure of the relation is a strict partial order.)

The latter is a standard assumption in the modelling of distributed systems (compare e.g. [21, Chapter 6.1]) effectively implied by the fact that there are no physical causal loops.

Additionally, we assume that a message will be only received by a peer, when the antecedent messages have already been received by the peer. This is a somewhat technical assumption to simplify the description of the system.

In a practical implementation a peer would buffer the set of messages that cannot yet be integrated. Note that this assumption is automatically implied if causal delivery is assumed.

We establish two properties under the above assumptions

- The integration algorithm never fails.
- Two peers having received the same set of messages will be in the same state.

The model assumptions are derived from Gomes et al.[7] and Shapiro et al.[23] with minor modifications required for WOOT.

theory *DistributedExecution*
  imports *IntegrateAlgorithm CreateAlgorithms HOL−Library.Product-Lexorder*
begin

type-synonym $'p$ *event-id* = $'p \times$ *nat*

datatype $('p,'s)$ *event* =
  Send ($'p$ *event-id*, $'s$) *message* |
  Receive $'p$ *event-id* ($'p$ *event-id*, $'s$) *message*

The type variable $'p$ denotes a unique identifier identifying a peer. We model each peer's history as a finite sequence of events, where each event is either the reception or broadcast of a message. The index of the event in a peer's history and its identifier form a pair uniquely identifying an event in a distributed execution of the framework. In the case of a reception, *Receive* s m indicated the reception of the message m sent at event s.

In the following we introduce the locale *dist-execution-preliminary* from which the *dist-execution* locale will inherit. The reason for the introduction of two locales is technical - in particular, it is not possible to interleave definitions and assumptions within the definition of a locale. The preliminary locale only introduces the assumption that the set of participating peers is finite.

locale *dist-execution-preliminary* =
  fixes *events* :: ($'p$ :: *linorder*) $\Rightarrow$ ($'p$,$'s$) *event list*
  — We introduce a locale fixing the sequence of events per peer.

  assumes *fin-peers*: *finite* (*UNIV* :: $'p$ *set*)
  — We are assuming a finite set of peers.

begin

fun *is-valid-event-id*
  where *is-valid-event-id* $(i,j) = (j <$ *length* (*events* i))

```
fun event-pred
  where event-pred (i,j) p  = (is-valid-event-id (i,j) ∧ p (events i ! j))

fun event-at
  where event-at i m = event-pred i ((=) m)

fun is-reception
  where
    is-reception i j = event-pred j (λe. case e of Receive s - ⇒ s = i | - ⇒ False)

fun happened-immediately-before where
  happened-immediately-before i j = (
     is-valid-event-id i ∧
     is-valid-event-id j ∧
     ((fst i = fst j ∧ Suc (snd i) = snd j) ∨ is-reception i j))
```

The *happened-immediately-before* describes immediate causal precedence:

- An events causally precedes the following event on the same peer.
- A message broadcast event causally precedes the reception event of it.

The transitive closure of this relation is the famous happened before relation introduced by Lamport[12].

In the *dist-execution* we will assume that the relation is acyclic - which implies that the transitive closure *happened-immediately-before*$^{++}$ is a strict partial order.

Each peer passes through a sequence of states, which may change after receiving a message. We denote the initial state of peer p as $(p, 0)$ and the state after event (p, i) as $(p, i + 1)$. Note that there is one more state per peer than events, since we are count both the initial and terminal state of a peer.

```
fun is-valid-state-id
  where is-valid-state-id (i,j) = (j ≤ length (events i))

fun received-messages
  where
    received-messages (i,j) = [m. (Receive - m) ← (take j (events i))]

fun state where state i = foldM integrate [] (received-messages i)
```

Everytime a peer receives a message its state is updated by integrating the message. The function *state* returns the state for a given state id.

end

The function *deps* computes the identifiers a message depends on.

```
fun extended-to-set :: 'a extended ⇒ 'a set
```

where
$extended\text{-}to\text{-}set$ $[\![i]\!] = \{i\}$ |
$extended\text{-}to\text{-}set$ - $= \{\}$

fun $deps :: ('id, 's)\ message \Rightarrow 'id\ set$
where
$deps\ (Insert\ (InsertMessage\ l\ \text{-}\ u\ \text{-})) = extended\text{-}to\text{-}set\ l \cup extended\text{-}to\text{-}set\ u$ |
$deps\ (Delete\ (DeleteMessage\ i)) = \{i\}$

locale $dist\text{-}execution = dist\text{-}execution\text{-}preliminary\ +$
assumes *no-data-corruption*:
$\bigwedge i\ s\ m.\ event\text{-}at\ i\ (Receive\ s\ m) \Longrightarrow event\text{-}at\ s\ (Send\ m)$
— A received message must also have been actually broadcast. Note that, we do not assume that a broadcast message will be received by all peers, similar to the modelling made by [7, Section 5.2].

assumes *at-most-once*:
$\bigwedge i\ j\ s\ m.$
$event\text{-}at\ i\ (Receive\ s\ m) \Longrightarrow$
$event\text{-}at\ j\ (Receive\ s\ m) \Longrightarrow$
$fst\ i = fst\ j \Longrightarrow i = j$
— A peer will never receive the same message twice. Note that this is something that can be easily implemented in the application layer, if the underlying transport mechanism does not guarantee it.

assumes *acyclic-happened-before*:
$acyclicP\ happened\text{-}immediately\text{-}before$
— The immediate causal precendence relation is acyclic, which implies that its closure, the *happened before* relation is a strict partial order.

assumes *semantic-causal-delivery*:
$\bigwedge m\ s\ i\ j\ i'.\ event\text{-}at\ (i,j)\ (Receive\ s\ m) \Longrightarrow i' \in deps\ m \Longrightarrow$
$\exists\, s'\ j'\ m'.\ event\text{-}at\ (i,j')\ (Receive\ s'\ (Insert\ m')) \wedge j' < j \wedge I\ m' = i'$
— A message will only be delivered to a peer, if its antecedents have already been delivered. (See beginning of this Section for the reason of this assumption).

assumes *send-correct*:
$\bigwedge m\ i.\ event\text{-}at\ i\ (Send\ m) \Longrightarrow$
$(\exists\, n\ \sigma.\ return\ m = state\ i \ggg= (\lambda s.\ create\text{-}insert\ s\ n\ \sigma\ i)) \vee$
$(\exists\, n.\ return\ m = state\ i \ggg= (\lambda s.\ create\text{-}delete\ s\ n))$
— A peer broadcasts messages by running the *create-insert* or *create-delete* algorithm on its current state. In the case of an insertion the new character is assigned the event id as its identifier. Note that, it would be possible to assume, a different choice for allocating unique identifiers to new W-characters. We choose the event id since it is automatically unique.

begin

Based on the assumptions above we show in Section 6:

- *Progress*: All reached states *state i* will be successful, i.e., the algorithm *integrate* terminates and does not fail.
- *Strong Eventual Consistency*: Any pair of states *state i* and *state j* which have been reached after receiving the same set of messages, i.e., *set* (*received-messages i*) = *set* (*received-messages j*) will be equal.

end

end

5 Formalized Proof

```
theory SortKeys
  imports Data HOL−Library.List-Lexorder HOL−Library.Product-Lexorder
begin

datatype sort-dir =
  Left |
  Right
derive linorder sort-dir

lemma sort-dir-less-def [simp]: (x < y) = (x = Left ∧ y = Right)
  ⟨proof⟩

datatype 'a sort-key =
  NonFinal ('a × sort-dir) 'a sort-key |
  Final 'a

type-synonym 'id position = 'id sort-key extended

fun embed-dir where embed-dir (x,Left) = (x, 0) | embed-dir (x,Right) = (x, Suc
(Suc 0))

lemma embed-dir-inj [simp]: (embed-dir x = embed-dir y) = (x = y)
  ⟨proof⟩

lemma embed-dir-mono [simp]: (embed-dir x < embed-dir y) = (x < y)
  ⟨proof⟩

fun sort-key-embedding :: 'a sort-key ⇒ ('a × nat) list
  where
   sort-key-embedding (NonFinal x y) = embed-dir x#(sort-key-embedding y) |
   sort-key-embedding (Final i) = [(i, Suc 0)]

lemma sort-key-embedding-injective:
  sort-key-embedding x = sort-key-embedding y ⟹ x = y
  ⟨proof⟩
```

instantiation *sort-key* :: (*ord*) *ord*
begin
definition *sort-key-less-eq-def* [*simp*]:
 $(x :: ('a :: ord)\ sort\text{-}key) \leq y \longleftrightarrow$
 $(sort\text{-}key\text{-}embedding\ x \leq sort\text{-}key\text{-}embedding\ y)$

definition *sort-key-less-def* [*simp*]:
 $(x :: ('a :: ord)\ sort\text{-}key) < y \longleftrightarrow$
 $(sort\text{-}key\text{-}embedding\ x < sort\text{-}key\text{-}embedding\ y)$

instance ⟨*proof*⟩
end

instantiation *sort-key* :: (*order*) *order*
begin
instance ⟨*proof*⟩
end

instantiation *sort-key* :: (*linorder*) *linorder*
begin
instance ⟨*proof*⟩
end

end

5.1 Definition of Ψ

theory *Psi*
 imports *SortKeys HOL−Eisbach.Eisbach*
begin

fun *extended-size* :: (*'a sort-key*) *extended* ⇒ *nat*
 where
 $extended\text{-}size\ [\![x]\!] = size\ x\ |$
 $extended\text{-}size\ \text{-} = 0$

lemma *extended-simps* [*simp*]:
 $(\vdash < x) = (x \neq \vdash)$
 $([\![x']\!] < [\![y']\!]) = (x' < y')$
 $[\![x']\!] < \dashv$
 $\neg([\![x']\!] < \vdash)$
 $\neg(\dashv < x)$
 $\vdash \leq x$
 $([\![x']\!] \leq [\![y']\!]) = ((x' :: 'a :: linorder) \leq y')$
 $x \leq \dashv$
 $\neg([\![x']\!] \leq \vdash)$
 $(\dashv \leq x) = (x = \dashv)$
 ⟨*proof*⟩

```
fun int-size where int-size (l,u) = max (extended-size l) (extended-size u)

lemma position-cases:
  assumes ⋀ y z. x = ⟦NonFinal (y,Left) z⟧ ⟹ p
  assumes ⋀ y z. x = ⟦NonFinal (y,Right) z⟧ ⟹ p
  assumes ⋀ y. x = ⟦Final y⟧ ⟹ p
  assumes x = ⊢ ⟹ p
  assumes x = ⊣ ⟹ p
  shows p
  ⟨proof⟩

fun derive-pos ::
  ('a :: linorder) × sort-dir ⇒ 'a sort-key extended ⇒ 'a sort-key extended
  where
    derive-pos h ⟦NonFinal x y⟧ =
      (if h < x then ⊣ else (if x < h then ⊢ else ⟦y⟧)) |
    derive-pos h ⟦Final x⟧ =
      (if fst h < x ∨ fst h = x ∧ snd h = Left then ⊣ else ⊢) |
    derive-pos - ⊢ = ⊢ |
    derive-pos - ⊣ = ⊣

lemma derive-pos-mono: x ≤ y ⟹ derive-pos h x ≤ derive-pos h y
  ⟨proof⟩

fun γ :: ('a :: linorder) position ⇒ sort-dir ⇒ 'a × sort-dir
  where
    γ ⟦NonFinal x y⟧ - = x |
    γ ⟦Final x⟧ d = (x,d) |
    γ ⊢ - = undefined |
    γ ⊣ - = undefined

fun derive-left  where
  derive-left (l, u) = (derive-pos (γ l Right) l, derive-pos (γ l Right) u)

fun derive-right where
  derive-right (l, u) = (derive-pos (γ u Left) l, derive-pos (γ u Left) u)

fun is-interval where is-interval (l,u) = (l < u)

fun elem where elem x (l,u) = (l < x ∧ x < u)

fun subset where subset (l,u) (l',u') = (l' ≤ l ∧ u ≤ u')

method interval-split for x :: ('a :: linorder) position × 'a position =
  (case-tac [!] x,
   rule-tac [!] position-cases [where x=fst x],
   rule-tac [!] position-cases [where x=snd x])

lemma derive-size:
```

$[\![Final\ i]\!] \le fst\ x \wedge is\text{-}interval\ x \Longrightarrow int\text{-}size\ (derive\text{-}left\ x) < int\text{-}size\ x$
$snd\ x \le [\![Final\ i]\!] \wedge is\text{-}interval\ x \Longrightarrow int\text{-}size\ (derive\text{-}right\ x) < int\text{-}size\ x$
⟨*proof*⟩

lemma *derive-interval*:
$[\![Final\ i]\!] \le fst\ x \Longrightarrow is\text{-}interval\ x \Longrightarrow is\text{-}interval\ (derive\text{-}left\ x)$
$snd\ x \le [\![Final\ i]\!] \Longrightarrow is\text{-}interval\ x \Longrightarrow is\text{-}interval\ (derive\text{-}right\ x)$
⟨*proof*⟩

function $\Psi :: ('a :: linorder)\ position \times 'a\ position \Rightarrow 'a \Rightarrow 'a\ sort\text{-}key$
where
$\Psi\ (l,u)\ i = Final\ i$
if $l < [\![Final\ i]\!] \wedge [\![Final\ i]\!] < u$ |
$\Psi\ (l,u)\ i = NonFinal\ (\gamma\ l\ Right)\ (\Psi\ (derive\text{-}left\ (l,u))\ i)$
if $[\![Final\ i]\!] \le l \wedge l < u$ |
$\Psi\ (l,u)\ i = NonFinal\ (\gamma\ u\ Left)\ (\Psi\ (derive\text{-}right\ (l,u))\ i)$
if $u \le [\![Final\ i]\!] \wedge l < u$ |
$\Psi\ (l,u)\ i = undefined$ **if** $u \le l$
⟨*proof*⟩

termination
⟨*proof*⟩

proposition *psi-elem*: $is\text{-}interval\ x \Longrightarrow elem\ [\![\Psi\ x\ i]\!]\ x$
⟨*proof*⟩

proposition *psi-mono*:
assumes $i1 < i2$
shows $is\text{-}interval\ x \Longrightarrow \Psi\ x\ i1 < \Psi\ x\ i2$
⟨*proof*⟩

proposition *psi-narrow*:
$elem\ [\![\Psi\ x'\ i]\!]\ x \Longrightarrow subset\ x\ x' \Longrightarrow \Psi\ x'\ i = \Psi\ x\ i$
⟨*proof*⟩

definition $preserve\text{-}order :: 'a :: linorder \Rightarrow 'a \Rightarrow 'b :: linorder \Rightarrow 'b \Rightarrow bool$
where $preserve\text{-}order\ x\ y\ u\ v \equiv (x < y \longrightarrow u < v) \wedge (x > y \longrightarrow u > v)$

proposition *psi-preserve-order*:
fixes $l\ l'\ u\ u'\ i\ i'$
assumes $elem\ [\![\Psi\ (l,\ u)\ i]\!]\ (l',u')$
assumes $elem\ [\![\Psi\ (l',\ u')\ i']\!]\ (l,\ u)$
shows $preserve\text{-}order\ i\ i'\ [\![\Psi\ (l,u)\ i]\!]\ [\![\Psi\ (l',\ u')\ i']\!]$
⟨*proof*⟩

end

5.2 Sorting

Some preliminary lemmas about sorting.

theory *Sorting*
  imports *Main HOL.List HOL−Library.Sublist*
begin

lemma *insort*:
  assumes $Suc\ l < length\ s$
  assumes $s\ !\ l < (v :: 'a :: linorder)$
  assumes $s\ !\ (l+1) > v$
  assumes *sorted-wrt* $(<)\ s$
  shows *sorted-wrt* $(<)\ ((take\ (Suc\ l)\ s)@v\#(drop\ (Suc\ l)\ s))$
⟨*proof*⟩

lemma *sorted-wrt-irrefl-distinct*:
  assumes *irreflp r*
  shows *sorted-wrt r xs* $\longrightarrow$ *distinct xs*
  ⟨*proof*⟩

lemma *sort-set-unique-h*:
  assumes *irreflp r* $\wedge$ *transp r*
  assumes $set\ (x\#xs) = set\ (y\#ys)$
  assumes $\forall\, z \in set\ xs.\ r\ x\ z$
  assumes $\forall\, z \in set\ ys.\ r\ y\ z$
  shows $x = y \wedge set\ xs = set\ ys$
  ⟨*proof*⟩

lemma *sort-set-unique-rel*:
  assumes *irreflp r* $\wedge$ *transp r*
  assumes $set\ x = set\ y$
  assumes *sorted-wrt r x*
  assumes *sorted-wrt r y*
  shows $x = y$
⟨*proof*⟩

lemma *sort-set-unique*:
  assumes $set\ x = set\ y$
  assumes *sorted-wrt* $(<)\ (map\ (f :: ('a \Rightarrow ('b :: linorder)))\ \ x)$
  assumes *sorted-wrt* $(<)\ (map\ f\ y)$
  shows $x = y$
  ⟨*proof*⟩

If two sequences contain the same element and strictly increasing with respect.

lemma *subseq-imp-sorted*:
  assumes *subseq s t*
  assumes *sorted-wrt p t*
  shows *sorted-wrt p s*

⟨*proof*⟩

If a sequence t is sorted with respect to a relation p then a subsequence will be as well.

fun *to-ord* **where** $\textit{to-ord}\ r\ x\ y = (\neg(r^{**}\ y\ x))$

lemma *trancl-idemp*: $r^{++++}\ x\ y = r^{++}\ x\ y$
⟨*proof*⟩

lemma *top-sort*:
fixes *rp*
assumes *acyclicP r*
shows $\textit{finite}\ s \longrightarrow (\exists\, l.\ \textit{set}\ l = s \wedge \textit{sorted-wrt}\ (\textit{to-ord}\ r)\ l \wedge \textit{distinct}\ l)$
⟨*proof*⟩

lemma *top-sort-eff*:
assumes $\textit{irreflp}\ p^{++}$
assumes *sorted-wrt* (*to-ord p*) *x*
assumes $i < \textit{length}\ x$
assumes $j < \textit{length}\ x$
assumes $(p^{++}\ (x\ !\ i)\ (x\ !\ j))$
shows $i < j$
⟨*proof*⟩

end

5.3 Consistency of sets of WOOT Messages

theory *Consistency*
imports *SortKeys Psi Sorting DistributedExecution*
begin

definition *insert-messages* :: $('\mathcal{I},\ '\Sigma)$ *message set* $\Rightarrow$ $('\mathcal{I},\ '\Sigma)$ *insert-message set*
where $\textit{insert-messages}\ M = \{x.\ \textit{Insert}\ x \in M\}$

lemma *insert-insert-message*:
$\textit{insert-messages}\ (M \cup \{\textit{Insert}\ m\}) = \textit{insert-messages}\ M \cup \{m\}$
⟨*proof*⟩

definition *delete-messages* :: $('a,\ 's)$ *message set* $\Rightarrow$ $'a$ *delete-message set*
where $\textit{delete-messages}\ M = \{x.\ \textit{Delete}\ x \in M\}$

fun *depends-on* **where** $\textit{depends-on}\ M\ x\ y = (x \in M \wedge y \in M \wedge I\ x \in \textit{deps}\ (\textit{Insert}\ y))$

definition *a-conditions* ::
$(('a :: \textit{linorder}),\ 's)$ *insert-message set* $\Rightarrow$ $('a\ \textit{extended} \Rightarrow\ 'a\ \textit{position}) \Rightarrow \textit{bool}$
where *a-conditions M a* = (
$a \vdash\ <\ a \dashv\ \wedge$

$(\forall\, m.\ m \in M \longrightarrow a\ (P\ m) < a\ (S\ m) \wedge$
$a\ [\![I\ m]\!] = [\![\Psi\ (a\ (P\ m),\ a\ (S\ m))\ (I\ m)]\!])$

definition *consistent* :: $('a$ :: *linorder*, $'s)$ *message set* $\Rightarrow$ *bool*
where *consistent* $M \equiv$
inj-on I (*insert-messages* M) $\wedge$
$(\bigcup$ (*deps* ‘ M) $\subseteq$ (I ‘ *insert-messages* M)) $\wedge$
wfP (*depends-on* (*insert-messages* M)) $\wedge$
$(\exists\, a.$ *a-conditions* (*insert-messages* M) $a)$

lemma *consistent-subset*:
assumes *consistent* N
assumes $M \subseteq N$
assumes $\bigcup$ (*deps* ‘ M) $\subseteq$ (I ‘ *insert-messages* M)
shows *consistent* M
⟨*proof*⟩

lemma *pred-is-dep*: $P\ m = [\![\ i\ ]\!] \longrightarrow i \in$ *deps* (*Insert* m)
⟨*proof*⟩

lemma *succ-is-dep*: $S\ m = [\![\ i\ ]\!] \longrightarrow i \in$ *deps* (*Insert* m)
⟨*proof*⟩

lemma *a-subset*:
fixes $M\ N\ a$
assumes $M \subseteq N$
assumes *a-conditions* (*insert-messages* N) a
shows *a-conditions* (*insert-messages* M) a
⟨*proof*⟩

definition *delete-maybe* :: $'\mathcal{I} \Rightarrow ('\mathcal{I},\ '\Sigma)$ *message set* $\Rightarrow\ '\Sigma \Rightarrow\ '\Sigma$ *option* **where**
delete-maybe $i\ D\ s$ = (*if Delete* (*DeleteMessage* i) $\in D$ *then None else Some* s)

definition *to-woot-character* ::
$('\mathcal{I},\ '\Sigma)$ *message set* $\Rightarrow ('\mathcal{I},\ '\Sigma)$ *insert-message* $\Rightarrow ('\mathcal{I},\ '\Sigma)$ *woot-character*
where
to-woot-character $D\ m$ = (
case m *of*
(*InsertMessage* $l\ i\ u\ s$) $\Rightarrow$ *InsertMessage* $l\ i\ u$ (*delete-maybe* $i\ D\ s$))

lemma *to-woot-character-keeps-i* [*simp*]: I (*to-woot-character* $M\ m$) = $I\ m$
⟨*proof*⟩

lemma *to-woot-character-keeps-i-lifted* [*simp*]:
I ‘ *to-woot-character* M ‘ X = I ‘ X
⟨*proof*⟩

lemma *to-woot-character-keeps-P* [*simp*]: P (*to-woot-character* $M\ m$) = $P\ m$
⟨*proof*⟩

lemma *to-woot-character-keeps-S* [*simp*]: *S* (*to-woot-character M m*) = *S m*
⟨*proof*⟩

lemma *to-woot-character-insert-no-eff*:
to-woot-character (*insert* (*Insert m*) *M*) = *to-woot-character M*
⟨*proof*⟩

definition *is-associated-string* ::
('*a*, '*s*) *message set* ⇒ ('*a* :: *linorder*, '*s*) *woot-character list* ⇒ *bool*
where *is-associated-string M s* ≡ (
consistent M ∧
set s = *to-woot-character M* ' (*insert-messages M*) ∧
(∀ *a*. *a-conditions* (*insert-messages M*) *a* ⟶
sorted-wrt (<) (*map a* (*ext-ids s*))))

fun *is-certified-associated-string* **where**
is-certified-associated-string M (*Inr v*) = *is-associated-string M v* |
is-certified-associated-string M (*Inl* -) = *False*

lemma *associated-string-unique*:
assumes *is-associated-string M s*
assumes *is-associated-string M t*
shows *s* = *t*
⟨*proof*⟩

lemma *is-certified-associated-string-unique*:
assumes *is-certified-associated-string M s*
assumes *is-certified-associated-string M t*
shows *s* = *t*
⟨*proof*⟩

lemma *empty-consistent*: *consistent* {}
⟨*proof*⟩

lemma *empty-associated*: *is-associated-string* {} []
⟨*proof*⟩

The empty set of messages is consistent and the associated string is the empty string.

end

5.4 Create Consistent

theory *CreateConsistent*
imports *CreateAlgorithms Consistency*
begin

lemma *nth-visible-inc'*:

assumes *sorted-wrt* (<) (*map a* (*ext-ids s*))
assumes *nth-visible s n* = *Inr i*
assumes *nth-visible s* (*Suc n*) = *Inr j*
shows *a i* < *a j*
⟨*proof*⟩

lemma *nth-visible-eff*:
assumes *nth-visible s n* = *Inr i*
shows *extended-to-set i* ⊆ *I* ‘ *set s*
⟨*proof*⟩

lemma *subset-mono*:
assumes *N* ⊆ *M*
shows *I* ‘ *insert-messages N* ⊆ *I* ‘ *insert-messages M*
⟨*proof*⟩

lemma *deps-insert*:
assumes ⋃ (*deps* ‘ *M*) ⊆ (*I* ‘ *insert-messages M*)
assumes *deps m* ⊆ *I* ‘ *insert-messages M*
shows ⋃ (*deps* ‘ (*M* ∪ {*m*})) ⊆ (*I* ‘ *insert-messages* (*M* ∪ {*m*}))
⟨*proof*⟩

lemma *wf-add*:
fixes *m* :: (′*a*,′*b*) *insert-message*
assumes *wfP* (*depends-on M*)
assumes ⋀*n*. *n* ∈ (*M* ∪ {*m*}) ⟹ *I m* ∉ *deps* (*Insert n*)
assumes *m* ∉ *M*
shows *wfP* (*depends-on* (*M* ∪ {*m*}))
⟨*proof*⟩

lemma *create-insert-p-s-ordered*:
assumes *is-associated-string N s*
assumes *a-conditions* (*insert-messages N*) *a*
assumes *Inr* (*Insert m*) = *create-insert s n σ new-id*
shows *a* (*P m*) < *a* (*S m*)
⟨*proof*⟩

lemma *create-insert-consistent*:
assumes *consistent M*
assumes *is-associated-string N s*
assumes *N* ⊆ *M*
assumes *Inr m* = *create-insert s n σ new-id*
assumes *new-id* ∉ *I* ‘ *insert-messages M*
shows *consistent* (*M* ∪ {*m*})
⟨*proof*⟩

lemma *bind-simp*: (*x* ⤜ (λ*l*. *y l*) = *Inr r*) ⟹ (*y* (*projr x*) = *Inr r*)
⟨*proof*⟩

lemma *create-delete-consistent*:
  assumes *consistent M*
  assumes *is-associated-string N s*
  assumes $N \subseteq M$
  assumes *Inr m* = *create-delete s n*
  shows *consistent* $(M \cup \{m\})$
⟨*proof*⟩

end

5.5 Termination Proof for *integrate-insert*

theory *IntegrateInsertCommute*
  imports *IntegrateAlgorithm Consistency CreateConsistent*
begin

In the following we show that *integrate-insert* terminates. Note that, this does not yet imply that the return value will not be an error state.

lemma *substr-simp* [*simp*]: *substr s l u* = *nths s* $\{k.\ l < Suc\ k \wedge Suc\ k < u\}$
⟨*proof*⟩

declare *substr.simps* [*simp del*]

Instead of simplifying *substr* with its definition we use *substr-simp* as a simplification rule. The right hand side of *substr-simp* is a better representation within proofs. However, we cannot directly define *substr* using the right hand side as it is not constructible term for Isabelle.

lemma *int-ins-loop-term-1*:
  assumes *isOK* (*mapM* (*concurrent w l u*) *t*)
  assumes $x \in$ *set* (*concat* (*projr* (*mapM* (*concurrent w l u*) *t*)))
  shows $x \in$ (*InString* ∘ *I*) ‘ (*set t*)
  ⟨*proof*⟩

lemma *fromSingleton-simp*: (*fromSingleton xs* = *Inr x*) = ([*x*] = *xs*)
  ⟨*proof*⟩

lemma *filt-simp*: ([*b*] = *filter p* [*0..<n*]) =
    $(p\ b \wedge b < n \wedge (\forall y < n.\ p\ y \longrightarrow b = y))$
  ⟨*proof*⟩

lemma *substr-eff*:
  assumes $x \in$ (*InString* ∘ *I*) ‘ *set* (*substr w l u*)
  assumes *isOK* (*idx w x*)
  shows $l <$ (*projr* (*idx w x*)) ∧ (*projr* (*idx w x*)) $< u$
⟨*proof*⟩

lemma *find-zip*:
  assumes *find* (*cond* ∘ *snd*) (*zip* (*p#v*) (*v@*[*s*])) = *Some* (*x*,*y*)
  assumes $v \neq []$

shows
cond y
$x \in set\ v \lor y \in set\ v$
$x = p \lor (x \in set\ v \land \neg(cond\ x))$
$y = s \lor (y \in set\ v)$
⟨*proof*⟩

fun *int-ins-measure′*
where
int-ins-measure′ $(m,w,p,s) = ($
do {
$l \leftarrow idx\ w\ p;$
$u \leftarrow idx\ w\ s;$
$assert\ (l < u);$
$return\ (u - l)$
})

fun *int-ins-measure*
where
int-ins-measure $(m,w,p,s) =$ *case-sum* $(\lambda e.\ 0)$ *id* (*int-ins-measure′* (m,w,p,s))

We show that during the iteration of *integrate-insert*, the arguments are decreasing with respect to *int-ins-measure*. Note, this means that the distance between the W-characters with identifiers p (resp. s) is decreasing.

lemma *int-ins-loop-term*:
assumes $idx\ w\ p = Inr\ l$
assumes $idx\ w\ s = Inr\ u$
assumes *mapM* (*concurrent* $w\ l\ u$) (*substr* $w\ l\ u$) $= Inr\ d$
assumes *concat* $d \neq []$
assumes *find* $((\lambda x.[\![I\ m]\!] < x \lor x = s) \circ snd)$
$(zip\ (p\#concat\ d)\ (concat\ d@[s])) = Some\ r$
shows *int-ins-measure* $(m,\ w,\ r) < u - l$
⟨*proof*⟩

lemma *assert-ok-simp* [*simp*]: $(assert\ p = Inr\ z) = p$ ⟨*proof*⟩

termination *integrate-insert*
⟨*proof*⟩

5.6 Integrate Commutes

locale *integrate-insert-commute* =
fixes M :: (′*a* :: *linorder*, ′*s*) *message set*
fixes a :: ′*a extended* ⇒ ′*a position*
fixes s :: (′*a*, ′*s*) *woot-character list*
assumes *associated-string-assm*: *is-associated-string M s*
assumes *a-conditions-assm*: *a-conditions* (*insert-messages M*) *a*
begin

lemma *dist-ext-ids*: *distinct* (*ext-ids s*)
⟨*proof*⟩

lemma *I-inj-on-S*:
$l < length\ s \wedge u < length\ s \wedge I(s\ !\ l) = I(s\ !\ u) \Longrightarrow l = u$
⟨*proof*⟩

lemma *idx-find*:
assumes $x < length\ (ext\text{-}ids\ s)$
assumes $ext\text{-}ids\ s\ !\ x = i$
shows $idx\ s\ i = Inr\ x$
⟨*proof*⟩

lemma *obtain-idx*:
assumes $x \in set\ (ext\text{-}ids\ s)$
shows $\exists\, i.\ idx\ s\ x = Inr\ i$
⟨*proof*⟩

lemma *sorted-a*:
assumes $idx\ s\ x = Inr\ l$
assumes $idx\ s\ y = Inr\ u$
shows $(l \le u) = (a\ x \le a\ y)$
⟨*proof*⟩

lemma *sorted-a-le*: $idx\ s\ x = Inr\ l \Longrightarrow idx\ s\ y = Inr\ u \Longrightarrow (l < u) = (a\ x < a\ y)$
⟨*proof*⟩

lemma *idx-intro-ext*: $i < length\ (ext\text{-}ids\ s) \Longrightarrow idx\ s\ (ext\text{-}ids\ s\ !\ i) = Inr\ i$
⟨*proof*⟩

lemma *idx-intro*:
assumes $i < length\ s$
shows $idx\ s\ [\![I\ (s\ !\ i)]\!] = Inr\ (Suc\ i)$
⟨*proof*⟩

end

locale *integrate-insert-commute-insert* = *integrate-insert-commute* +
fixes *m*
assumes *consistent-assm*: $consistent\ (M \cup \{Insert\ m\})$
assumes *insert-assm*: $Insert\ m \notin M$
assumes *a-conditions-assm-2*:
$a\text{-}conditions\ (insert\text{-}messages\ (M \cup \{Insert\ m\}))\ a$
begin

definition *invariant* **where**
$invariant\ pm\ sm = (pm \in set\ (ext\text{-}ids\ s) \wedge sm \in set\ (ext\text{-}ids\ s) \wedge$
$subset\ (a\ pm,\ a\ sm)\ (a\ (P\ m),\ a\ (S\ m)) \wedge$
$elem\ (a\ [\![I\ m]\!])\ (a\ pm,\ a\ sm))$

fun *is-concurrent* **where**
  $\textit{is-concurrent pm sm } x = (x \in \textit{set } s \wedge$
   $\textit{subset } (a\ pm,\ a\ sm)\ (a\ (P\ x),\ a\ (S\ x)) \wedge$
   $\textit{elem } (a\ [\![I\ x]\!])\ (a\ pm,\ a\ sm))$

lemma *no-id-collision*: $I\ m \notin I$ ‘ *insert-messages* M
⟨*proof*⟩

lemma *not-deleted*: $\textit{to-woot-char } m = \textit{to-woot-character } M\ m$
⟨*proof*⟩

lemma *invariant-imp-sorted*:
  assumes $\textit{Suc } l < \textit{length } (\textit{ext-ids } s)$
  assumes $a(\textit{ext-ids } s\ !\ l) < a\ [\![I\ m]\!] \wedge a\ [\![I\ m]\!] < a(\textit{ext-ids } s\ !\ (l+1))$
  shows $\textit{sorted-wrt } (<)\ (\textit{map } a\ (\textit{ext-ids } ((\textit{take } l\ s)@\textit{to-woot-char } m\#\textit{drop } l\ s)))$
⟨*proof*⟩

lemma *no-self-dep*: $\neg\ \textit{depends-on } (\textit{insert-messages } M \cup \{m\})\ m\ m$
⟨*proof*⟩

lemma *pred-succ-order*:
  $m' \in (\textit{insert-messages } M \cup \{m\}) \Longrightarrow a(P\ m') < a\ [\![I\ m']\!] \wedge a(S\ m') > a\ [\![I\ m']\!]$
  ⟨*proof*⟩

lemma *find-dep*:
  assumes $\textit{Insert } m' \in (M \cup \{\textit{Insert } m\})$
  assumes $i \in \textit{deps } (\textit{Insert } m')$
  shows $[\![i]\!] \in \textit{set } (\textit{ext-ids } s)$
⟨*proof*⟩

lemma *find-pred*:
  $m' \in (\textit{insert-messages } M \cup \{m\}) \Longrightarrow P\ m' \in \textit{set } (\textit{ext-ids } s)$
  ⟨*proof*⟩

lemma *find-succ*:
  $m' \in (\textit{insert-messages } M \cup \{m\}) \Longrightarrow S\ m' \in \textit{set } (\textit{ext-ids } s)$
  ⟨*proof*⟩

fun *is-certified-associated-string′* **where**
  $\textit{is-certified-associated-string}'\ (\textit{Inr } v) = ($
    $\textit{set } v = \textit{to-woot-character } (M \cup \{\textit{Insert } m\})$ ‘
      $(\textit{insert-messages } (M \cup \{\textit{Insert } m\})) \wedge$
    $\textit{sorted-wrt } (<)\ (\textit{map } a\ (\textit{ext-ids } v))) \mid$
  $\textit{is-certified-associated-string}'\ (\textit{Inl -}) = \textit{False}$

lemma *integrate-insert-final-step*:
  assumes $\textit{invariant pm sm}$
  assumes $\textit{idx } s\ pm = \textit{Inr } l$

assumes *idx s sm* = *Inr* (*Suc l*)
shows *is-certified-associated-string'* (*Inr* (*take l s*@(*to-woot-char m*)#*drop l s*))
⟨*proof*⟩

lemma *concurrent-eff*:
assumes *idx s pm* = *Inr l*
assumes *idx s sm* = *Inr u*
obtains *d* **where** *mapM* (*concurrent s l u*) (*substr s l u*) = *Inr d* ∧
set (*concat d*) = *InString* ' *I* ' {*x*. *is-concurrent pm sm x*}
⟨*proof*⟩

lemma *concurrent-eff-2*:
assumes *invariant pm sm*
assumes *is-concurrent pm sm x*
shows *preserve-order* ⟦*I x*⟧ ⟦*I m*⟧ (*a* ⟦*I x*⟧) (*a* ⟦*I m*⟧)
⟨*proof*⟩

lemma *concurrent-eff-3*:
assumes *idx s pm* = *Inr l*
assumes *idx s sm* = *Inr u*
assumes *Suc l* < *u*
shows {*x*. *is-concurrent pm sm x*} ≠ {}
⟨*proof*⟩

lemma *integrate-insert-result-helper*:
invariant pm sm ⟹ *m'* = *m* ⟹ *s'* = *s* ⟹
is-certified-associated-string' (*integrate-insert m' s' pm sm*)
⟨*proof*⟩

lemma *integrate-insert-result*:
is-certified-associated-string' (*integrate-insert m s* (*P m*) (*S m*))
⟨*proof*⟩
end

lemma *integrate-insert-result*:
assumes *consistent* (*M* ∪ {*Insert m*})
assumes *Insert m* ∉ *M*
assumes *is-associated-string M s*
shows *is-certified-associated-string* (*M* ∪ {*Insert m*}) (*integrate-insert m s* (*P* *m*) (*S m*))
⟨*proof*⟩

locale *integrate-insert-commute-delete* = *integrate-insert-commute* +
fixes *m*
assumes *consistent-assm*: *consistent* (*M* ∪ {*Delete m*})
begin

fun *delete* :: ('*a*, '*s*) *woot-character* ⇒ ('*a*, '*s*) *woot-character*
where *delete* (*InsertMessage p i u* -) = *InsertMessage p i u None*

definition *delete-only-m* :: ('*a*, '*s*) *woot-character* ⇒ ('*a*, '*s*) *woot-character*
  where *delete-only-m x* = (*if DeleteMessage* (*I x*) = *m then delete x else x*)

lemma *set-s*: *set s* = *to-woot-character M* ' *insert-messages M*
  ⟨*proof*⟩

lemma *delete-only-m-effect*:
  delete-only-m (*to-woot-character M x*) = *to-woot-character* (*M* ∪ {*Delete m*}) *x*
  ⟨*proof*⟩

lemma *integrate-delete-result*:
  is-certified-associated-string (*M* ∪ {*Delete m*}) (*integrate-delete m s*)
⟨*proof*⟩
end

lemma *integrate-delete-result*:
  assumes *consistent* (*M* ∪ {*Delete m*})
  assumes *is-associated-string M s*
  shows *is-certified-associated-string* (*M* ∪ {*Delete m*}) (*integrate-delete m s*)
⟨*proof*⟩

fun *is-delete* :: (('*a*, '*s*) *message*) ⇒ *bool*
  where
    is-delete (*Insert m*) = *False* |
    is-delete (*Delete m*) = *True*

proposition *integrate-insert-commute*:
  assumes *consistent* (*M* ∪ {*m*})
  assumes *is-delete m* ∨ *m* ∉ *M*
  assumes *is-associated-string M s*
  shows *is-certified-associated-string* (*M* ∪ {*m*}) (*integrate s m*)
  ⟨*proof*⟩

end

5.7 Strong Convergence

theory *StrongConvergence*
  imports *IntegrateInsertCommute CreateConsistent HOL.Finite-Set DistributedExecution*
begin

lemma (**in** *dist-execution*) *happened-before-same*:
  assumes $i < j$
  assumes $j <$ *length* (*events k*)
  shows (*happened-immediately-before*)$^{++}$ (*k*,*i*) (*k*,*j*)
⟨*proof*⟩

definition *make-set* **where** *make-set* $(k :: nat)\ p = \{x.\ \exists j.\ p\ j\ x \wedge j < k\}$

lemma *make-set-nil* [*simp*]: *make-set 0 p* = {} ⟨*proof*⟩

lemma *make-set-suc* [*simp*]: *make-set* (*Suc k*) *p* = *make-set k p* ∪ {*x. p k x*}
⟨*proof*⟩

lemma (**in** *dist-execution*) *received-messages-eff*:
assumes *is-valid-state-id* (*i*,*j*)
shows *set* (*received-messages* (*i*,*j*)) = *make-set j* (λ*k x.* (∃ *s. event-at* (*i*, *k*) (*Receive s x*)))
⟨*proof*⟩

lemma (**in** *dist-execution*) *finite-valid-event-ids*:
finite {*i. is-valid-event-id i*}
⟨*proof*⟩

lemma (**in** *dist-execution*) *send-insert-id-1*:
state i ⤜ (λ*s. create-insert s n σ i*) = *Inr* (*Insert m*) ⟹ *I m* = *i*
⟨*proof*⟩

lemma (**in** *dist-execution*) *send-insert-id-2*:
state i ⤜ (λ*s. create-delete s n*) = *Inr* (*Insert m*) ⟹ *False*
⟨*proof*⟩

lemma (**in** *dist-execution*) *send-insert-id*:
event-at i (*Send* (*Insert m*)) ⟹ *I m* = *i*
⟨*proof*⟩

lemma (**in** *dist-execution*) *recv-insert-once*:
event-at (*i*,*j*) (*Receive s* (*Insert m*)) ⟹ *event-at* (*i*,*k*) (*Receive t* (*Insert m*)) ⟹ *j* = *k*
⟨*proof*⟩

proposition *integrate-insert-commute′*:
fixes *M m s*
assumes *consistent M*
assumes *is-delete m* ∨ *m* ∉ *T*
assumes *m* ∈ *M*
assumes *T* ⊆ *M*
assumes *deps m* ⊆ *I* ‘ *insert-messages T*
assumes *is-certified-associated-string T s*
shows *is-certified-associated-string* (*T* ∪ {*m*}) (*s* ⤜ (λ*t. integrate t m*))
⟨*proof*⟩

lemma *foldM-rev*: *foldM f s* (*li*@[*ll*]) = *foldM f s li* ⤜ (λ*t. f t ll*)
⟨*proof*⟩

lemma (**in** *dist-execution*) *state-is-associated-string′*:

fixes *i M*
assumes $j \leq$ *length* (*events i*)
assumes *consistent M*
assumes *make-set j* (λk *m*. $\exists$ *s*. *event-at* (*i*,*k*) (*Receive s m*)) $\subseteq$ *M*
shows *is-certified-associated-string* (*make-set j* (λk *m*. $\exists$ *s*. *event-at* (*i*,*k*) (*Receive s m*))) (*state* (*i*,*j*))
⟨*proof*⟩

lemma (**in** *dist-execution*) *sent-before-recv*:
assumes *event-at* (*i*,*k*) (*Receive s m*)
assumes $j <$ *length* (*events i*)
assumes $k < j$
shows *event-at s* (*Send m*) $\wedge$ *happened-immediately-before*$^{++}$ *s* (*i*,*j*)
⟨*proof*⟩

lemma (**in** *dist-execution*) *irrefl-p*: *irreflp* (*happened-immediately-before*$^{++}$)
⟨*proof*⟩

lemma (**in** *dist-execution*) *new-messages-keep-consistency*:
assumes *consistent M*
assumes *event-at i* (*Send m*)
assumes *set* (*received-messages i*) $\subseteq$ *M*
assumes $i \notin I$ ‘ *insert-messages M*
shows *consistent* (*insert m M*)
⟨*proof*⟩

lemma (**in** *dist-execution*) *sent-messages-consistent*:
consistent {*m*. ($\exists$ *i*. *event-at i* (*Send m*))}
⟨*proof*⟩

lemma (**in** *dist-execution*) *received-messages-were-sent*:
assumes *is-valid-state-id* (*i*,*j*)
shows *make-set j* (λk *m*. ($\exists$ *s*. *event-at* (*i*, *k*) (*Receive s m*))) $\subseteq$ {*m*. $\exists$ *i*. *event-at i* (*Send m*)}
⟨*proof*⟩

lemma (**in** *dist-execution*) *state-is-associated-string*:
assumes *is-valid-state-id i*
shows *is-certified-associated-string* (*set* (*received-messages i*)) (*state i*)
⟨*proof*⟩

end

6 Strong Eventual Consistency

theory *SEC*
imports *StrongConvergence*
begin

In the following theorem we establish that all reached states are successful. This implies with the unconditional termination property (Section 5.5) of it that the integration algorithm never fails.

theorem (**in** *dist-execution*) *no-failure*:
  fixes i
  assumes *is-valid-state-id* i
  shows *isOK* (*state* i)
  ⟨*proof*⟩

The following theorem establishes that any pair of peers having received the same set of updates, will be in the same state.

theorem (**in** *dist-execution*) *strong-convergence*:
  assumes *is-valid-state-id* i
  assumes *is-valid-state-id* j
  assumes *set* (*received-messages* i) $=$ *set* (*received-messages* j)
  shows *state* i $=$ *state* j
  ⟨*proof*⟩

As we noted in Section 4.7, we have not assumed eventual delivery, but a corollary of this theorem with the eventual delivery assumption implies eventual consistency. Since finally all peer would have received all messages, i.e., an equal set.

7 Code generation

export-code *integrate create-insert create-delete* **in** *Haskell*
  module-name *WOOT* **file-prefix** *code*

8 Proof Outline

In this section we outline and motivate the approach we took to prove the strong eventual consistency of WOOT.

While introducing operation-based CRDTs Shapiro et al. also establish [24][Theorem 2.2]. If the following two conditions are met:

- Concurrent operations commute, i.e., if a pair of operations m_1, m_2 is concurrent with respect to the order induced by the happened-before relation, and they are both applicable to a state s, then the message m_1 (resp. m_2) is still applicable on the state reached by applying m_2 (resp. m_1) on s and the resulting states are equal.
- Assuming causal delivery, the messages are applicable.

Then the CRDT has strong convergence. The same authors extend the above result in [23, Proposition 2.2] to more general delivery orders $\xrightarrow{d}$

(weaker than the one induced by the happened-before relation), i.e., two messages may be causally dependent but concurrent with respect to $\xrightarrow{d}$. Assuming operations that are concurrent with respect to $\xrightarrow{d}$ commute, and messages are applicable, when the delivery order respects $\xrightarrow{d}$ then again the CRDT has strong convergence.

A key difficulty of the consistency proof of the WOOT framework is that the applicability condition for the WOOT framework has three constraints:

1. Dependencies must be met.
2. Identifiers must be distinct.
3. The order must be consistent, i.e. the predecessor W-character must appear before the successor W-character in the state an insert message is being integrated.

The first constraint is a direct consequence of the semantic causal delivery order. The uniqueness of identifiers can be directly established by analyzing the implementation of the message creation algorithms. Alternatively, Gomes et al. [7] use an axiomatic approach, where they require the underlying network protocol to deliver messages with unique identifiers. They provide a formal framework in Isabelle/HOL that can be used to show consistency of arbitrary CRDTs. Their results could be used to establish constraints 1 and 2.

The last constraint is the most intricate one, and forces us to use a different method to establish the strong eventual consistency. The fact that the order constraint is fulfilled is a consequence of the consistency property. But the current fundamental lemmas require applicability of the operations in the first place to establish consistency, which would result in a circular argument.

Zeller et. al. actually predict the above circumstance in the context of state-based CRDTs [27]:

> In theory it could even be the case that there are two reachable states for which the merge operation does not yield the correct result, but where the two states can never be reached in the same execution.

Because of the above, we treat WOOT as a distributed message passing algorithm and show convergence by establishing a global invariant, which is maintained during the execution of the framework. The invariant captures that the W-characters appear in the same order on all peers. It has strong convergence as a consequence, in the special case, when peers have received the same set of updates. It also implies that the generated messages will be applicable.

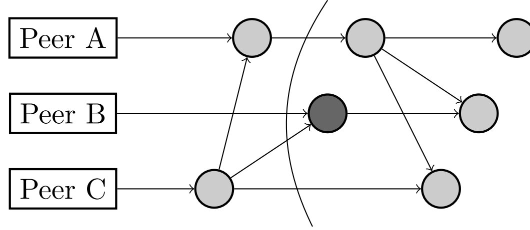

Figure 3: Example state graph, where the consistency is established left of the bend curve.

In Figure 3, we exemplify an induction step in a proof over the execution of the framework. The invariant is established for all states left of the dashed lines, and we show that it remains true if we include the state, drawn in dark gray. Note that induction proceeds in an order consistent with the happened-before relation.

The technique we are using is to define a relation *is-associated-string* from a set of messages to the final state their application leads to. Crucially, that relation can be defined in a message-order independent way. We show that it correctly models the behaviour of Algorithm *integrate* by establishing that applying the integration algorithm to the associated string of a set M leads to the associated string of the set $M \cup \{m\}$ in Proposition *integrate-insert-commute*.

We also show that at most one s fulfills *is-associated-string M s*, which automatically implies commutativity (cf. Lemma *associated-string-unique*).

Note that the domain of the relation *is-associated-string* consists of the sets of messages that we call *consistent*. We show that, in every state of a peer, the set of received messages will be consistent. The main ingredient required for the definition of a consistent set of messages as the relation *is-associated-string* are *sort keys* associated to the W-characters, which we will explain in the following Section.

8.1 Sort Keys

There is an implicit sort key, which is deterministically computable, using the immutable data associated to a W-character and the data of the W-characters it (transitively) depends on.

We show that Algorithm *integrate* effectively maintains the W-characters ordered with respect to that sort key, which is the reason we can construct the mapping *is-associated-string* in a message-order independent way. An alternative viewpoint would be to see Algorithm *integrate-insert* as an optimized version of a more mundane algorithm, that just inserts the W-characters using this implicit sort key.

Since the sort key is deterministically computable using the immutable data associated to a W-character and the data of the W-characters it (transitively) depends on, all peers could perform this computation independently, which leads to the conclusion that the W-characters will be ordered consistently across all peers.

The construction relies on a combinator Ψ that computes the sort key for a W-character, and which requires as input:

- The unique identifier associated to a W-character.
- The sort keys of the predecessor/successor W-characters.

Its values are elements of a totally ordered space.

Note that the predecessor (resp. successor) W-character of a W-character is the W-character that was immediately before (resp. after) it at the time it was inserted. Like its unique identifier, it is immutable data associated with that W-character. Sometimes a W-character is inserted at the beginning (resp. end) of the string. For those W-characters, we use the special smallest (resp. largest) sort keys, denoted by $\vdash$ (resp. $\dashv$) as predecessor (resp. successor). These keys themselves are never associated to a W-character.

We will write $\Psi\ (l,\ u)\ i$ for the value computed by the combinator for a W-character with identifier i, assuming the sort key of its predecessor (resp. successor) is l (resp. u).

For example, the sort key for a W-character with identifier i inserted in an empty string (hence its predecessor is $\vdash$ and its successor is $\dashv$) will be $\Psi\ (\vdash,\ \dashv)\ i$. A W-character inserted between that character and the end of the string, with identifier j, would be assigned the sort key $\Psi\ ([\![\Psi\ (\vdash,\ \dashv)\ i]\!],\ \dashv)\ j$.

The sort key needs to fulfill a couple of properties, to be useful:

There should never be a pair of W-characters with the same sort key. Note, if this happens, even if those W-characters were equal or ordered consistently, we would not be able to insert a new W-character between those W-characters.

Since the W-characters have themselves unique identifiers, a method to insure the above property is to require that Ψ be injective with respect to the identifier of the W-character it computes a sort key for, i.e., $\Psi\ (l,\ u)\ i = \Psi\ (l',\ u')\ i' \Longrightarrow i = i'$.

Another essential property is that the W-characters with predecessor having the sort key l and successor having the sort key u should have a sort key that is between l and u, such that the W-character is inserted between the preceding and succeeding W-character, i.e., $l < \Psi\ (l,u)\ i < u$.

This latter property ensures intention preservation, i.e. the inserted W-character will be placed at the place the user intended.

If we review function *concurrent*, then we see that the algorithm compares W-characters by identifier, in the special case, when the inserted W-character is compared to a W-character whose predecessor and successor are outside of the range it is to be inserted in. A careful investigation, leads to the conclusion that:

If $l \leq l' < \Psi\ (l,u)\ i < u' \leq u$ then $\Psi(l,u)\ i$ can be compared with $\Psi(l',u')\ i'$ by comparing i with i', i.e.:

- $i < i' \Longrightarrow \Psi\ (l,u)\ i < \Psi(l',u')\ i'$

In Section 5.1 we show that a combinator Ψ with the above properties can be constructed (cf. Propositions *psi-narrow psi-mono psi-elem*). Using the sort keys we can define the notion of a consistent set of messages as well as the relation *is-associated-string* in a message-order independent way.

8.2 Induction

We have a couple of criteria that define a consistent set of messages:

- Each insert message in the set has a unique identifier.
- If a message depends on another message identifier, a message with that identifier will be present. Note that for insert messages, these are the predecessor/successor W-characters if present. For delete messages it is the corresponding insert message.
- The dependencies form a well-order, i.e., there is no dependency cycle.
- It is possible to assign sort keys to each insert message, such that the assigned sort key for each insert message is equal to the value returned by the Ψ for it, using the associated sort keys of its predecessor and successors, i.e., $a\ (P\ m) < a\ (S\ m) \wedge a\ [\![I\ m]\!] = [\![\Psi\ (a\ (P\ m),\ a\ (S\ m))\ (I\ m)]\!]$. Note that we also require that sort key of the predecessor is smaller than the sort key of the successor.

The relation *is-associated-string* is then defined by ordering the insert messages according to the assigned sort keys above and marking W-characters, for which there are delete messages as deleted.

The induction proof (Lemma *dist-execution.sent-messages-consistent*) over the states of the framework is straight forward: Using Lemma *top-sort* we find a possible order of the states consistent with the happened before relation. The induction invariant is that the set of generated messages by all peers is consistent (independent of whether they have been received by all peers (yet)). The latter also implies that the subset a peer has received in any of those states is consistent, using the additional fact that each messages

dependencies will be delivered before the message itself (see also Lemma *consistent-subset* and Proposition *integrate-insert-commute'*). For the induction step, we rely on the results from Section 5.4 that any additional created messages will keep the set of messages consistent and that the peers' states will be consistent with the (consistent subset of) messages they received (Lemma *dist-execution.state-is-associated-string'*).

end

9 Example

theory *Example*
 imports *SEC*
begin

In this section we formalize the example from Figure 2 for a possible run of the WOOT framework with three peers, each performing an edit operation. We verify that it fulfills the conditions of the locale *dist-execution* and apply the theorems.

```
datatype example-peers
  = PeerA
  | PeerB
  | PeerC
derive linorder example-peers

fun example-events :: example-peers ⇒ (example-peers, char) event list where
  example-events PeerA = [
      Send (Insert (InsertMessage ⊢ (PeerA, 0) ⊣ CHR ''B'')),
      Receive (PeerA, 0) (Insert (InsertMessage ⊢ (PeerA, 0) ⊣ CHR ''B'')),
      Receive (PeerB, 0) (Insert (InsertMessage ⊢ (PeerB, 0) ⊣ CHR ''A'')),
      Receive (PeerC, 1) (Insert (InsertMessage ⟦(PeerA, 0)⟧ (PeerC, 1) ⊣ CHR
''R''))
  ] |
  example-events PeerB = [
      Send (Insert (InsertMessage ⊢ (PeerB, 0) ⊣ CHR ''A'')),
      Receive (PeerB, 0) (Insert (InsertMessage ⊢ (PeerB, 0) ⊣ CHR ''A'')),
      Receive (PeerA, 0) (Insert (InsertMessage ⊢ (PeerA, 0) ⊣ CHR ''B'')),
      Receive (PeerC, 1) (Insert (InsertMessage ⟦(PeerA, 0)⟧ (PeerC, 1) ⊣ CHR
''R''))
  ] |
  example-events PeerC = [
      Receive (PeerA, 0) (Insert (InsertMessage ⊢ (PeerA, 0) ⊣ CHR ''B'')),
      Send (Insert (InsertMessage ⟦(PeerA, 0)⟧ (PeerC, 1) ⊣ CHR ''R'')),
      Receive (PeerC, 1) (Insert (InsertMessage ⟦(PeerA, 0)⟧ (PeerC, 1) ⊣ CHR
''R'')),
      Receive (PeerB, 0) (Insert (InsertMessage ⊢ (PeerB, 0) ⊣ CHR ''A''))
  ]
```

The function *example-events* returns the sequence of events that each peer

evaluates. We instantiate the preliminary context by showing that the set of peers is finite.

interpretation *example*: *dist-execution-preliminary example-events*
⟨*proof*⟩

To prove that the *happened-before* relation is acyclic, we provide an order on the state that is consistent with it, i.e.:

- The assigned indicies for successive states of the same peer are increasing.
- The assigned index of a state receiving a message is larger than the assigned index of the messages source state.

```
fun witness-acyclic-events :: example-peers event-id ⇒ nat
  where
    witness-acyclic-events (PeerA, 0) = 0 |
    witness-acyclic-events (PeerB, 0) = 1 |
    witness-acyclic-events (PeerA, (Suc 0)) = 2 |
    witness-acyclic-events (PeerB, (Suc 0)) = 3 |
    witness-acyclic-events (PeerC, 0) = 4 |
    witness-acyclic-events (PeerC, (Suc 0)) = 5 |
    witness-acyclic-events (PeerC, (Suc (Suc 0))) = 6 |
    witness-acyclic-events (PeerC, (Suc (Suc (Suc 0)))) = 7 |
    witness-acyclic-events (PeerA, (Suc (Suc 0))) = 8 |
    witness-acyclic-events (PeerA, (Suc (Suc (Suc 0)))) = 9 |
    witness-acyclic-events (PeerB, (Suc (Suc 0))) = 8 |
    witness-acyclic-events (PeerB, (Suc (Suc (Suc 0)))) = 9 |
    witness-acyclic-events (PeerA, (Suc (Suc (Suc (Suc n))))) = undefined |
    witness-acyclic-events (PeerB, (Suc (Suc (Suc (Suc n))))) = undefined |
    witness-acyclic-events (PeerC, (Suc (Suc (Suc (Suc n))))) = undefined
```

To prove that the created messages make sense, we provide the edit operation that results with it. The first function is the inserted letter and the second function is the position the letter was inserted.

```
fun witness-create-letter :: example-peers event-id ⇒ char
  where
    witness-create-letter (PeerA, 0) = CHR ''B'' |
    witness-create-letter (PeerB, 0) = CHR ''A'' |
    witness-create-letter (PeerC, Suc 0) = CHR ''R''
```

```
fun witness-create-position :: example-peers event-id ⇒ nat
  where
    witness-create-position (PeerA, 0) = 0 |
    witness-create-position (PeerB, 0) = 0 |
    witness-create-position (PeerC, Suc 0) = 1
```

To prove that dependencies of a message are received before a message,

we provide the event id as well as the message, when the peer received a messages dependency.

fun *witness-deps-received-at* :: *example-peers event-id* ⇒ *example-peers event-id* ⇒ *nat*
where
witness-deps-received-at (*PeerA*, *Suc* (*Suc* (*Suc 0*))) (*PeerA*, *0*) = *1* |
witness-deps-received-at (*PeerB*, *Suc* (*Suc* (*Suc 0*))) (*PeerA*, *0*) = *2* |
witness-deps-received-at (*PeerC*, *Suc* (*Suc 0*)) (*PeerA*, *0*) = *0*

fun *witness-deps-received-is* :: *example-peers event-id* ⇒ *example-peers event-id* ⇒ (*example-peers event-id*, *char*) *insert-message*
where
witness-deps-received-is (*PeerA*, *Suc* (*Suc* (*Suc 0*))) (*PeerA*, *0*) = (*InsertMessage* ⊢ (*PeerA*, *0*) ⊣ *CHR* ''*B*'') |
witness-deps-received-is (*PeerB*, *Suc* (*Suc* (*Suc 0*))) (*PeerA*, *0*) = (*InsertMessage* ⊢ (*PeerA*, *0*) ⊣ *CHR* ''*B*'') |
witness-deps-received-is (*PeerC*, *Suc* (*Suc 0*)) (*PeerA*, *0*) = (*InsertMessage* ⊢ (*PeerA*, *0*) ⊣ *CHR* ''*B*'')

lemma *well-order-consistent*:
fixes *i j*
assumes *example.happened-immediately-before i j*
shows *witness-acyclic-events i* < *witness-acyclic-events j*
⟨*proof*⟩

Finally we show that the *example-events* meet the assumptions for the distributed execution context.

interpretation *example*: *dist-execution example-events*
⟨*proof*⟩

As expected all peers reach the same final state.

lemma
example.state (*PeerA*, *4*) = *Inr* [
InsertMessage ⊢ (*PeerA*, *0*) ⊣ (*Some CHR* ''*B*''),
InsertMessage ⊢ (*PeerB*, *0*) ⊣ (*Some CHR* ''*A*''),
InsertMessage ⟦(*PeerA*, *0*)⟧ (*PeerC*, *1*) ⊣ (*Some CHR* ''*R*'')]
example.state (*PeerA*, *4*) = *example.state* (*PeerB*, *4*)
example.state (*PeerB*, *4*) = *example.state* (*PeerC*, *4*)
⟨*proof*⟩

We can also derive the equivalence of states using the strong convergence theorem. For example the set of received messages in the third state of peer A and B is equivalent, even though they were not received in the same order:

lemma
example.state (*PeerA*, *3*) = *example.state* (*PeerB*, *3*)
⟨*proof*⟩

Similarly we can conclude that reached states are successful.

lemma
isOK (*example.state* (*PeerC*, *4*))
⟨*proof*⟩

end

References

[1] M. Ahmed-Nacer, C.-L. Ignat, G. Oster, H.-G. Roh, and P. Urso. Evaluating CRDTs for real-time document editing. In *Symposium on Document Engineering (DocEng)*, pages 103–112. ACM, 2011.

[2] L. Briot, P. Urso, and M. Shapiro. High responsiveness for group editing crdts. In *International Conference on Supporting Group Work (GROUP)*, pages 51–60. ACM, 2016.

[3] R. Brown, S. Cribbs, C. Meiklejohn, and S. Elliott. Riak DT map: a composable, convergent replicated dictionary. In *Workshop on Principles and Practice of Eventual Consistency*, page 1. ACM, 2014.

[4] R. Dallaway. WOOT model for Scala and JavaScript via Scala.js. https://github.com/d6y/wootjs, 2016. Accessed: 2017-01-25.

[5] C. A. Ellis and S. J. Gibbs. Concurrency control in groupware systems. In *ACM SIGMOD Record*, volume 18, pages 399–407. ACM, 1989.

[6] V. Emanouilov. Collaborative rich text editor. https://github.com/kroky/woot, 2016. Accessed: 2017-01-25.

[7] V. B. F. Gomes, M. Kleppmann, D. P. Mulligan, and A. R. Beresford. Verifying strong eventual consistency in distributed systems. *Proceedings of the ACM on Programming Languages (PACMPL)*, 1(OOPSLA), 2017.

[8] R. Kaplan. A real time collaboration toy project based on WOOT. https://github.com/ryankaplan/woot-collaborative-editor, 2016. Accessed: 2017-01-25.

[9] G. Klein, T. Nipkow, D. von Oheimb, C. Pusch, and M. Strecker. Java source and bytecode formalizations in isabelle: μjava.

[10] A. D. Kshemkalyani and M. Singhal. *Distributed Computing: Principles, Algorithms, and Systems*. Cambridge University Press, 2011.

[11] S. Kumawat and A. Khunteta. A survey on operational transformation algorithms: Challenges, issues and achievements. *International Journal of Computer Applications*, 3(12):30–38, 2010.

[12] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, July 1978.

[13] M. Letia, N. Preguiça, and M. Shapiro. Consistency without concurrency control in large, dynamic systems. *ACM SIGOPS Operating Systems Review*, 44(2):29–34, 2010.

[14] D. Li and R. Li. An admissibility-based operational transformation framework for collaborative editing systems. *Computer Supported Cooperative Work (CSCW)*, 19(1):1–43, 2010.

[15] B. Nédelec, P. Molli, A. Mostefaoui, and E. Desmontils. LSEQ: an adaptive structure for sequences in distributed collaborative editing. In *Symposium on Document Engineering (DocEng)*, pages 37–46. ACM, 2013.

[16] T. Olson. Real time group editor without operational transformation. https://github.com/TGOlson/woot-haskell, 2016. Accessed: 2017-01-25.

[17] G. Oster, P. Molli, P. Urso, and A. Imine. Tombstone transformation functions for ensuring consistency in collaborative editing systems. In *International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom)*, pages 1–10. IEEE, 2006.

[18] G. Oster, P. Urso, P. Molli, and A. Imine. Real time group editors without operational transformation. Technical Report RR-5580, INRIA, 2005.

[19] G. Oster, P. Urso, P. Molli, and A. Imine. Data consistency for P2P collaborative editing. In *Conference on Computer Supported Cooperative Work (CSCW)*, pages 259–268. ACM, 2006.

[20] N. Preguica, J. M. Marques, M. Shapiro, and M. Letia. A commutative replicated data type for cooperative editing. In *International Conference on Distributed Computing Systems (ICDCS)*, pages 395–403. IEEE, 2009.

[21] M. Raynal. *Distributed Algorithms for Message-Passing Systems*. Springer, 2013.

[22] H.-G. Roh, J.-S. Kim, J. Lee, and S. Maeng. Optimistic operations for replicated abstract data types. Technical report, 2009.

[23] M. Shapiro, N. Preguiça, C. Baquero, and M. Zawirski. A comprehensive study of Convergent and Commutative Replicated Data Types. Research Report RR-7506, Inria – Centre Paris-Rocquencourt ; INRIA, Jan. 2011.

[24] M. Shapiro, N. Preguiça, C. Baquero, and M. Zawirski. Conflict-free replicated data types. In *International Conference on Stabilization, Safety, and Security of Distributed Systems (SSS)*, pages 386–400. Springer-Verlag, 2011.

[25] S. Weiss, P. Urso, and P. Molli. Wooki: a P2P wiki-based collaborative writing tool. In *International Conference on Web Information Systems Engineering*, pages 503–512. Springer, 2007.

[26] S. Weiss, P. Urso, and P. Molli. Logoot: A scalable optimistic replication algorithm for collaborative editing on P2P networks. In *International Conference on Distributed Computing Systems (ICDCS)*, pages 404–412. IEEE, 2009.

[27] P. Zeller, A. Bieniusa, and A. Poetzsch-Heffter. Formal specification and verification of crdts. In E. Ábrahám and C. Palamidessi, editors, *Formal Techniques for Distributed Objects, Components, and Systems - 34th IFIP WG 6.1 International Conference, FORTE 2014, Held as Part of the 9th International Federated Conference on Distributed Computing Techniques, DisCoTec 2014, Berlin, Germany, June 3-5, 2014. Proceedings*, volume 8461 of *Lecture Notes in Computer Science*, pages 33–48. Springer, 2014.