

Decomposing Opacity (Appendix)

Mohsen Lesani Jens Palsberg
Computer Science Department
University of California, Los Angeles
{lesani, palsberg}@ucla.edu

Contents

1	Histories	2
2	Opacity	5
3	Markability	7
4	Marking Theorem	10
5	Synchronization Object Types	23
6	Marking TL2	37
7	Marking DSTM (visible reads)	57
8	Marking NORec	59
9	The Cost of Read Validation	61

1 Histories

Strings. We use $\|s\|$ to denote the size of the string s . If s_1 and s_2 are strings, we write $s_1 \in s_2$ iff s_1 is a subsequence of s_2 . For example, $bd \in abcde$. Let s be an isogram (i.e. contains no repeating occurrence of the alphabet.) For any $s_1, s_2 \in s$, we write $s_1 \triangleleft_s s_2$ iff the last element of s_1 occurs before the first element of s_2 in s . For example, $ab \triangleleft_{abcde} de$. We use $s(i)$ to denote the i^{th} element of s .

Method calls and events. Let O denote the set of objects, n denote the set of method names, $Trans$ denote the set of transactions, V denote the set of values and $Label$ denote the set of labels. We use l, R and W as labels. The set of invocation events is $Inv = \{inv(l \triangleright o.n_T(v)) \mid l \in Label, o \in O, n \in N, T \in Trans, v \in V\}$. The set of response events is $Res = \{ret(l \triangleright v) \mid l \in Label, v \in V \cup \{\mathbb{A}, \mathbb{C}\}\}$. (\mathbb{A} and \mathbb{C} are used later to denote abortion and commitment of transactions.) The set of events is $Ev = Inv \cup Res$. We will use the term completed method call to denote a sequence of an invocation event followed by the matching response event (with the same label). We use $l \triangleright o.n_T(v):v$ to denote the completed method call $inv(l \triangleright o.n_T(v)) \cdot ret(l \triangleright v)$.

Operations on event sequences. Let E and E' be event sequences. We use $E \cdot E'$ to denote the concatenation of E and E' . For a transaction T , we use $E|T$ to denote the subsequence of all events of T in E . For an object o , we use $E|o$ to denote the subsequence of all events of o in E . Let $Sequential$ be the set of sequences of completed method calls possibly followed by an invocation event. A transaction T is sequential in a sequence of events E if $E|T$ is sequential. **Execution history.** An execution history is a sequence of events where each invocation event has a unique label and every transaction is sequential. Let $History$ denote the set of execution histories. We say label l is in X and write $l \in X$ if there is an invocation event with label l in X . Let $Labels(X)$ denote the set of labels in X . Let $Trans(X)$ denote the set of transactions in X . As the labels are unique in a history, the following functions on $Labels(X)$ are defined. The functions $obj_X, name_X, trans_X, arg1_X, arg2_X, retv_X$ map labels to the receiving object, the method name, the transaction identifier, the first and the second argument, and the return value associated with the labels. Similarly, iEv and rEv functions on $Labels(X)$ map labels to the invocation and the response events associated with the labels.

A history X is equivalent to a history X' , $X \equiv X'$, if one is a permutation of the other one that is only the events are reordered but the components of the events (including the argument and return values) are preserved.

Real-time relations. For an execution history X , we define the real-time relations $\prec_X, \preceq_X, \sim_X, \lesssim_X$ on $Labels(X)$ as follows: First, $l_1 \prec_X l_2$ iff $rEv(l_1) \triangleleft_X iEv(l_2)$. $l_1 \preceq_X l_2$ iff $l_1 \prec_X l_2 \vee l_1 = l_2$. Second, $l_1 \sim_X l_2$ iff $l_1 \not\prec_X l_2 \wedge l_2 \not\prec_X l_1$. Third, $l_1 \lesssim_X l_2$ iff $l_1 \prec_X l_2 \vee l_1 \sim_X l_2$.

From the definition of $Sequential$ we have that $X \in Sequential$ iff $\forall l, l' \in X: l \preceq_X l' \vee l' \prec_X l$. For an execution history X , we define the real-time relations \preccurlyeq_X and \preceqslant_X as follows. First, $T \preccurlyeq_X T'$ iff $X|T \triangleleft_X X|T'$. Second, $T \preceqslant_X T'$ iff $T \preccurlyeq_X T' \vee T = T'$.

We now define shared memory and transaction histories.

Transactional Memory. The *transactional memory* is a singleton object mem that encapsulates a set of locations where each location, $i \in I$, $I = \{1, \dots, m\}$ encapsulates a value v . The object mem has five methods $init_t()$, $read_t(i)$, $write_t(i, v)$, $commit_t()$ and $abort_t()$. The parameter t is the invoking transaction identifier. The method call $init_t()$ initializes t and returns ok . The method call $read_t(i)$ returns the value of location i or aborts t and returns \mathbb{A} . The method $write_t(i, v)$ writes v to location i and returns ok or aborts t and returns \mathbb{A} . The method $commit_t()$ tries to commit transaction t . If t is successfully committed, it returns \mathbb{C} ; otherwise, it returns \mathbb{A} . The method $abort_t()$ aborts t and returns \mathbb{A} . The object mem can be implicit, that is $read_t(i)$ abbreviates $mem.read_T(i)$. The values $ok, \mathbb{A}, \mathbb{C}$ are reserved values that are used to denote successful completion of writes and abortion and commitment of transactions respectively.

Transaction History. A *transaction history* H is an execution history such that $H|mem = H_{Init} \cdot H'$ with the following conditions. H_{Init} is the following history that initializes every location to v_0 . $H_{Init} =$

$l_{0i} \triangleright \text{init}_{T_0}() \cdot l_{00} \triangleright \text{write}_{T_0}(1, v_0):\text{ok} \cdot \dots \cdot l_{0m} \triangleright \text{write}_{T_0}(m, v_0):\text{ok} \cdot l_{0c} \triangleright \text{commit}_{T_0}:\mathbb{C}$. For every $T \in H'$, the history $H'|T$ is a prefix of $e.e'$. The event sequence e is the initialization method call $l \triangleright \text{init}_T()$ (for some l), and then a sequence of reads $l \triangleright \text{read}_T(i):v$ and writes $l \triangleright \text{write}_T(i, v)$ (for some l, i , and v). The event sequence e' is one of the following sequences (for some l, i , and v): (1) $\text{inv}(l \triangleright \text{read}_T(i)), \text{ret}(l \triangleright \mathbb{A})$, (2) $\text{inv}(l \triangleright \text{write}_T(i, v)), \text{ret}(l \triangleright \mathbb{A})$, (3) $\text{inv}(l \triangleright \text{commit}_T()), \text{ret}(l \triangleright \mathbb{C})$, (4) $\text{inv}(l \triangleright \text{commit}_T()), \text{ret}(l \triangleright \mathbb{A})$, or (5) $\text{inv}(l \triangleright \text{abort}_T()), \text{ret}(l \triangleright \mathbb{A})$. Let $T\text{History}$ denote the set of transaction histories. Let $\text{Trans}(H)$ denote the set of transactions of H . The projection of H on i , written $H|i$, denotes the subsequence of history H that contains exactly the events on location i . For a TM algorithm specification π , let $\mathbb{H}(\pi)$ denote the set of complete transaction histories that π results.

Now, we present a set of basic lemmas about execution orders.

Lemma 1 (XASym) *For every execution history X and method calls l and l' , if $l \prec_X l'$ then $\neg(l' \prec_X l) \wedge \neg(l' \sim_X l) \wedge \neg(l' = l)$*

Lemma 2 (XTrans) *For every execution history X and method calls l, l' and l'' , if $l \prec_X l'$ and $l' \prec l''$ then $l \prec_X l''$*

Lemma 3 (XXTrans) *For every execution history X and method calls l_1, l_2, l_3 , and l_4 , if $l_1 \prec_X l_2$, $l_2 \lesssim_X l_3$, and $l_3 \prec_X l_4$ then $l_1 \prec_X l_4$*

Lemma 4 (XTotal) *For every execution history X and method calls l and l' , if $l \in X$, and $l' \in X$, then $(l \prec_X l') \vee (l' \prec_X l) \vee (l \sim_X l') \vee (l = l')$*

Lemma 5 (X2X) *For every execution history X and method calls l and l' , if $l \prec_X l'$ then $l \in X$, and $l' \in X$.*

Lemma 6 (XI2X) *For every execution history X and method calls l, l' , and l'' if $l \prec_X l'$ and $\text{inv}(l') \triangleleft_X \text{inv}(l'')$ then $l \prec_X l''$.*

Lemma 7 (RX2X) *For every execution history X and method calls l, l' , and l'' if $\text{ret}(l) \triangleleft_X \text{ret}(l')$ and $l' \prec_X l''$ then $l \prec_X l''$.*

Proof Sketches.

Lemma 1:

We Assume

(1) $l \prec_X l'$

From [1] and definition of \sim_X , we have

(2) $\neg(l' \sim_X l)$

From [1], we have

(3) $\text{rEv}(l) \triangleleft_X \text{iEv}(l')$

As X is a valid history, we have

(4) $\text{iEv}(l) \triangleleft_X \text{rEv}(l)$

(5) $\text{iEv}(l') \triangleleft_X \text{rEv}(l')$

From [4], [3], and [5], we have

(6) $\text{iEv}(l) \triangleleft_X \text{rEv}(l')$

From [6], we have

(7) $\neg(\text{rEv}(l') \triangleleft_X \text{iEv}(l))$

From [7], and definition of \prec_X , we have

$$(9) \neg(l' \prec_X l)$$

From [3] and [7], we have

$$(9) \neg(l' = l)$$

Lemma 2:

Straightforward from the definition of \prec_X .

Lemma 3:

We have

$$(1) l_1 \prec_X l_2$$

$$(2) l_3 \prec_X l_4$$

$$(3) l_2 \sim_X l_3$$

From [1], we have

$$(4) rEv(l_1) \triangleleft_X iEv(l_2)$$

From [2], we have

$$(5) rEv(l_3) \triangleleft_X iEv(l_4)$$

From [3], we have

$$(6) \neg(l_3 \prec_X l_2)$$

From [6], we have

$$(7) \neg(rEv(l_3) \triangleleft_X iEv(l_2))$$

From [7], we have

$$(8) iEv(l_2) \triangleleft_X rEv(l_3)$$

From [4], [8], and [5], we have

$$(9) rEv(l_1) \triangleleft_X iEv(l_4)$$

From [9], we have

$$l_1 \prec_X l_4$$

Lemma 4:

Straightforward from the definition of \prec_X and \sim_X .

Lemma 5:

Straightforward from the definition of \prec_X .

Lemma 6:

Straightforward from the definition of \prec_X and \triangleleft_X .

Lemma 7:

Straightforward from the definition of \prec_X and \triangleleft_X .

2 Opacity

In this section, we present a formal definition of opacity. Opacity of a TM algorithm is defined in two steps. First, it is defined what it means for a transaction history to be opaque which is called final-state-opacity. Then, a TM algorithm is defined to be opaque if every transaction history of every source program running on top of that TM algorithm is final-state-opaque.

FinalStateOpaque is defined in Figure 1. First, we present some preliminary definitions. We use T prefix before some of the terms for transactions to avoid confusion with the terms for concurrent objects. We say that a transaction history is *transaction sequential* if it is a sequence of transactions. A transaction T is *committed* or *aborted* in a transaction history H if there is respectively a commit or abort response event for T in H . A *completed* transaction is either committed or aborted. A *live* transaction is a transaction that is not completed. A transaction history is *complete* if all its transactions are completed. A *pending* transaction has a pending event and a *commit-pending* transaction has a commit pending event. An *extension* of a history is obtained by committing or aborting its commit-pending transactions and aborting the other live transactions. For a TM algorithm specification π , let $\mathbb{H}(\pi)$ denote the set of complete transaction histories that π results.

If H is a transaction history and p is a predicate on transaction identifiers, we define $filter(H, p)$ to be the subsequence of H that contains the events of transactions T for which $p(T)$ is true. The *visible history* for a transaction T in a sequential transaction history S , $Visible(S, T)$, is the sequence of committed transactions before T in S and T itself. The *sequential specification* of a location i , $SeqSpec(i)$, is the set of sequential histories of read and write method calls on location i where every read returns the value given as the argument to the latest preceding write (regardless of transaction identifiers). It is essentially the sequential specification of a register. *Transactional sequential specification* is the set of complete sequential transaction histories S that for every transaction T and location i , $Visible(S, T)|i$ is a member of the sequential specification of i . A transaction history H is *final-state-opaque* if there is an equivalent sequential transaction history S for an extension of H such that S is real-time-preserving and a member of transactional sequential specification. The sequential history S is called the justifying history. In other words, every correct concurrent execution is indistinguishable from a correct sequential execution.

$$\begin{aligned}
TReads(H) &= \{R \mid R \in H \wedge obj_H(R) = mem \wedge name_H(R) = read \wedge retv_H(R) \neq \mathbb{A}\} \\
TWrites(H) &= \{W \mid W \in H \wedge obj_H(W) = mem \wedge name_H(W) = write \wedge retv_H(W) \neq \mathbb{A}\} \\
Commits(H) &= \{C \mid C \in H \wedge obj_H(C) = mem \wedge name_H(C) = commit\} \\
Trans(H) &= \{T \mid \exists l \in H : trans_H(l) = T\} \\
TSequential &= \{S \in THistory \mid \preceq_S \text{ is a total order of } Trans(S)\} \\
Committed(H) &= \{T \mid \exists l \in Commits(H) \wedge retv_H(l) = \mathbb{C}\} \\
Aborted(H) &= \{T \mid \exists l \in H : obj_H(l) = mem \wedge trans_H(l) = T \wedge retv_H(l) = \mathbb{A}\} \\
Completed(H) &= Committed(H) \cup Aborted(H) \\
Live(H) &= Trans(H) \setminus Completed(H) \\
TComplete &= \{H \in THistory \mid \forall T \in Trans(H) : T \in Completed(H)\} \\
CommitPending(H) &= \{T \in Live(H) \mid \exists l \in H : obj_H(l) = mem \wedge trans_H(l) = T \wedge name_H(l) = commit\} \\
TExtension(H) &= \{H' \in THistory \mid \exists H'' : H' = H \cdot H'' \\
&\quad Trans(H'') \subseteq Trans(H) \wedge \forall T : \|H''\|T\| \leq 1 \wedge \\
&\quad Live(H) \setminus CommitPending(H) \subseteq Aborted(H') \wedge \\
&\quad CommitPending(H) \subseteq Completed(H')\} \\
Visible(S, T) &= filter(S, \lambda T'. (T' = T) \vee ((T' \prec_S T) \wedge T' \in Committed(S))) \\
NoWriteBetween_S(W, R) &= \forall W' \in TWrites(S) : W' \preceq_S W \vee R \prec_S W' \\
SeqSpec(i) &= \{S \in Sequential \mid \forall R \in TReads(S) : \exists W \in TWrites(S) : \\
&\quad W \prec_S R \wedge NoWriteBetween_S(W, R) \wedge \\
&\quad retv_S(R) = arg2_S(W)\} \\
TSeqSpec &= \{S \in TSequential \cap TComplete \mid \forall T \in S : \forall i \in I : \\
&\quad (Visible(S, T) \mid i) \in SeqSpec(i)\} \\
FinalStateOpaque &= \{H \in THistory \mid \exists H' \in TExtension(H) : \exists S \in TSequential : \\
&\quad H' \equiv S \wedge \preceq_{H'} \subseteq \preceq_S \wedge S \in TSeqSpec\}
\end{aligned}$$

Figure 1: *FinalStateOpaque*

3 Markability

In this section, we define markability for general histories.

First, we present some preliminary definitions in Figure 2. (We use T prefix before some of the terms for transactions to avoid confusion with similar terms that used for concurrent objects.) A transaction T is *committed* or *aborted* in a transaction history H if there is respectively a commit or abort response event for T in H . A *completed* transaction is either committed or aborted. A *live* transaction is a transaction that is not completed. A *pending* transaction has a pending event and a *commit-pending* transaction has a commit pending event. An *extension* of a history is obtained by committing or aborting its commit-pending transactions and aborting the other live transactions.

A *local* read is a read that is preceded by a write by the same transaction to the same location. Intuitively, a local read should read a value that is previously written by the same transaction and hence the name. A *global read* is a read that is not local. A *local write* is a write that precedes a write by the same transaction to the same location. A local write is overwritten by the same transaction and hence the name. A *global write* is a write that is not local. The *writers of i* are the committed transactions that write to location i .

Markability is defined in Figure 3. A *marking* \sqsubseteq of a transaction history is the union of the following relations on the set of transactions and the global reads.

- The *effect order*: The set of transactions is totally ordered by \sqsubseteq . In other words, \sqsubseteq is total, antisymmetric and transitive on the set of transactions.
- The *access orders*: For each global read R from a location i , R and every writer of i are ordered by \sqsubseteq . In other words, \sqsubseteq totally orders every global read R from a location i with respect to writers of i and is antisymmetric.

The *write-observation* property is comprised of the two properties: *local write-observation* and *global write-observation*. Local write-observation requires that every local read R from a location i returns the value written by the last write before it in the same transaction to i . Global write-observation requires that the value that every global read R from a location i returns is the value written by the global write of the last pre-accessor transaction to i . We remind that pre-accessors of R are the writers of i that are ordered before R in the access order and the last pre-accessor is the one that is greatest in the effect order.

The *Read-preservation* property requires that for every global read R from location i by transaction T , there is no writer transaction T' of i such that T' is marked between T and R (i.e. T' accesses i after R and takes effect before T), or similarly, T' is marked between R and T (i.e. T' takes effect after T and accesses i before R).

The *real-time-preservation* property requires that if T is before T' in the real-time order, then T takes effect before T' as well.

A transaction history is *final-state-markable* if and only if there exists a marking for an extension of it that is write-observant, read-preserving, and real-time-preserving.

$$\begin{aligned}
Committed(H) &= \{T \mid \exists l \in H: obj_H(l) = mem \wedge trans_H(l) = T \wedge retv_H(l) = \mathbb{C}\} \\
Aborted(H) &= \{T \mid \exists l \in H: obj_H(l) = mem \wedge trans_H(l) = T \wedge retv_H(l) = \mathbb{A}\} \\
Completed(H) &= Committed(H) \cup Aborted(H) \\
Live(H) &= Trans(H) \setminus Completed(H) \\
CommitPending(H) &= \{T \in Live(H) \mid \exists l \in H: obj_H(l) = mem \wedge trans_H(l) = T \wedge \\
&\quad name_H(l) = commit\} \\
TExtension(H) &= \{H' \in THistory \mid \exists H'': H' = H \cdot H'' \\
&\quad Trans(H'') \subseteq Trans(H) \wedge \forall T: \|H''|T\| \leq 1 \wedge \\
&\quad Live(H) \setminus CommitPending(H) \subseteq Aborted(H') \wedge \\
&\quad CommitPending(H) \subseteq Completed(H')\} \\
TReads(H) &= \{R \mid R \in H \wedge obj_H(R) = mem \wedge name_H(R) = read \wedge retv_H(R) \neq \mathbb{A}\} \\
TWrites(H) &= \{W \mid W \in H \wedge obj_H(W) = mem \wedge name_H(W) = write \wedge retv_H(W) \neq \mathbb{A}\} \\
LocalTReads(H) &= \{R \mid R \in TReads(H) \wedge \exists W \in TWrites(H): \\
&\quad trans_H(R) = trans_H(W) \wedge arg1_H(R) = arg1_H(W) \wedge W \prec_H R\} \\
GlobalTReads(H) &= TReads(H) \setminus LocalTReads(H) \\
LocalTWrites(H) &= \{W \mid W \in TWrites(H) \wedge \exists W' \in TWrites(H): \\
&\quad trans_H(W) = trans_H(W') \wedge arg1_H(W) = arg1_H(W') \wedge W \prec_H W'\} \\
GlobalTWrites(H) &= TWrites(H) \setminus LocalTWrites(H) \\
Writers_H(i) &= \{T \in Trans(H) \mid \exists l \in TWrites(H): arg1_H(l) = i \wedge \\
&\quad trans_H(l) = T \wedge T \in Committed(H)\}
\end{aligned}$$

Figure 2: Basic Definitions

$$\begin{aligned}
\text{Marking}(H) &= \{ \sqsubseteq \mid \\
&\quad \forall T1, T2, T3 \in \text{Trans}(H): \\
&\quad\quad (T1 \sqsubseteq T2 \vee T2 \sqsubseteq T1) \wedge \\
&\quad\quad (T1 \sqsubseteq T2 \wedge T2 \sqsubseteq T1) \Rightarrow (T1 = T2) \wedge \\
&\quad\quad (T1 \sqsubseteq T2) \wedge (T2 \sqsubseteq T3) \Rightarrow (T1 \sqsubseteq T3) \wedge \\
&\quad \forall R, T: \text{Let } i = \text{arg1}_H(R): (R \in \text{GlobalTRead}(H) \wedge T \in \text{Writers}_H(i)) \Rightarrow \\
&\quad\quad (R \sqsubseteq T \vee T \sqsubseteq R) \wedge \\
&\quad\quad (R \sqsubseteq T \Rightarrow \neg T \sqsubseteq R) \wedge (T \sqsubseteq R \Rightarrow \neg R \sqsubseteq T) \} \\
\text{NoWriteBetween}_H(W, R) &\Leftrightarrow \\
&\quad \forall W' \in \text{TWrites}(H): W' \preceq_H W \vee R \prec_H W' \\
\text{LocalWriteObs}(H) &\Leftrightarrow \\
&\quad \forall R \in \text{LocalTReads}(H): \text{Let } T = \text{trans}_H(R), i = \text{arg1}_H(R), H' = H|T|i: \\
&\quad\quad \exists W \in \text{TWrites}(H'): W \prec_{H'} R \wedge \text{NoWriteBetween}_{H'}(W, R) \wedge \text{retv}_{H'}(R) = \text{arg2}_{H'}(W) \\
\text{NoWriterBetween}_{H,i}(x, \sqsubseteq, x') &\Leftrightarrow \\
&\quad \forall T \in \text{Writers}_H(i): T \sqsubseteq x \vee x' \sqsubseteq T \\
\text{LastPreAccessor}_{H,\sqsubseteq}(T', R) &\Leftrightarrow \text{Let } i = \text{arg1}_H(R), T = \text{trans}_H(R): \\
&\quad T' \in \text{Writers}_H(i) \wedge T' \neq T \wedge T' \sqsubset R \wedge \text{NoWriterBetween}_{H,i}(T', \sqsubseteq, R) \\
\text{GlobalWriteObs}(H, \sqsubseteq) &\Leftrightarrow \\
&\quad \forall R \in \text{GlobalTReads}(H): \exists W \in \text{GlobalTWrites}(H): \text{Let } T' = \text{trans}_H(W): \\
&\quad\quad \text{LastPreAccessor}_{H,\sqsubseteq}(T', R) \wedge \text{arg1}_H(R) = \text{arg1}_H(W) \wedge \text{retv}_H(R) = \text{arg2}_H(W) \\
\text{WriteObs}(H, \sqsubseteq) &\Leftrightarrow \\
&\quad \text{LocalWriteObs}(H) \wedge \text{GlobalWriteObs}(H, \sqsubseteq) \\
\text{ReadPres}(H, \sqsubseteq) &\Leftrightarrow \\
&\quad \forall R \in \text{GlobalTReads}(H): \text{Let } i = \text{arg1}_H(R), T = \text{trans}_H(R): \\
&\quad\quad \text{NoWriterBetween}_{H,i}(R, \sqsubseteq, T) \wedge \text{NoWriterBetween}_{H,i}(T, \sqsubseteq, R) \\
\text{RealTimePres}(H, \sqsubseteq) &\Leftrightarrow \\
&\quad \preceq_H \subseteq \sqsubseteq \\
\text{FinalStateMarkable} &= \{ H \in \text{THistory} \mid \exists H' \in \text{TExtension}(H): \exists \sqsubseteq \in \text{Marking}(H'): \\
&\quad\quad \text{ReadPres}(H', \sqsubseteq) \wedge \text{WriteObs}(H', \sqsubseteq) \wedge \text{RealTimePres}(H', \sqsubseteq) \}
\end{aligned}$$

Figure 3: *FinalStateMarkable*

4 Marking Theorem

In this section, we prove the marking theorem.

For the sake of brevity, we use the shorthand notation

$$\exists l = o.n_T(v_1):v_2 \in X$$

for

$$\exists l \in X: \text{obj}_X(l) = o \wedge \text{name}_X(l) = n \wedge \text{trans}_X(l) = T \wedge \text{arg1}_X(l) = v_1 \wedge \text{retv}_X(l) = v_2$$

and similarly for universal quantification.

We also use W, R to denote labels.

Lemma 8 *For all $S \in T\text{Sequential}$, $T \in S$, $S' = \text{Visible}(S, T)$, and $T', T'' \in S'$, we have $T' \preceq_{S'} T'' \iff T' \preceq_S T''$.*

Proof.

$$\begin{aligned} & T' \preceq_{S'} T'' \\ \iff & S'|T' \triangleleft_{S'} S'|T'' \vee T' = T'' \\ \iff & S|T' \triangleleft_{S'} S|T'' \vee T' = T'' \\ \iff & S|T' \triangleleft_S S|T'' \vee T' = T'' \\ \iff & T' \preceq_S T'' \end{aligned}$$

In these four steps we apply:

- 1) the definition of $\preceq_{S'}$,
- 2) that the definition of $\text{Visible}(S, T)$ implies both $S'|T' = S|T'$ and $S'|T'' = S|T''$,
- 3) $S' \in S$, and
- 4) the definition of \preceq_S . □

Lemma 9 For all $S \in TSequential$, $T \in S$, $i \in I$, $v, v' \in V$, $R = read_T(i):v \in GlobalReads(S)$, $S' = Visible(S, T)$, $T' \in S'$, and $W' = write_{T'}(i, v') \in GlobalWrites(S)$, we have

$$NoWriteBetween_{(S'|i)}(W', R) \iff NoWriterBetween_{S,i}(T', \preceq_S, T)$$

Proof.

$$\begin{aligned}
& NoWriteBetween_{(S'|i)}(W', R) \\
\iff & \forall W'' \in Writes(S'|i): W'' \preceq_{(S'|i)} W' \vee R \preceq_{(S'|i)} W'' \\
\iff & \forall T'' \in S'|i: \forall i' \in I: \forall v'' \in V: \forall W'' = write_{T''}(i', v'') \in S'|i: W'' \preceq_{(S'|i)} W' \vee R \preceq_{(S'|i)} W'' \\
\iff & \forall T'' \in S'|i: \forall v'' \in V: \forall W'' = write_{T''}(i, v'') \in S'|i: W'' \preceq_{(S'|i)} W' \vee R \preceq_{(S'|i)} W'' \\
\iff & \forall T'' \in S': \forall v'' \in V: \forall W'' = write_{T''}(i, v'') \in S': W'' \preceq_{S'} W' \vee R \preceq_{S'} W'' \\
\iff & \forall T'' \in S': \forall v'' \in V: \forall W'' = write_{T''}(i, v'') \in S': T'' \preceq_{S'} T' \vee T \preceq_{S'} T'' \\
\iff & \forall T'' \in S': \forall v'' \in V: \forall W'' = write_{T''}(i, v'') \in S': T'' \preceq_S T' \vee T \preceq_S T'' \\
\iff & \forall T'' \in S': \forall v'' \in V: \forall W'' = write_{T''}(i, v'') \in S': T'' \preceq_S T \Rightarrow T'' \preceq_S T' \\
\iff & \forall T'' \in S: \forall v'' \in V: \forall W'' = write_{T''}(i, v'') \in S: \\
& \quad [[(T'' = T) \vee (T'' \preceq_S T \wedge T'' \in Committed(S))] \wedge [T'' \preceq_S T]] \Rightarrow T'' \preceq_S T' \\
\iff & \forall T'' \in S: \forall v'' \in V: \forall W'' = write_{T''}(i, v'') \in S: \\
& \quad (T'' \in Committed(S) \wedge T'' \preceq_S T) \Rightarrow T'' \preceq_S T' \\
\iff & \forall T'' \in Writers_{SS}(i): T'' \preceq_S T \Rightarrow T'' \preceq_S T' \\
\iff & \forall T'' \in Writers_{SS}(i): T'' \preceq_S T' \vee T \preceq_S T'' \\
\iff & NoWriterBetween_{S,i}(T', \preceq_S, T)
\end{aligned}$$

In these twelve steps, we apply:

- 1) the definition of *NoWriteBetween*,
- 2) the definition of *Writes*,
- 3) the definition of projection $S'|i$,
- 4) R, W' and W'' access location i ,
- 5) $S' \in TSequential$ and $R \in GlobalReads(S')$ and $W' \in GlobalWrites(S')$ (that are concluded from $S \in TSequential$, $R \in GlobalReads(S)$, $W' \in GlobalWrites(S)$ and $S' = Visible(S, T)$.),
- 6) Lemma 8,
- 7) Boolean logic and that \preceq_S is total,
- 8) the definition of *Visible*,
- 9) logical simplification,
- 10) the definition of *Writers*,
- 11) Boolean logic and that \preceq_S is total, and
- 12) the definition of *NoWriterBetween*. □

Lemma 10 $TSequential \subset Sequential$

Proof. Straightforward from definitions of $TSequential$, $THistory$ and $Sequential$. \square

Lemma 11 $\forall i \in I: \forall v, v' \in V: \forall T, T' \in Trans: \text{if } R = read_T(i):v, W = write_{T'}(i,v), W' = write_T(i,v'), S \in TSequential, W \prec_S R, NoWriteBetween_S(W, R) \text{ and } W' \prec_S R, \text{ then } T = T'.$

Proof. Suppose (1) $S \in TSequential$, (2) $W \prec_S R$, (3) $NoWriteBetween_S(W, R)$ and (4) $W' \prec_S R$. From [1] and Lemma 10, we have (5) $S \in Sequential$. From [4] and [5], we have (6) $\neg(R \prec_S W')$. From [3] we have (7) $W' \preceq_S W \vee R \prec_S W'$. From [6] and [7], we have (8) $W' \preceq_S W$. From [2] and [8], we have (9) $W' \preceq_S W \preceq_S R$. From [9], [1], and that W' and R are by T and W is by T' , we have $T = T'$. \square

Lemma 12 *Suppose $S \in TSequential$. We have:*

$$\begin{aligned}
& \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in LocalReads(S): \\
& \quad \exists T' \in Visible(S, T): \exists W = write_{T'}(i, v) \in Visible(S, T): \\
& \quad \quad W \prec_{(Visible(S, T) \mid i)} R \wedge NoWriteBetween_{(Visible(S, T) \mid i)}(W, R) \\
\iff & S \in LocalTSeqSpec
\end{aligned}$$

Proof. Suppose $S \in TSequential$. Thus, from Lemma 10, we have $S \in Sequential$. Let $S' = Visible(S, T)$. From $S \in TSequential$ and Lemma 8, we have $S' \in TSequential$. Thus, from Lemma 10, we have $S' \in Sequential$. From the definition of *Visible*, we have $S'|T = S|T$.

$$\begin{aligned}
& \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in LocalReads(S): \\
& \quad \exists T' \in S': \exists W = write_{T'}(i, v) \in S': \\
& \quad \quad W \prec_{(S' \mid i)} R \wedge NoWriteBetween_{(S' \mid i)}(W, R) \\
\iff & \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in LocalReads(S): \\
& \quad \exists v' \in V: \exists W' = write_T(i, v') \in S: W' \prec_S R \wedge \\
& \quad \exists T' \in S': \exists W = write_{T'}(i, v) \in S': \\
& \quad \quad W \prec_{(S' \mid i)} R \wedge NoWriteBetween_{(S' \mid i)}(W, R) \\
\iff & \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in LocalReads(S): \\
& \quad \exists v' \in V: \exists W' = write_T(i, v') \in S': W' \prec_S R \wedge \\
& \quad \exists T' \in S': \exists W = write_{T'}(i, v) \in S': \\
& \quad \quad W \prec_{(S' \mid i)} R \wedge NoWriteBetween_{(S' \mid i)}(W, R) \\
\iff & \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in LocalReads(S): \\
& \quad \exists v' \in V: \exists W' = write_T(i, v') \in S': W' \prec_{S'} R \wedge \\
& \quad \exists T' \in S': \exists W = write_{T'}(i, v) \in S': \\
& \quad \quad W \prec_{(S' \mid i)} R \wedge NoWriteBetween_{(S' \mid i)}(W, R) \\
\iff & \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in LocalReads(S): \\
& \quad \exists v' \in V: \exists W' = write_T(i, v') \in S': W' \prec_{(S' \mid i)} R \wedge \\
& \quad \exists W = write_T(i, v) \in S': \\
& \quad \quad W \prec_{(S' \mid i)} R \wedge NoWriteBetween_{(S' \mid i)}(W, R) \\
\iff & \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in LocalReads(S): \\
& \quad \exists W = write_T(i, v) \in S': \\
& \quad \quad W \prec_{(S' \mid i)} R \wedge NoWriteBetween_{(S' \mid i)}(W, R) \\
\iff & \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in LocalReads(S): \\
& \quad \exists W = write_T(i, v) \in S: \\
& \quad \quad W \prec_{(S' \mid i)} R \wedge NoWriteBetween_{(S' \mid i)}(W, R)
\end{aligned}$$

$$\begin{aligned}
&\iff \forall T \in S: \forall i \in I: \forall v \in V: \forall R = \text{read}_T(i): v \in \text{LocalReads}(S): \\
&\quad \exists W = \text{write}_T(i, v) \in S: \\
&\quad \quad W \prec_{S'} R \wedge \text{NoWriteBetween}_{(S' | i)}(W, R) \\
&\iff \forall T \in S: \forall i \in I: \forall v \in V: \forall R = \text{read}_T(i): v \in \text{LocalReads}(S): \\
&\quad \exists W = \text{write}_T(i, v) \in S: \\
&\quad \quad W \prec_S R \wedge \text{NoWriteBetween}_{(S' | i)}(W, R) \\
&\iff \forall T \in S: \forall i \in I: \forall v \in V: \forall R = \text{read}_T(i): v \in \text{LocalReads}(S): \\
&\quad \exists W = \text{write}_T(i, v) \in S: \\
&\quad \quad W \prec_S R \wedge \forall W' \in \text{Writes}(S' | i): W' \preceq_{(S' | i)} W \vee R \prec_{(S' | i)} W' \\
&\iff \forall T \in S: \forall i \in I: \forall v \in V: \forall R = \text{read}_T(i): v \in \text{LocalReads}(S): \\
&\quad \exists W = \text{write}_T(i, v) \in S: \\
&\quad \quad W \prec_S R \wedge \neg \exists W' \in \text{Writes}(S' | i): \neg(W' \preceq_{(S' | i)} W) \wedge \neg(R \prec_{(S' | i)} W') \\
&\iff \forall T \in S: \forall i \in I: \forall v \in V: \forall R = \text{read}_T(i): v \in \text{LocalReads}(S): \\
&\quad \exists W = \text{write}_T(i, v) \in S: \\
&\quad \quad W \prec_S R \wedge \neg \exists W' \in \text{Writes}(S' | i): W \prec_{(S' | i)} W' \prec_{(S' | i)} R \\
&\iff \forall T \in S: \forall i \in I: \forall v \in V: \forall R = \text{read}_T(i): v \in \text{LocalReads}(S): \\
&\quad \exists W = \text{write}_T(i, v) \in S: \\
&\quad \quad W \prec_S R \wedge \neg \exists v' \in V: \exists W' = \text{write}_T(i, v'): W \prec_{(S' | i)} W' \prec_{(S' | i)} R \\
&\iff \forall T \in S: \forall i \in I: \forall v \in V: \forall R = \text{read}_T(i): v \in \text{LocalReads}(S): \\
&\quad \exists W = \text{write}_T(i, v) \in S: \\
&\quad \quad W \prec_S R \wedge \neg \exists v' \in V: \exists W' = \text{write}_T(i, v'): W \prec_{(S | i)} W' \prec_{(S | i)} R \\
&\iff \forall T \in S: \forall i \in I: \forall v \in V: \forall R = \text{read}_T(i): v \in \text{LocalReads}(S): \\
&\quad \exists W = \text{write}_T(i, v) \in S: \\
&\quad \quad W \prec_S R \wedge \neg \exists W' \in \text{Writes}(S | i): W \prec_{(S | i)} W' \prec_{(S | i)} R \\
&\iff \forall T \in S: \forall i \in I: \forall v \in V: \forall R = \text{read}_T(i): v \in \text{LocalReads}(S): \\
&\quad \exists W = \text{write}_T(i, v) \in S: \\
&\quad \quad W \prec_S R \wedge \forall W' \in \text{Writes}(S | i): \neg(W \prec_{(S | i)} W') \vee \neg(W' \prec_{(S | i)} R) \\
&\iff \forall T \in S: \forall i \in I: \forall v \in V: \forall R = \text{read}_T(i): v \in \text{LocalReads}(S): \\
&\quad \exists W = \text{write}_T(i, v) \in S: W \prec_S R \wedge \\
&\quad \quad \forall W' \in \text{Writes}(S | i): W' \preceq_{(S | i)} W \vee R \prec_{(S | i)} W' \\
&\iff \forall T \in S: \forall i \in I: \forall v \in V: \forall R = \text{read}_T(i): v \in \text{LocalReads}(S): \\
&\quad \exists W = \text{write}_T(i, v) \in S | T | i: W \prec_{S | T | i} R \wedge \\
&\quad \quad \forall W' \in \text{Writes}(S | T | i): W' \preceq_{(S | T | i)} W \vee R \prec_{(S | T | i)} W' \\
&\iff \forall T \in S: \forall i \in I: \forall v \in V: \forall R = \text{read}_T(i): v \in \text{LocalReads}(S): \\
&\quad \exists W = \text{write}_T(i, v) \in S | T | i: \\
&\quad \quad W \prec_{S | T | i} R \wedge \text{NoWriteBetween}_{(S | T | i)}(W, R) \\
&\iff S \in \text{LocalTSeqSpec}
\end{aligned}$$

In these twenty steps, we apply: 1) the definition of *LocalReads*,

- 2) the definition of *Visible*,
- 3) $S'|T = S|T$ and that both W' and R are by T ,
- 4) that both W' and R are on i ,
- 5) Lemma 11,
- 6) duplicate conjunction,
- 7) the definition of *Visible*,
- 8) that both R and W are on i ,
- 9) $S'|T = S|T$ and that both R and W are by T ,
- 10) the definition of *NoWriteBetween*,
- 11) first-order logic,
- 12) $(S' | i) \in \textit{Sequential}$,
- 13) from $(S' | i) \in \textit{TSequential}$, R and W are by transaction T and W' is between them, we have W' is by T ,
- 14) $S'|T = S|T$,
- 15) from $(S | i) \in \textit{TSequential}$, R and W are by transaction T and W' is between them, we have W' is by T .
- 16) first-order logic,
- 17) $(S | i) \in \textit{Sequential}$,
- 18) $(S | i) \in \textit{Sequential}$, $\textit{trans}_H(R) = \textit{trans}_H(W) = T$ and $\textit{arg1}_H(R) = \textit{arg1}_H(W) = i$,
- 19) the definition of *NoWriteBetween*,
- 20) the definition of *LocalTSeqSpec*.

□

Lemma 13 *Suppose $S \in TSequential \cap TComplete$. We have:*

$$\begin{aligned}
& S \in TSeqSpec \\
\iff & S \in LocalTSeqSpec \wedge \\
& \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in GlobalReads(S): \\
& \exists T' \in Committed(S): \exists W = write_{T'}(i, v) \in GlobalWrites(S): \\
& (T' \ll_S T) \wedge NoWriterBetween_{S,i}(T', \underline{\ll}_S T)
\end{aligned}$$

Proof. Suppose $S \in TSequential \cap TComplete$. From $S \in TSequential$ and Lemma 8, we have $Visible(S, T) \in TSequential$.

$$\begin{aligned}
& S \in TSeqSpec \\
\iff & \forall T \in S: \forall i \in I: (Visible(S, T) \mid i) \in SeqSpec(i) \\
\iff & \forall T \in S: \forall i \in I: \\
& \quad \forall T'' \in (Visible(S, T) \mid i): \forall v \in V: \forall R = read_{T''}(i):v \in (Visible(S, T) \mid i): \\
& \quad \exists T' \in (Visible(S, T) \mid i): \exists W = write_{T'}(i, v) \in (Visible(S, T) \mid i): \\
& \quad W \prec_{(Visible(S, T) \mid i)} R \wedge NoWriteBetween_{(Visible(S, T) \mid i)}(W, R) \\
\iff & \forall T \in S: \forall i \in I: \\
& \quad \forall T'' \in Visible(S, T): \forall v \in V: \forall R = read_{T''}(i):v \in Visible(S, T): \\
& \quad \exists T' \in Visible(S, T): \exists W = write_{T'}(i, v) \in Visible(S, T): \\
& \quad W \prec_{(Visible(S, T) \mid i)} R \wedge NoWriteBetween_{(Visible(S, T) \mid i)}(W, R) \\
\iff & \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in S: \\
& \quad \exists T' \in Visible(S, T): \exists W = write_{T'}(i, v) \in Visible(S, T): \\
& \quad W \prec_{(Visible(S, T) \mid i)} R \wedge NoWriteBetween_{(Visible(S, T) \mid i)}(W, R) \\
\iff & \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in LocalReads(S): \\
& \quad \exists T' \in Visible(S, T): \exists W = write_{T'}(i, v) \in Visible(S, T): \\
& \quad W \prec_{(Visible(S, T) \mid i)} R \wedge NoWriteBetween_{(Visible(S, T) \mid i)}(W, R) \\
& \wedge \\
& \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in GlobalReads(S): \\
& \quad \exists T' \in Visible(S, T): \exists W = write_{T'}(i, v) \in Visible(S, T): \\
& \quad W \prec_{(Visible(S, T) \mid i)} R \wedge NoWriteBetween_{(Visible(S, T) \mid i)}(W, R) \\
\iff & S \in LocalTSeqSpec \wedge \\
& \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in GlobalReads(S): \\
& \quad \exists T' \in Visible(S, T): \exists W = write_{T'}(i, v) \in Visible(S, T): \\
& \quad W \prec_{(Visible(S, T) \mid i)} R \wedge NoWriteBetween_{(Visible(S, T) \mid i)}(W, R)
\end{aligned}$$

$$\begin{aligned}
&\iff S \in LocalTSeqSpec \wedge \\
&\quad \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in GlobalReads(S): \\
&\quad \quad \exists T' \in Visible(S,T): \exists W = write_{T'}(i,v) \in Visible(S,T): \\
&\quad \quad \quad W \prec_{Visible(S,T)} R \wedge NoWriteBetween_{(Visible(S,T) \mid i)}(W,R) \\
&\iff S \in LocalTSeqSpec \wedge \\
&\quad \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in GlobalReads(S): \\
&\quad \quad \exists T' \in Visible(S,T): \exists W = write_{T'}(i,v) \in Visible(S,T): \\
&\quad \quad \quad T' \prec_{Visible(S,T)} T \wedge NoWriteBetween_{(Visible(S,T) \mid i)}(W,R) \\
&\iff S \in LocalTSeqSpec \wedge \\
&\quad \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in GlobalReads(S): \\
&\quad \quad \exists T' \in Visible(S,T): \exists W = write_{T'}(i,v) \in Visible(S,T): \\
&\quad \quad \quad T' \prec_S T \wedge NoWriteBetween_{(Visible(S,T) \mid i)}(W,R) \\
&\iff S \in LocalTSeqSpec \wedge \\
&\quad \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in GlobalReads(S): \\
&\quad \quad \exists T' \in Visible(S,T): \exists W = write_{T'}(i,v) \in GlobalWrites(S): \\
&\quad \quad \quad T' \prec_S T \wedge NoWriteBetween_{(Visible(S,T) \mid i)}(W,R) \\
&\iff S \in LocalTSeqSpec \wedge \\
&\quad \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in GlobalReads(S): \\
&\quad \quad \exists T' \in Visible(S,T): \exists W = write_{T'}(i,v) \in GlobalWrites(S): \\
&\quad \quad \quad T' \prec_S T \wedge NoWriterBetween_{S,i}(T', \prec_S T) \\
&\iff S \in LocalTSeqSpec \wedge \\
&\quad \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in GlobalReads(S): \\
&\quad \quad \exists T' \in Visible(S,T): \exists W = write_{T'}(i,v) \in GlobalWrites(S): \\
&\quad \quad \quad (T' \prec_S T) \wedge T' \in Committed(S) \wedge NoWriterBetween_{S,i}(T', \preceq_S T) \\
&\iff S \in LocalTSeqSpec \wedge \\
&\quad \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in GlobalReads(S): \\
&\quad \quad \exists T' \in Committed(S): \exists W = write_{T'}(i,v) \in GlobalWrites(S): \\
&\quad \quad \quad (T' \prec_S T) \wedge NoWriterBetween_{S,i}(T', \preceq_S T)
\end{aligned}$$

In these thirteen steps, we apply:

- 1) the definition of $TSeqSpec$ and $S \in TSequential \cap TComplete$,
- 2) the definition of $SeqSpec(i)$,
- 3) R and W access location i ,
- 4) that we can choose $T'' = T$,
- 5) $Reads(S) = LocalReads(S) \cup GlobalReads(S)$,
- 6) Lemma 12,
- 7) that R and W are both on location i
- 8) that R and W are by transactions T and T' respectively, $Visible(S,T) \in TSequential$, and $R \in GlobalReads(Visible(S,T))$ (because $R \in GlobalReads(R)$ and $Visible(S,T) \mid T = S \mid T$),
- 9) Lemma 8,
- 10) $T' \prec_S T$ and $NoWriteBetween_{(Visible(S,T) \mid i)}(W,R)$,
- 11) Lemma 9,

- 12) $T' \in \text{Visible}(S, T)$ and $(T' \prec_S T)$, and
- 13) the definition of $\text{Visible}(S, T)$.

□

Lemma 14 (Invariance) *If $H \equiv H'$, then $\text{Marking}(H) = \text{Marking}(H')$ and $\text{ReadPres}(H) = \text{ReadPres}(H')$ and $\text{WriteObs}(H) = \text{WriteObs}(H')$.*

Proof. Immediate from the definitions of Marking , ReadPres , and WriteObs . □

Lemma 15 $\forall H \in \text{THistory}: \forall \sqsubseteq \in \text{Marking}(H): \exists S \in \text{TSequential}: H \equiv S \wedge \preceq_H \subseteq \preceq_S \wedge \preceq_S \subseteq \sqsubseteq$.

Proof. Let $H \in \text{THistory}$ and let $\sqsubseteq \in \text{Marking}(H)$. We have that \sqsubseteq is a total order of Trans so we can choose a permutation π on $1..n$ such that $\forall i, j \in 1..n: (i < j) \Leftrightarrow (T_{\pi(i)} \sqsubseteq T_{\pi(j)})$. Define: $S = H|T_{\pi(1)}, \dots, H|T_{\pi(n)}$. It is straightforward to prove that $S \in \text{TSequential} \wedge H \equiv S \wedge \preceq_H \subseteq \preceq_S \wedge \preceq_S \subseteq \sqsubseteq$. □

Lemma 16 *Suppose $\sqsubseteq \in \text{Marking}(H) \wedge p_2 \notin \text{Writers}_H(i)$. If $\text{NoWriterBetween}_{H,i}(T_1, \sqsubseteq, p_2)$ and $\text{NoWriterBetween}_{H,i}(p_2, \sqsubseteq, T_3)$, then $\text{NoWriterBetween}_{H,i}(T_1, \sqsubseteq, T_3)$.*

Proof.

$$\begin{aligned}
& \text{NoWriterBetween}_{H,i}(T_1, \sqsubseteq, p_2) \wedge \text{NoWriterBetween}_{H,i}(p_2, \sqsubseteq, T_3) \\
\iff & \forall T \in \text{Writers}_H(i): (T \sqsubseteq T_1 \vee p_2 \sqsubseteq T) \wedge (T \sqsubseteq p_2 \vee T_3 \sqsubseteq T) \\
\iff & \forall T \in \text{Writers}_H(i): (T \sqsubseteq T_1 \wedge (T \sqsubseteq p_2 \vee T_3 \sqsubseteq T)) \vee \\
& \qquad (p_2 \sqsubseteq T \wedge T \sqsubseteq p_2) \vee (p_2 \sqsubseteq T \wedge T_3 \sqsubseteq T) \\
\implies & \forall T \in \text{Writers}_H(i): (T \sqsubseteq T_1) \vee (T_3 \sqsubseteq T) \\
\iff & \text{NoWriterBetween}_{H,i}(T_1, \sqsubseteq, T_3)
\end{aligned}$$

The first step uses the definition of NoWriterBetween . The second step uses \wedge distribution over \vee . The third step simplifies the first disjunct using conjunction elimination, eliminates the second disjunct using $p_2 \notin \text{Writers}_H(i)$ and simplifies the third disjunct using conjunction elimination. The fourth step uses the definition of NoWriterBetween . □

Lemma 17 *Suppose $S \in TSequential \cap TComplete$. We have:*

$$S \in TSeqSpec \iff S \in FinalStateMarkable \quad (4.1)$$

Proof. Let $S \in TSequential \cap TComplete$. From Lemma 13, the definition of *FinalStateMarkable*, and $S \in TComplete$, we have that we must prove:

$$\begin{aligned} & S \in LocalTSeqSpec \wedge \\ & \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in GlobalReads(S): \\ & \exists T' \in Committed(S): \exists W = write_{T'}(i, v) \in GlobalWrites(S): \\ & \quad (T' \ll_S T) \wedge NoWriterBetween_{S,i}(T', \ll_S, T) \\ \iff & \exists \sqsubseteq \in Marking(S): \ll_S \subseteq \sqsubseteq \wedge \sqsubseteq \in ReadPres(S) \wedge \sqsubseteq \in WriteObs(S) \end{aligned}$$

From the definition of *WriteObs* and *LastPreAccessor* we have that:

$$\begin{aligned} & \sqsubseteq \in WriteObs(S) \\ \iff & S \in LocalTSeqSpec \wedge \\ & \forall T \in Trans: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in GlobalReads(S): \\ & \exists T' \in Trans: \exists W = write_{T'}(i, v) \in GlobalWrites(S): \\ & \quad T' \in Writers_S(i) \wedge T' \neq T \wedge T' \sqsubset R \wedge NoWriterBetween_{S,i}(T', \sqsubseteq, R) \\ \iff & S \in LocalTSeqSpec \wedge \\ & \forall T \in Trans: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in GlobalReads(S): \\ & \exists T' \in Trans: \exists W = write_{T'}(i, v) \in GlobalWrites(S): \\ & \quad T' \in Committed(S) \wedge T' \neq T \wedge T' \sqsubset R \wedge NoWriterBetween_{S,i}(T', \sqsubseteq, R) \end{aligned}$$

We are now ready to prove the two directions of the equivalence.

\Rightarrow :

Assume that

$$\begin{aligned} & S \in LocalTSeqSpec \wedge \\ & \forall T \in S: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in GlobalReads(S): \\ & \exists T' \in Committed(S): \exists W = write_{T'}(i, v) \in GlobalWrites(S): \\ & \quad (T' \ll_S T) \wedge NoWriterBetween_{S,i}(T', \ll_S, T) \end{aligned}$$

Define:

$$\begin{aligned} p_1 \sqsubset p_2 & \iff (p_1 \ll_S p_2) \vee \\ & \quad (trans_S(p_1) \ll_S p_2) \vee \\ & \quad (p_1 \ll_S trans_S(p_2)) \\ p_1 \sqsubseteq p_2 & \iff p_1 \sqsubset \vee p_2 p_1 = p_2 \end{aligned}$$

We show that

$$\begin{aligned} & \sqsubseteq \in Marking(S) \wedge \\ & \ll_S \subseteq \sqsubseteq \wedge \sqsubseteq \in ReadPres(S) \wedge \\ & S \in LocalTSeqSpec \wedge \\ & \forall T \in Trans: \forall i \in I: \forall v \in V: \forall R = read_T(i):v \in GlobalReads(S): \\ & \exists T' \in Trans: \exists W = write_{T'}(i, v) \in GlobalWrites(S): \\ & \quad T' \in Committed(S) \wedge T' \neq T \wedge T' \sqsubset R \wedge NoWriterBetween_{S,i}(T', \sqsubseteq, R) \end{aligned}$$

It is straightforward to prove $\sqsubseteq \in \text{Marking}(S)$ and $\preceq_S \subseteq \sqsubseteq$, $\sqsubseteq \in \text{ReadPres}(S)$. Additionally, the first conjunct of $\text{WriteObs}(S)$ (that is, $S \in \text{LocalTSeqSpec}$) is immediate from the assumption. So, we still need to prove the second conjunct of $\text{WriteObs}(S)$.

Let $T \in \text{Trans}$, $i \in I$, $v \in V$, $R = \text{read}_T(i):v \in \text{GlobalReads}(S)$. From the assumption (the left-hand side), we have that we can find (1) $T' \in \text{Committed}(S)$ and (2) $W = \text{write}_{T'}(i,v) \in \text{GlobalWrites}(S)$ such that (3) $(T' \prec_S T)$ and (4) $\text{NoWriterBetween}_{S,i}(T', \preceq_S, T)$. Let us now prove each conjunct of $T' \neq T \wedge T' \sqsubseteq R \wedge \text{NoWriterBetween}_{S,i}(T', \sqsubseteq, R)$ in turn.

From [3] and that \preceq_S is a total order of $\text{Trans}(S)$, we have (5) $T' \neq T$. From [3] and the definition of \sqsubseteq , we have $T' \sqsubseteq R$. From [4] and $\preceq_S \subseteq \sqsubseteq$, we have (6) $\text{NoWriterBetween}_{S,i}(T', \sqsubseteq, T)$. From $T \preceq_S T$ and the definition of \sqsubseteq , we have (7) $R \sqsubseteq T$. From [6], [7] and the definition of \sqsubseteq and transitivity of \preceq_S , we have $\text{NoWriterBetween}_{S,i}(T', \sqsubseteq, R)$.

\Leftarrow :

Assume the right-hand side and choose $\sqsubseteq \in \text{Marking}(S)$ such that:

$$\begin{aligned} & \preceq_S \subseteq \sqsubseteq \quad \wedge \quad \sqsubseteq \in \text{ReadPres}(S) \quad \wedge \\ & S \in \text{TLocalSeqSpec} \quad \wedge \\ & \forall T \in \text{Trans}: \forall i \in I: \forall v \in V: \forall R = \text{read}_T(i):v \in \text{GlobalReads}(S): \\ & \exists T' \in \text{Committed}(S): \exists W = \text{write}_{T'}(i,v) \in \text{GlobalWrites}(S): \\ & \quad T' \neq T \quad \wedge \quad T' \sqsubseteq R \quad \wedge \quad \text{NoWriterBetween}_{S,i}(T', \sqsubseteq, R) \end{aligned}$$

We show that

$$\begin{aligned} & S \in \text{LocalTSeqSpec} \quad \wedge \\ & \forall T \in S: \forall i \in I: \forall v \in V: \forall R = \text{read}_T(i):v \in \text{GlobalReads}(S): \\ & \exists T' \in \text{Committed}(S): \exists W = \text{write}_{T'}(i,v) \in \text{GlobalWrites}(S): \\ & \quad (T' \prec_S T) \quad \wedge \quad \text{NoWriterBetween}_{S,i}(T', \preceq_S, T) \end{aligned}$$

The first conjunct (of the left-hand side), $S \in \text{LocalTSeqSpec}$, is immediate from the assumption. From the assumption we have (1) $\preceq_S \subseteq \sqsubseteq$, (2) $\sqsubseteq \in \text{ReadPres}(S)$. Let $T \in \text{Trans}$, $i \in I$, $v \in V$, $R = \text{read}_T(i):v \in \text{GlobalReads}(S)$. From the above property of \sqsubseteq , we have that we can find (3) $T' \in \text{Committed}(S)$ and (4) $W = \text{write}_{T'}(i,v) \in \text{GlobalWrites}(S)$ such that (5) $T' \neq T$ and (6) $T' \sqsubseteq R$ and (7) $\text{NoWriterBetween}_{S,i}(T', \sqsubseteq, R)$. From [1], that \sqsubseteq is a total order on $\text{Trans}(S)$ ($\sqsubseteq \in \text{Marking}(S)$), and that \preceq_S is a total order on $\text{Trans}(S)$ ($S \in \text{TSequential}$), we have (8) $\forall T, T' \in \text{Trans}: T' \sqsubseteq T \Rightarrow T' \preceq_S T$.

First we prove $T' \prec_S T$. From [2], we have (9) $\text{NoWriterBetween}_{S,i}(T, \sqsubseteq, R)$. From [3] and [4], we have (10) $T' \in \text{Writers}_S(i)$. From [9] and [10], we have (11) $T' \sqsubseteq T \vee R \sqsubseteq T'$. From [6], $T' \neq R$ and \sqsubseteq is a total order on $\{R\} \cup \text{Writers}_S(i)$ ($\sqsubseteq \in \text{Marking}(S)$), we have (12) $R \not\sqsubseteq T'$. From [11] and [12], we have (13) $T' \sqsubseteq T$. From [8] and [13], we have (14) $T' \preceq_S T$. From [14] and [5], we have $T' \prec_S T$.

Second, we prove $\text{NoWriterBetween}_{S,i}(T', \preceq_S, T)$. From [2], we have (15) $\text{NoWriterBetween}_{S,i}(R, \sqsubseteq, T)$. From $R \notin \text{Writers}_S(i)$, [7], [15], and Lemma 16, we have (16) $\text{NoWriterBetween}_{S,i}(T', \sqsubseteq, T)$. From [16] and [8] we have $\text{NoWriterBetween}_{S,i}(T', \preceq_S, T)$. \square

Theorem 18 (Marking) $FinalStateOpaque = FinalStateMarkable$.

Proof.

$$\begin{aligned}
& FinalStateOpaque \\
= & \{H \in THistory \mid \exists H' \in TExtension(H): \exists S \in TSequential: \\
& \quad H' \equiv S \wedge \preceq_{H'} \subseteq \preceq_S \wedge S \in TSeqSpec\} \\
= & \{H \in THistory \mid \exists H' \in TExtension(H): \exists S \in TSequential: \\
& \quad H' \equiv S \wedge \preceq_{H'} \subseteq \preceq_S \wedge S \in FinalStateMarkable\} \\
= & \{H \in THistory \mid \exists H' \in TExtension(H): \exists S \in TSequential: H' \equiv S \wedge \preceq_{H'} \subseteq \preceq_S \wedge \\
& \quad \exists \sqsubseteq \in Marking(S): \preceq_S \subseteq \sqsubseteq \wedge \sqsubseteq \in ReadPres(S) \cap WriteObs(S)\} \\
= & \{H \in THistory \mid \exists H' \in TExtension(H): \exists S \in TSequential: H' \equiv S \wedge \preceq_{H'} \subseteq \preceq_S \wedge \\
& \quad \exists \sqsubseteq \in Marking(H'): \preceq_S \subseteq \sqsubseteq \wedge \sqsubseteq \in ReadPres(H') \cap WriteObs(H')\} \\
= & \{H \in THistory \mid \exists H' \in TExtension(H): \exists \sqsubseteq \in Marking(H'): \\
& \quad \sqsubseteq \in ReadPres(H') \cap WriteObs(H') \wedge \\
& \quad \exists S \in TSequential: H' \equiv S \wedge \preceq_{H'} \subseteq \preceq_S \wedge \preceq_S \subseteq \sqsubseteq \} \\
= & \{H \in THistory \mid \exists H' \in TExtension(H): \exists \sqsubseteq \in Marking(H'): \\
& \quad \preceq_{H'} \subseteq \sqsubseteq \wedge \sqsubseteq \in ReadPres(H') \cap WriteObs(H') \wedge \\
& \quad \exists S \in TSequential: H' \equiv S \wedge \preceq_{H'} \subseteq \preceq_S \wedge \preceq_S \subseteq \sqsubseteq \} \\
= & \{H \in THistory \mid \exists H' \in TExtension(H): \exists \sqsubseteq \in Marking(H'): \\
& \quad \preceq_{H'} \subseteq \sqsubseteq \wedge \sqsubseteq \in ReadPres(H') \cap WriteObs(H')\} \\
= & Markable
\end{aligned}$$

In these eight steps we apply:

- 1) the definition of $FinalStateOpaque$,
- 2) Lemma 17 and $S \in TComplete$ (because $H' \in TExtension(H)$ and $H' \equiv S$),
- 3) the definition of $FinalStateMarkable$ and $S \in TComplete$,
- 4) Lemma 14,
- 5) logical rearrangement,
- 6) transitivity of \subseteq ,
- 7) Lemma 15, and
- 8) the definition of $FinalStateMarkable$. □

5 Synchronization Object Types

In this subsection, we first define the semantics of basic and linearizable objects. Then, we define the interface and the sequential specifications of the following abstract object types: register, lock, try-lock, counter, set and map. For each abstract object type, we define concrete synchronization object types. We define the following synchronization object types: basic register, atomic register, atomic cas register, lock, try-lock, strong counter, basic set and basic map. For each synchronization object type, we present lemmas that characterize the properties of its execution histories. Please see the end of this section for notes on the proof of the lemmas that we present in this subsection.¹

Basic and Linearizable Object Types

The abstract type of each object o specifies the sequential specification of o , denoted by $SeqSpec(o)$, that is the prefix-closed set of correct sequential histories of o . In the following subsections, we will consider several synchronization object types and define their sequential specifications.

We consider two concurrent types: basic and linearizable. Linearizable objects comply with their sequential specification in every concurrent execution. Basic objects, on the other hand, comply with their sequential specification if they are accessed sequentially.

Definition 1 (Basic Object Semantics) *Every sequential execution on a basic object is an execution in its sequential specification. The semantics of a basic object o , $\mathbb{H}_B(o)$, is a set of histories that is constrained as follows:*

$$\mathbb{H}_B(o) \cap Sequential \subseteq SeqSpec(o) \quad (5.1)$$

Definition 2 (Linearizable Object Semantics) *An execution history X is linearizable for an object o iff there is an indistinguishable sequential history L that is in the sequential specification of o and is real-time-preserving. L is a linearization and \prec_L is a linearization order of X . The semantics of a linearizable object o , $\mathbb{H}_L(o)$, is defined as the following set of execution and linearization pairs.*

$$\mathbb{H}_L(o) = \{(X, L) \mid X \equiv L \wedge L \in SeqSpec(o) \wedge \prec_X \subseteq \prec_L\} \quad (5.2)$$

Lemma 19 (X2L) *For every linearization L of an execution history X on object o and method calls l and l' , if $l \prec_X l'$ then $l \prec_L l'$.*

Lemma 20 (LASym) *For every linearization L of an execution history X on object o and method calls l and l' , if $l \prec_L l'$ then $\neg(l' \prec_L l) \wedge \neg(l = l')$*

Lemma 21 (LTrans) *For every linearization L of an execution history X on object o and method calls l , l' , and l'' , if $l \prec_L l'$ and $l' \prec_L l''$ then $l \prec_L l''$.*

Lemma 22 (LTotal) *For every linearization L of an execution history X on object o and method calls l and l' , if $l \in X$ and $l' \in X$ then $(l \prec_L l') \vee (l' \prec_L l) \vee (l = l')$*

Lemma 23 (L2X) *For every linearization L of an execution history X on object o and method calls l and l' , if $(l \prec_L l')$ then $l \in X$, $l' \in X$, and l and l' are both on o .*

Lemma 24 (XLTrans) *For every linearization L of an execution history X on object o and method calls l_1 , l_2 , l_3 , and l_4 , if $l_1 \prec_X l_2$, $l_2 \prec_L l_3$, $l_3 \prec_X l_4$, then $l_1 \prec_X l_4$*

¹ In this subsection, we use \forall and \exists as a notational convenience. $\forall l: p$ can be rewritten as $\bigwedge_{(l \in Labels(X))} p(X)$ and $\exists l: p$ can be rewritten as $\bigvee_{(l \in Labels(X))} p(X)$.

Register

Register. A register reg is an object that encapsulates a value and supports *read* and *write* methods. The method call $reg.read()$ returns the current encapsulated value of reg . The method call $reg.write(v)$ overwrites the encapsulated value of reg with v .

Definition 3 *The sequential specification of register reg is the set of sequential histories of read and write method calls on reg where every read returns the argument of the latest preceding write (regardless of thread identifiers). (Note that it is assumed that a write method call initializes the register before other methods are invoked.) The sequential specification of a register r , $SeqSpec(r)$, is defined as follows:*

$$isXRead_{X,r}(l_R) = l_R \in X \wedge obj_X(l_R) = r \wedge name_X(l_R) = read \quad (5.3)$$

$$isXWrite_{X,r}(l_W) = l_W \in X \wedge obj_X(l_W) = r \wedge name_X(l_W) = write \quad (5.4)$$

$$NoWriteBetween_{X,r}(l_W, l_R) = \forall l'_W: isXWrite_{X,r}(l'_W) \Rightarrow (l'_W \preceq_X l_W \vee l_R \prec_X l'_W) \quad (5.5)$$

$$isXWriter_{X,r}(l_W, l_R) = isXWrite_{X,r}(l_W) \wedge \quad (5.6)$$

$$l_W \prec_X l_R \wedge$$

$$NoWriteBetween_{X,r}(l_W, l_R)$$

$$Legal(r) = \{S \mid \forall l_R: isXRead_{S,r}(l_R) \Rightarrow \quad (5.7)$$

$$\exists l_W: isXWriter_{S,r}(l_W, l_R) \wedge$$

$$retv_S(l_R) = arg1_S(l_W)\}$$

$$SeqSpec(r) = \{S \mid S|r = S \wedge S \in Sequential \cap Legal(r)\} \quad (5.8)$$

Basic Register. A basic register is a basic instance of the register type.

Let *BasicRegister* denote the type of basic registers.

Lemma 25 *In every sequential execution on a basic register, every read reads the value that the latest preceding write writes. Formally,*

$$\forall reg \in BasicRegister: \forall X \in \mathbb{H}_B(reg): X \in Sequential \Rightarrow \quad (5.9)$$

$$\forall l_R: isXRead_{X,reg}(l_R) \Rightarrow$$

$$\exists l_W: isXWriter_{X,reg}(l_W, l_R) \wedge$$

$$retv_X(l_R) = arg1_X(l_W)$$

Two concurrent read method calls on a register do not conflict. Thus, basic registers can maintain consistency even when the execution involves concurrent read method calls. Let us define

$$isXRaceFree_{X,r}(l) = \forall l_w: isXWrite_{X,r}(l_w) \Rightarrow l_w \preceq_X l \vee l \prec_X l_w \quad (5.10)$$

$$isXSequentiallyWritten_r(X) = \forall l \in X: isXWrite_{X,r}(l) \Rightarrow isXRaceFree_{X,r}(l) \quad (5.11)$$

A method call is race-free if and only if there is no write method call that executes concurrent to it. An execution is sequentially-written if and only if every pair of write method calls on it are ordered in the execution order or in other words, every write method call on it is race-free.

Definition 4 (Basic Register Semantics) *An execution history on a basic register is in the semantics of the basic register if and only if it is not sequentially-written or it is sequentially-written and every race-free*

read reads the value that the latest preceding write writes. The semantics of a basic register r , $\mathbb{H}_B(r)$, is defined as follows.

$$\begin{aligned} \mathbb{H}_B(r) = \{X \mid X|_o = X \wedge & \tag{5.12} \\ & isXSequentiallyWritten_r(X) \Rightarrow \\ & \forall l_r: isXRead_{X,r}(l_r) \wedge isXRaceFree_{X,r}(l_r) \Rightarrow \\ & \exists l_w: isXWriter_{X,r}(l_w, l_r) \wedge \\ & \quad retv_X(l_r) = arg1_X(l_w) \} \end{aligned}$$

Note that if an execution is not sequentially-written, reads may return arbitrary values. Similarly, racy reads may return arbitrary values.

Note that this definition satisfies the constraint of Definition 1.

Lemma 26 (BReg) *In every sequentially-written execution on a basic register, every race-free read reads the value that the latest preceding write writes. Formally,*

$$\begin{aligned} \forall reg \in BasicRegister: \forall X \in \mathbb{H}_B(reg): isXSequentiallyWritten_r(X) \Rightarrow & \tag{5.13} \\ \forall l_R: isXRead_{X,reg}(l_R) \wedge isXRaceFree_{X,r}(l_R) \Rightarrow & \\ \exists l_W: isXWriter_{X,reg}(l_W, l_R) \wedge & \\ \quad retv_X(l_R) = arg1_X(l_W) & \end{aligned}$$

Atomic Register. An atomic register is a linearizable instance of the register type.

Let *AtomicRegister* denote the type of atomic registers.

Let us define

$$LNoWriteBetween_{X,L,r}(l_W, l_R) = \forall l'_W: isXWrite_{X,r}(l'_W) \Rightarrow (l'_W \preceq_L l_W \vee l_R \prec_L l'_W) \tag{5.14}$$

$$\begin{aligned} isLWriter_{X,L,r}(l_W, l_R) = isXWrite_{X,r}(l_W) \wedge & \tag{5.15} \\ l_W \prec_L l_R \wedge & \\ LNoWriteBetween_{X,L,r}(l_W, l_R) & \end{aligned}$$

Lemma 27 (AReg) *In every execution on an atomic register, every read reads the value written by the last write linearized before it. Formally,*

$$\begin{aligned} \forall r \in AtomicRegister: \forall (X, L) \in \mathbb{H}_L(r): & \tag{5.16} \\ \forall l_R: isXRead_{X,r}(l_R) \Rightarrow & \\ \exists l_W: isLWriter_{X,L,r}(l_W, l_R) \wedge & \\ \quad retv_X(l_R) = arg1_X(l_W) & \end{aligned}$$

CAS (Compare-And-Swap) Register

A CAS register is an object that encapsulates a value and supports the *cas* method in addition to *read* and *write* methods. The method call $r.cas(v_1, v_2)$ updates the value of the register to v_2 and returns *true* if the current value of the register is v_1 . It returns *false* otherwise.

A *successful write* is either a *write* method call or a successful *cas* method call. The *written value* of a successful write is its first argument, if it is a *write* method call or is its second argument, if it is a *cas* method call.

Definition 5 *The sequential specification of cas register reg is the set of sequential histories of read, write and cas method calls on reg with the following two conditions. Every read returns the written value of the latest preceding successful write (regardless of thread identifiers). (Note that it is assumed that a write method call initializes the register before other methods are invoked.) Every cas with the first argument v_1 returns true if the written value of the latest preceding successful write is v_1 and returns false otherwise.*

Atomic CAS Register. An atomic CAS register is a linearizable instance of CAS register type.

Let *AtomicCASRegister* denote the type of Atomic CAS registers.

Let us define

$$isXCAS_{X,r}(l_W) = l_W \in X \wedge obj_X(l_W) = r \wedge name_X(l_W) = cas \quad (5.17)$$

$$isXCWrite_{X,r}(l_W) = isXWrite(l_W) \vee (isXCAS(l_W) \wedge retv_X(l_W) = true) \quad (5.18)$$

$$writtenValue_X(l_W) = \begin{cases} arg1_X(l_W) & \text{if } name_X(l_W) = write \\ arg2_X(l_W) & \text{if } name_X(l_W) = cas \end{cases} \quad (5.19)$$

$$LNoWriteBetween_{X,L,r}(l_W, l_R) = \forall l'_W: isXCWrite_{X,r}(l'_W) \Rightarrow (l'_W \preceq_L l_W \vee l_R \prec_L l'_W) \quad (5.20)$$

$$isLCWriter_{X,L,r}(l_W, l_R) = isXCWrite_{X,r}(l_W) \wedge \\ l_W \prec_L l_R \wedge \\ LNoWriteBetween_{X,L,r}(l_W, l_R) \quad (5.21)$$

Lemma 28 (CASRegRead) *In every execution on an atomic cas register, every read returns the value the last successful write linearized before it writes. Formally,*

$$\forall r \in AtomicCASRegister: \forall (X, L) \in \mathbb{H}_L(r): \quad (5.22) \\ \forall l_R: isXRead_{X,r}(l_R) \Rightarrow \\ \exists l_W: isLCWriter_{X,L,r}(l_W, l_R) \wedge \\ retv_X(l_R) = arg1_X(l_W)$$

Lemma 29 (CASRegCAS) *In every execution on an atomic cas register, every cas returns true if its first argument is equal to the argument of the last successful write linearized before it and returns false otherwise. Formally,*

$$\forall reg \in AtomicCASRegister: \forall (X, Reg) \in \mathbb{H}_L(reg): \quad (5.23) \\ \forall l_C, l_W: \\ isXCAS_{X,reg}(l_C) \wedge \\ isLCWriter_{X,Reg,reg}(l_W, l_R) \\ \Rightarrow \\ (writtenValue_X(l_W) = arg1_X(l_C) \Rightarrow retv_X(l_C) = true) \wedge \\ (\neg(writtenValue_X(l_W) = arg1_X(l_C)) \Rightarrow retv_X(l_C) = false)$$

Lock

Abstract lock. An abstract lock l is an object that encapsulates a state, acquired \mathbb{A} or released \mathbb{R} , and supports the following methods: *lock*: The method call $l.lock()$ changes the state from \mathbb{R} to \mathbb{A} . *unlock*: The method call $l.unlock()$ changes the state from \mathbb{A} to \mathbb{R} . *read*: The method call $l.read()$ returns *true* if the state of *lock* is \mathbb{A} and *false* otherwise. The method calls *lock* and *unlock* are mutating method calls. The method call *read* is an accessor method call.

Definition 6 *The sequential specification of a lock l is the set of sequential histories L of lock, unlock, and read method calls on l where the sub-history of L for mutating methods is an alternating sequence of lock and unlock methods and every read method call in L returns true if the last mutating method call before it in L is a lock and returns false otherwise.*

Lock. A lock is a linearizable instance of the abstract lock type.

Let $Lock$ denote the type of locks.

Now, we present some preliminary definitions and then lemmas about locks.

$$isXLock_{X,lo}(l) = \tag{5.24}$$

$$l \in X \wedge obj_X(l) = lo \wedge name_X(l) = lock$$

$$isXUnlock_{X,lo}(l) = \tag{5.25}$$

$$l \in X \wedge obj_X(l) = lo \wedge name_X(l) = unlock$$

$$isXRead_{X,lo}(l) = \tag{5.26}$$

$$l \in X \wedge obj_X(l) = lo \wedge name_X(l) = read$$

The common usage protocol for locks is that a thread unlocks a lock only if it has already acquired it. Many languages including Java enforce this property of programs by runtime checks. We capture this property as follows.

Definition 7 *A history is owner-respecting for a lock if every thread in the history releases the lock only after it has already acquired it.*

$$isXOwnerRespecting_{lo}(X) = \tag{5.27}$$

$$\forall l: isXUnlock_{X,lo}(l) \Rightarrow$$

$$\exists l': isXLock_{X,lo}(l') \wedge$$

$$thread_X(l') = thread_X(l) \wedge$$

$$l' \prec_X l \wedge$$

$$\forall l'': (isXUnlock_{X,lo}(l'') \wedge thread_X(l'') = thread_X(l)) \Rightarrow (l'' \prec_X l' \vee l \preceq_X l'')$$

Lemma 30 *If l is a lock, X is an owner-respecting history of l and L is the linearization of X , then the sub-history of L for mutating method calls is a sequence of pairs of lock and unlock method calls by the same thread (possibly followed by a lock method call).*

Lemma 31 (Lock) *In an owner-respecting execution for a lock l , if a lock method call by a thread T_1 is linearized before an unlock method call by a thread T_2 , then an unlock method call by T_1 is linearized before a lock method call by T_2 . Formally,*

$$\forall o \in Lock: \forall (X, L) \in \mathbb{H}_L(o): \forall l_{l1}, l_{u2}: \tag{5.28}$$

$$(isXOwnerRespecting_o(X) \wedge$$

$$isXLock_{X,o}(l_{l1}) \wedge$$

$$isXUnlock_{X,o}(l_{u2}) \wedge$$

$$l_{l1} \prec_L l_{u2}) \Rightarrow$$

$$\exists l_{u1}, l_{l2}:$$

$$isXUnlock_{X,o}(l_{u1}) \wedge thread_X(l_{l1}) = thread_X(l_{u1}) \wedge$$

$$isXLock_{X,o}(l_{l2}) \wedge thread_X(l_{l2}) = thread_X(l_{u2}) \wedge$$

$$l_{u1} \prec_L l_{l2}$$

Lemma 32 (LockReadL) *In an owner-respecting execution for a lock l , if a read method call that returns false is linearized before an unlock method call by a thread T , then the read method call is linearized before a lock method call by T . Formally,*

$$\begin{aligned}
& \forall o \in \text{Lock}: \forall (X, L) \in \mathbb{H}_L(o): \forall l_{u1}, l_{r2}: & (5.29) \\
& \quad (isXOwnerRespecting_o(X) \wedge \\
& \quad isXRead_{X,o}(l_{r2}) \wedge \text{retv}_X(l_{r2}) = \text{false} \\
& \quad isXUnlock_{X,o}(l_{u1}) \wedge \\
& \quad l_{r2} \prec_L l_{u1}) \Rightarrow \\
& \exists l_{l1}: \\
& \quad isXLock_{X,o}(l_{l1}) \wedge \text{thread}_X(l_{l1}) = \text{thread}_X(l_{u1}) \wedge \\
& \quad l_{r2} \prec_L l_{l1}
\end{aligned}$$

Lemma 33 (LockReadR) *In an owner-respecting execution for a lock l , if a lock method call by a thread T is linearized before a read method call that returns false, then an unlock method call by T is linearized before the read method call. Formally,*

$$\begin{aligned}
& \forall o \in \text{Lock}: \forall (X, L) \in \mathbb{H}_L(o): \forall l_{l1}, l_{r2}: & (5.30) \\
& \quad (isXOwnerRespecting_o(X) \wedge \\
& \quad isXLock_{X,o}(l_{l1}) \wedge \\
& \quad isXRead_{X,o}(l_{r2}) \wedge \text{retv}_X(l_{r2}) = \text{false} \\
& \quad l_{l1} \prec_L l_{r2}) \Rightarrow \\
& \exists l_{u1}: \\
& \quad isXUnlock_{X,o}(l_{u1}) \wedge \text{thread}_X(l_{l1}) = \text{thread}_X(l_{u1}) \wedge \\
& \quad l_{u1} \prec_L l_{r2}
\end{aligned}$$

Lemma 34 (LockReadM) *In an owner-respecting execution for a lock l , every read method call that is linearized between a pair of matching lock and unlock method calls returns true. Formally,*

$$\begin{aligned}
& \forall o \in \text{Lock}: \forall (X, L) \in \mathbb{H}_L(o): \forall l_{l1}, l_{u1}, l_{r2}: & (5.31) \\
& \quad (isXOwnerRespecting_o(X) \wedge \\
& \quad isXLock_{X,o}(l_{l1}) \wedge \\
& \quad isXUnlock_{X,o}(l_{u1}) \wedge \\
& \quad \text{thread}_X(l_{l1}) = \text{thread}_X(l_{u1}) \wedge \\
& \quad \forall l'_{u1}: (isXUnlock_{X,o}(l'_{u1}) \wedge \text{thread}_X(l_{l1}) = \text{thread}_X(l'_{u1})) \Rightarrow (l'_{u1} \prec_X l_{l1} \vee l_{u1} \preceq_X l'_{u1}) \\
& \quad isXRead_{X,o}(l_{r2}) \wedge \\
& \quad l_{l1} \prec_L l_{r2} \wedge l_{r2} \prec_L l_{u1}) \\
& \Rightarrow \\
& \quad \text{retv}_X(l_{r2}) = \text{true}
\end{aligned}$$

Try-Lock

Abstract Try-lock. A try-lock l is an object that encapsulates an abstract state, acquired \mathbb{A} or released \mathbb{R} , and in addition to *lock*, *unlock* and *read* methods, it supports the *trylock* method. If the state of the *lock* is \mathbb{R} , $l.\text{trylock}()$ changes it to \mathbb{A} and returns *true*. Otherwise, it returns *false*.

We call a *lock* method call or a successful *tryLock* method call, a *successful lock* method call. We call a *lock* method call, successful *tryLock* method call or *unlock* method call, a *mutating* method call.

Definition 8 *The sequential specification of a try-lock l is the set of sequential histories L of lock, unlock, read and tryLock method calls on l with the following conditions: The last mutating method call before a successful lock method call is an unlock method call. Similarly, the last mutating method call before an unlock method call is a successful lock method call. A tryLock method call returns true if the latest preceding mutating method call is an unlock and returns false otherwise. Similarly, A read method call returns true if the latest preceding mutating method call is a successful lock and returns false otherwise.*

Try-Lock. A try-lock is a linearizable instance of the abstract try-lock type.

Let *TryLock* denote the type of try-locks.

Similar to the *Lock* type, after some preliminary definitions, we define the owner-respecting histories and state the *TryLock* type lemmas.

$$isXTryLock_{X,o}(l) = \tag{5.32}$$

$$l \in X \wedge obj_X(l) = o \wedge name_X(l) = tryLock$$

$$isXTLock_{X,o}(l) = \tag{5.33}$$

$$isXLock_{X,o}(l) \vee (isXTryLock_{X,o}(l) \wedge retv_X(l) = true)$$

The intuition for owner-respecting histories remains the same. A history is owner-respecting for a try-lock if every thread in the history releases the lock only after it has already acquired it. The minor difference from the prior definition for locks is that the acquisition of a try-lock is either by a *lock* method call or a successful *tryLock* method call.

$$isXTOwnerRespecting_o(X) = \tag{5.34}$$

$$\forall l: isXUnlock_{X,o}(l) \Rightarrow$$

$$\exists l': isXTLock_{X,o}(l') \wedge$$

$$thread_X(l') = thread_X(l) \wedge$$

$$l' \prec_X l \wedge$$

$$\forall l'': (isXUnlock_{X,o}(l'') \wedge thread_X(l'') = thread_X(l)) \Rightarrow l'' \prec_X l' \vee l \preceq_X l''$$

Lemma 35 *If l is a try-lock, X is an owner-respecting history of l and L is the linearization of X , then the sub-history of L for mutating method calls is a sequence of pairs of successful lock and unlock method calls by the same thread (possibly followed by a successful lock method call).*

Lemma 36 (TryLock) *In an owner-respecting execution for a try-lock l , if a successful lock method call by a thread T_1 is linearized before an unlock method call by a thread T_2 , then an unlock method call by T_1*

is linearized before a successful lock method call by T_2 . Formally,

$$\begin{aligned}
\forall o \in \text{TryLock}: \forall (X, L) \in \mathbb{H}_L(o): \forall l_{l_1}, l_{u_2}: & \quad (5.35) \\
(isXTOwnerRespecting_o(X) \wedge & \\
isXTLock_{X,o}(l_{l_1}) \wedge & \\
isXUnlock_{X,o}(l_{u_2}) \wedge & \\
l_{l_1} \prec_L l_{u_2}) \Rightarrow & \\
\exists l_{u_1}, l_{l_2}: & \\
isXUnlock_{X,o}(l_{u_1}) \wedge thread_X(l_{l_1}) = thread_X(l_{u_1}) \wedge & \\
isXTLock_{X,o}(l_{l_2}) \wedge thread_X(l_{l_2}) = thread_X(l_{u_2}) \wedge & \\
l_{u_1} \prec_L l_{l_2} &
\end{aligned}$$

Lemma 37 (TryLockReadL) *In an owner-respecting execution for a try-lock l , a read method call that returns false is linearized before if an unlock method call by a thread T then the read method call is linearized before a successful lock method call by T . Formally,*

$$\begin{aligned}
\forall o \in \text{TryLock}: \forall (X, L) \in \mathbb{H}_L(o): \forall l_{u_1}, l_{r_2}: & \quad (5.36) \\
(isXTOwnerRespecting_o(X) \wedge & \\
isXRead_{X,o}(l_{r_2}) \wedge retv_X(l_{r_2}) = false & \\
isXUnlock_{X,o}(l_{u_1}) \wedge & \\
l_{r_2} \prec_L l_{u_1}) \Rightarrow & \\
\exists l_{l_1}: & \\
isXTLock_{X,o}(l_{l_1}) \wedge thread_X(l_{l_1}) = thread_X(l_{u_1}) \wedge & \\
l_{r_2} \prec_L l_{l_1} &
\end{aligned}$$

Lemma 38 (TryLockReadR) *In an owner-respecting execution for a try-lock l , if a successful lock method call by a thread T is linearized before a read method call that returns false, then an unlock method call by T is linearized before the read method call. Formally,*

$$\begin{aligned}
\forall o \in \text{TryLock}: \forall (X, L) \in \mathbb{H}_L(o): \forall l_{l_1}, l_{r_2}: & \quad (5.37) \\
(isXTOwnerRespecting_o(X) \wedge & \\
isXTLock_{X,o}(l_{l_1}) \wedge & \\
isXRead_{X,o}(l_{r_2}) \wedge retv_X(l_{r_2}) = false & \\
l_{l_1} \prec_L l_{r_2}) \Rightarrow & \\
\exists l_{u_1}: & \\
isXUnlock_{X,o}(l_{u_1}) \wedge thread_X(l_{l_1}) = thread_X(l_{u_1}) \wedge & \\
l_{u_1} \prec_L l_{r_2} &
\end{aligned}$$

Lemma 39 (TryLockReadM) *In an owner-respecting execution for a try-lock l , every read method call*

that is linearized between a pair of matching successful and unlock method calls returns true. Formally,

$$\begin{aligned}
& \forall o \in \text{TryLock}: \forall (X, L) \in \mathbb{H}_L(o): \forall l_{l1}, l_{u1}, l_{r2}: & (5.38) \\
& \quad (\text{isXOwnerRespecting}_o(X) \wedge \\
& \quad \text{isXTLock}_{X,o}(l_{l1}) \wedge \\
& \quad \text{isXUnlock}_{X,o}(l_{u1}) \wedge \\
& \quad \text{thread}_X(l_{l1}) = \text{thread}_X(l_{u1}) \wedge \\
& \quad \forall l'_{u1}: (\text{isXUnlock}_{X,o}(l'_{u1}) \wedge \text{thread}_X(l_{l1}) = \text{thread}_X(l'_{u1})) \Rightarrow (l'_{u1} \prec_X l_{l1} \vee l_{u1} \preceq_X l'_{u1}) \\
& \quad \text{isXRead}_{X,o}(l_{r2}) \wedge \\
& \quad l_{l1} \prec_L l_{r2} \wedge l_{r2} \prec_L l_{u1}) \\
& \Rightarrow \\
& \quad \text{retv}_X(l_{r2}) = \text{true}
\end{aligned}$$

Seq-Lock

Abstract seq-lock. A seq-lock l is an object that encapsulates a number and an abstract state, acquired \mathbb{A} or released \mathbb{R} . It supports the *read*, *compareAndLock* and *incAndUnlock* methods. The method call $l.\text{read}()$ returns the pair of the encapsulated number and *true* if the state of *lock* is \mathbb{A} and *false* otherwise. The method call $l.\text{compareAndLock}(n)$ compares the the encapsulated number with n and if they are equal, changes the state from \mathbb{R} to \mathbb{A} and returns *true*. Otherwise, it does not change the state of the seq-lock and returns *false*. The method call $l.\text{incAndUnlock}()$ increments the encapsulated number and changes the state from \mathbb{A} to \mathbb{R} .

A successful *compareAndLock* and *incAndUnlock* are mutating method calls. The method call *read* is an accessor method call.

Definition 9 *The sequential specification of a seq-lock l is the set of sequential histories L of read, compareAndLock, and incAndUnlock method calls on l with the following conditions:*

Every read method call returns the pair of the number of incAndUnlock method calls before it and true if the last mutating method call before it is a successful compareAndLock and false otherwise.

A compareAndLock method call returns true if the last mutating method call before it is an incAndUnlock method call and the number of incAndUnlock method calls before it is equal to its argument. It returns false otherwise.

The last mutating method call before an incAndUnlock method call is a successful compareAndLock method call.

Seq-Lock. A seq-lock is a linearizable instance of the abstract seq-lock type.

Let *SeqLock* denote the type of seq-locks.

Counter

Abstract Counter: A counter c is an object that encapsulates a number and supports the following two methods: The method call $c.\text{read}()$ returns the current value of c . The method call $c.\text{iaf}()$ increments the value of c and returns the incremented value.

Definition 10 *The sequential specification of a counter c is the set of sequential histories of read and iaf method calls on c where every method call returns the number of iaf method calls before it (including the method call itself). Note that it is assumed that the initial value of the counter is zero.*

Strong Counter. A strong counter is a linearizable instance of abstract counter type.

Let $SCounter$ denote the type of strong counters.

Lemma 40 (SCounter) *The return value of every method call that is linearized before an iaf method call is smaller than the return value of the iaf method call. Formally,*

$$\begin{aligned} \forall c \in SCounter: \forall (X, C) \in \mathbb{H}_L(c): \forall l, l': & \\ l \in X \wedge l' \in X \wedge name_X(l') = iaf \wedge l \prec_C l' & \\ \Rightarrow & \\ retv_X(l) < retv_X(l') & \end{aligned} \quad (5.39)$$

Set

A set s is an object that represents a set of values and supports the following methods: *add*: The method call $s.add(v)$ adds value v to set s . *contains*: The method call $s.contains(v)$ returns *true* if v is a member of s and *false* otherwise.

Definition 11 *The sequential specification of a set s is the set of sequential histories of add and contains method calls on s where every contains method call returns true if there is a preceding add method call with the same argument, and returns false otherwise. Note that it is assumed that the set is initially empty.*

Basic Set. A basic set is a basic instance of set type.

Let $BasicSet$ denote the type of basic sets.

Let us define

$$isXContains_{X,s}(l) = \quad (5.40)$$

$$l \in X \wedge obj_X(l) = s \wedge name_X(l) = contains$$

$$isXAdd_{X,s}(l) = \quad (5.41)$$

$$l \in X \wedge obj_X(l) = s \wedge name_X(l) = add$$

Lemma 41 (BasicSetContains) *In every sequential execution on a basic set, for every contains method call that returns true, there is a preceding add method call with the same argument. Formally,*

$$\forall s \in BasicSet: \forall X \in \mathbb{H}_B(s): X \in Sequential \Rightarrow \quad (5.42)$$

$$\forall l_c: isXContains_{X,s}(l_c) \wedge retv_X(l_c) = true \Rightarrow$$

$$\exists l_a: isXAdd_{X,s}(l_a) \wedge$$

$$arg1(l_a) = arg1(l_c) \wedge l_a \prec_X l_c$$

Lemma 42 (BasicSetAdd) *In every sequential execution on a basic set, every contains method call that succeeds an add method call with the same argument returns true. Formally,*

$$\forall s \in BasicSet: \forall X \in \mathbb{H}_B(s): X \in Sequential \Rightarrow \quad (5.43)$$

$$\forall l_c, l_a:$$

$$isXContains_{X,s}(l_c) \wedge$$

$$isXAdd_{X,s}(l_a) \wedge$$

$$arg1(l_a) = arg1(l_c) \wedge l_a \prec_X l_c$$

$$\Rightarrow$$

$$retv_X(l_c) = true$$

Map

A map m is an object that represents a mapping from a set of keys to a set of values and supports the following methods: *put*: The method call $m.put(k, v)$ adds or updates the mapping of the key k to the value v ($v \neq \perp$) in the map m . *get*: The method call $m.get(k)$ returns the value that the map m associates with the key k . It returns \perp if m does not map k .

Definition 12 *The sequential specification of a map m is the set of sequential histories of put and get method calls on m where every get method call returns \perp if there is no preceding put method call with the same key argument; otherwise it returns the second argument of the latest preceding put method call with the same key argument. Note that it is assumed that the map is initially empty.*

Basic Map. A basic set is a basic instance of map type.

Let *BasicMap* denote the type of basic maps.

Let us define

$$isXGet_{X,m}(l) = \tag{5.44}$$

$$l \in X \wedge obj_X(l) = m \wedge name_X(l) = get$$

$$isXPut_{X,m}(l) = \tag{5.45}$$

$$l \in X \wedge obj_X(l) = m \wedge name_X(l) = put$$

$$isXPutter_{X,m}(l_p, l_g) \Leftrightarrow \tag{5.46}$$

$$isXPut_{X,m}(l_p) \wedge arg1_X(l_p) = arg1_X(l_g) \wedge l_p \prec_X l_g \wedge \tag{5.47}$$

$$\forall l'_p: isXPut_{X,m}(l'_p) \wedge arg1_X(l'_p) = arg1_X(l_g) \Rightarrow (l'_p \preceq_X l_p \vee l_g \prec_X l'_p) \tag{5.48}$$

Lemma 43 (BasicMapGet) *In every sequential execution on a basic map, the return value of every get method call that does not return \perp is equal to the value argument of the latest preceding put method call with the same key argument. Formally,*

$$\forall m \in BasicMap: \forall X \in \mathbb{H}_B(m): X \in Sequential \Rightarrow \tag{5.49}$$

$$\forall l_g: isXGet_{X,m}(l_g) \wedge \neg(retv_X(l_g) = \perp) \Rightarrow$$

$$\exists l_p: isPutter_{X,m}(l_p, l_g) \wedge$$

$$arg2_X(l_p) = retv_X(l_g)$$

Lemma 44 (BasicMapPut) *In every sequential execution on a basic map, for every get method call g , if p is the latest preceding put method call with the same key argument then the return value of g is equal to the value argument of p . Formally,*

$$\forall m \in BasicMap: \forall X \in \mathbb{H}_B(m): X \in Sequential \Rightarrow \tag{5.50}$$

$$\forall l_g, l_p:$$

$$isXGet_{X,m}(l_g) \wedge$$

$$isPutter_{X,m}(l_p, l_g) \wedge$$

$$\Rightarrow$$

$$retv_X(l_g) = arg2_X(l_p)$$

Proof Sketches.

Lemma 19:

Straightforward from $\prec_X \subseteq \prec_L$.

Lemma 20:

We have

$$(1) \ l \prec_L l'$$

From [1], we have

$$(2) \ rEv(l) \triangleleft_L iEv(l')$$

From the well-formedness of the history O , we have

$$(3) \ iEv(l) \triangleleft_L rEv(l)$$

$$(4) \ iEv(l') \triangleleft_L rEv(l')$$

From [3], [2] and [4], we have

$$(5) \ iEv(l) \triangleleft_L rEv(l')$$

From [5], we have

$$(6) \ \neg(rEv(l') \triangleleft_L iEv(l))$$

From [2] and [6], we have

$$(7) \ \neg(l' = l)$$

From the definition of \prec_X on [6], we have

$$(8) \ \neg(l' \prec_L l)$$

The conclusion is

$$[8] \text{ and } [7]$$

Lemma 21:

Straightforward from the fact that L is a member of sequential specification and a sequential specification is a set of sequential histories and the execution order is total in sequential histories.

Lemma 22:

Straightforward from the fact that L is a member of sequential specification and a sequential specification is a set of sequential histories and the execution order is total in sequential histories.

We have

$$(1) \ l \in X$$

$$(2) \ l' \in X$$

$$(3) \ X \equiv L$$

$$(4) \ L \in SeqSpec(o)$$

From [4], we have

$$(5) \ L \in Sequential$$

From [3], [1] and [2], we have

$$(6) \ l \in L$$

$$(7) \ l' \in L$$

From [4], [6] and [7], we have

$$l \prec_L l' \vee l' \prec_L l \vee l = l'$$

Lemma 23:

Straightforward from the fact that L is equivalent to X .

We have

$$(1) \ X \equiv L$$

(2) $L \in SeqSpec(o)$

(3) $l \prec_L l'$

From [3], we have

(4) $l \in L$

(5) $l' \in L$

From [2] on [4] and [5], we have

(6) $obj_L(l) = o$

(7) $obj_L(l') = o$

From [1] on [4] and [5], we have

$l \in X$

$l' \in X$

From [1] on [6] and [7], we have

$obj_X(l) = o$

$obj_X(l') = o$

Lemma 24:

Using L2X and XTotal, we have four cases:

Case: $l \prec l'$

Straightforward from XTrans.

Case: $l \sim l'$

Straightforward from XXTrans.

Case: $l' \prec l$

Straightforward from X2L and LASym.

Case: $l' = l$

Straightforward from LASym.

Lemma 25:

Derived from the semantics of basic objects (Definition 1) and the sequential specification of register (Definition 3).

Lemma 26:

Derived from the semantics of basic register (Definition 4).

Lemma 27:

This is a restatement of Theorem 3 from the original definition of linearizability []. Derivable from the semantics of linearizable objects (Definition 2) and the sequential specification of register (Definition 3).

Lemma 28:

Derivable from the semantics of linearizable objects (Definition 2) and the sequential specification of cas register (Definition 5).

Lemma 29:

Derivable from the semantics of linearizable objects (Definition 2) and the sequential specification of cas register (Definition 5).

Lemma 30:

Derivable from the semantics of linearizable objects (Definition 2), the sequential specification of the lock

(Definition 6), the owner-respecting property (Definition 7), and that the sub-history for each thread is sequential (from the definition of execution histories).

Lemma 31:

Derived from Lemma 30.

Lemma 32:

Derived from Lemma 30 and the sequential specification of lock (Definition 6).

Lemma 33:

Derived from Lemma 30 and the sequential specification of lock (Definition 6).

Lemma 34:

Derived from Lemma 30 and the sequential specification of lock (Definition 6).

Lemma 35:

Derivable from the semantics of linearizable objects (Definition 2), the sequential specification of the lock (Definition 8), the owner-respecting property (Definition 35), and that the sub-history for each thread is sequential (from the definition of execution histories).

Lemma 36:

Derived from Lemma 35.

Lemma 37:

Derived from Lemma 35 and the sequential specification of try-lock (Definition 8).

Lemma 38:

Derived from Lemma 35 and the sequential specification of try-lock (Definition 8).

Lemma 39:

Derived from Lemma 35 and the sequential specification of try-lock (Definition 8).

Lemma 40:

Derivable from the semantics of linearizable objects (Definition 2), the sequential specification of counter (Definition 10).

Lemma 41:

Derivable from the semantics of basic objects (Definition 1), the sequential specification of set (Definition 11).

Lemma 42:

Derivable from the semantics of basic objects (Definition 1), the sequential specification of set (Definition 11).

Lemma 43:

Derivable from the semantics of basic objects (Definition 1), the sequential specification of set (Definition 12).

Lemma 44:

Derivable from the semantics of basic objects (Definition 1), the sequential specification of set (Definition 12).

6 Marking TL2

<pre> <i>reg</i>: BasicRegister[[<i>I</i>]], <i>ver</i>: AtomicRegister[[<i>I</i>]], <i>lock</i>: TryLock[[<i>I</i>]], <i>clock</i>: SCounter, </pre>	<pre> <i>rver</i>: ThreadLocal BasicRegister, <i>rset</i>: ThreadLocal BasicSet, <i>wset</i>: ThreadLocal BasicMap, <i>lset</i>: ThreadLocal BasicSet </pre>
<pre> def <i>init</i>_{<i>t</i>}() I01 ▷ <i>snap</i> = <i>clock.read</i>(), I02 ▷ <i>rver</i>[<i>t</i>].<i>write</i>(<i>snap</i>), I03 ▷ return <i>ok</i>, </pre>	<pre> def <i>commit</i>_{<i>t</i>}() C01 ▷ foreach (<i>i</i> ∈ <i>wset</i>[<i>t</i>]) C02_{<i>i</i>} ▷ <i>locked</i> = <i>lock</i>[<i>i</i>].<i>trylock</i>(), if (¬<i>locked</i>) C03_{<i>i</i>} ▷ <i>lset.add</i>(<i>i</i>) else C04_{<i>i</i>} ▷ foreach (<i>j</i> ∈ <i>lset</i>) C05_{<i>ij</i>} ▷ <i>lock</i>[<i>j</i>].<i>unlock</i>(), C06_{<i>i</i>} ▷ return ⊆, </pre>
<pre> def <i>read</i>_{<i>t</i>}(<i>i</i>) R01 ▷ <i>pv</i> = <i>wset</i>[<i>t</i>].<i>get</i>(<i>i</i>), if (<i>pv</i> ≠ ⊥) R02 ▷ return <i>pv</i>, </pre>	<pre> C07 ▷ <i>wver</i> = <i>clock.iaf</i>(), C08 ▷ <i>sver</i> = <i>rver</i>[<i>t</i>].<i>read</i>(), if (<i>wver</i> ≠ <i>sver</i> + 1) C09 ▷ foreach (<i>i</i> ∈ <i>rset</i>[<i>t</i>]) C10_{<i>i</i>} ▷ <i>l</i> = <i>lock</i>[<i>i</i>].<i>read</i>(), C11_{<i>i</i>} ▷ <i>s</i> = <i>ver</i>[<i>i</i>].<i>read</i>(), if (¬(¬<i>l</i> ∧ <i>s</i> ≤ <i>sver</i>)) C12_{<i>i</i>} ▷ foreach (<i>j</i> ∈ <i>lset</i>) C13_{<i>ij</i>} ▷ <i>lock</i>[<i>j</i>].<i>unlock</i>(), C14_{<i>i</i>} ▷ return ⊆, </pre>
<pre> R03 ▷ <i>s</i>₁ = <i>ver</i>[<i>i</i>].<i>read</i>(), R04 ▷ <i>v</i> = <i>reg</i>[<i>i</i>].<i>read</i>(), R05 ▷ <i>l</i> = <i>lock</i>[<i>i</i>].<i>read</i>(), R06 ▷ <i>s</i>₂ = <i>ver</i>[<i>i</i>].<i>read</i>(), R07 ▷ <i>sver</i> = <i>rver</i>[<i>t</i>].<i>read</i>(), if (¬(¬<i>l</i> ∧ <i>s</i>₁ = <i>s</i>₂ ∧ <i>s</i>₂ ≤ <i>sver</i>)) R08 ▷ return ⊆, </pre>	<pre> C15 ▷ foreach ((<i>i</i>, <i>v</i>) ∈ <i>wset</i>[<i>t</i>]) C16_{<i>i</i>} ▷ <i>reg</i>[<i>i</i>].<i>write</i>(<i>v</i>), C17_{<i>i</i>} ▷ <i>ver</i>[<i>i</i>].<i>write</i>(<i>wver</i>), C18_{<i>i</i>} ▷ <i>lock</i>[<i>i</i>].<i>unlock</i>(), </pre>
<pre> R09 ▷ <i>rver</i>[<i>t</i>].<i>add</i>(<i>i</i>), R10 ▷ return <i>v</i>, {R03 → R04, R04 → R05, R05 → R06}, </pre>	<pre> C19 ▷ return ⊆, {C01 → C07, C10 → C11, C09 → C15, C16 → C17, C17 → C18}, </pre>
<pre> def <i>write</i>_{<i>t</i>}(<i>i</i>, <i>v</i>) W01 ▷ <i>wset</i>[<i>t</i>].<i>put</i>(<i>i</i>, <i>v</i>), W02 ▷ return <i>ok</i>, </pre>	
<pre> def <i>abort</i>_{<i>t</i>}() A01 ▷ return ⊆, </pre>	

Figure 4: TL2 Algorithm Specification

Atomic register, try-lock and strong counter are linearizable object types and basic register, basic set and basic map are basic object types. (At a high level, for every execution on a linearizable object, there is an equivalent sequential execution that complies with the sequential specification of the object. On the other hand, a basic object complies with its sequential specification only if it is accessed sequentially.) TL2 uses the following base objects: Value registers *reg*: an array of basic registers. Version registers *ver*: an array of atomic registers with the initial value 0. Locks *lock*: an array of try-locks that are initially released. The arrays are of size memory location count $|I|$.² Global version clock *clock*: a strong counter with the initial value 0. A strong counter provides two methods in its interface: *iaf* (inc-and-fetch) that increments the counter and returns the counter value and *read* that returns the counter value. Read version *rver*: a thread-local basic register. Read set *rset*: a thread-local basic set that is initially \emptyset . Write set *wset*: a thread-local basic map that is initially \emptyset . Lock set *lset*: a thread-local basic set that is initially \emptyset . As relaxed execution may reorder program statements, any order that is not implied by the data or control dependencies but is required for the correctness of the algorithm is explicitly declared at the end of each method definition. The values *ok*, \mathbb{A} , \mathbb{C} are reserved to denote successful completion of writes and abortion and commitment of transactions respectively.

TL2 is a deferred-update TM algorithm. A value that a transaction t writes to a location is buffered in the write set $wset[t]$ at $W01$ and is written back to register $reg[i]$ at $C16_i$ while t is committing. TL2 records a version in the register $ver[i]$ for the value stored in the register $reg[i]$. The version register $ver[i]$ is updated to ascending numbers at $C17_i$ after new values are written back to $reg[i]$ at $C16_i$. The try-lock $lock[i]$ is used for exclusive access to the registers for location i . At commit, the lock $lock[i]$ of each location i in the write set $wset[t]$ is acquired at $C01$ to $C06$. (If a lock cannot be acquired, the previously acquired locks are released at $C05$ and the transaction is aborted at $C06$.) Then, a new snapshot number is read from *clock* at $C07$. Then, for each location in the read set $rset[t]$, first $lock[i]$ and then $ver[i]$ are read at $C10_i$ and $C11_i$ and the read is validated. (If a read is not validated, the acquired locks are released at $C13$ and the transaction is aborted at $C14$.) Finally, the value buffered for each location i in $wset[t]$ is written back at $C15_i$ to $C18_i$. For each pair in the write set $wset[t]$, the following three operations are executed in order. First, the buffered value is written back to $reg[i]$, then $ver[i]$ is updated, and then $lock[i]$ is released. In the *init* method, each transaction t reads the current snapshot version from *clock* at $I01$ and writes it to the read version register $rver[t]$ at $I02$. The read version is read at $R07$ and $C08$ to validate the read values. To read a location i , a transaction reads $ver[i]$, $reg[i]$, $lock[i]$ and again $ver[i]$ in order at $R03$ to $R06$ and then validates the read. (If the validation fails, the transaction is aborted.) Finally, i is added to the read set $rset[t]$ and the read value is returned.

²As observed by previous work [3], in the original TL2 paper, the authors maintain the version number and the lock bit of every location in the same memory word, thus, the order of reading the lock and the version register in the commit method is ambiguous. In our specification, we treat the lock and the version as separate registers and make the orders explicit.

Notation. Let us remind the notation. Consider an execution history H . We use $l_1 \prec_H l_2$ to denote that l_1 is executed before l_2 . We use $l_1 \sim_H l_2$ to denote that l_1 is executed concurrently to l_2 . We use $l_1 \lesssim_H l_2$ to denote that l_1 is executed before or concurrently to l_2 . We use \prec_{clock} , $\prec_{ver[i]}$ and $\prec_{lock[i]}$ to denote the linearization order of $clock$, $ver[i]$ and $lock[i]$ respectively.

A label c_1c_2 is a call string that denotes a method call labeled c_2 that is executed in the body of the method call labeled c_1 .

We use $initOf_H(T)$ and $commitOf_H(T)$ to denote the *init* and *commit* method calls of transaction T in history H .

Marking Relation. Now, we define the marking relation for TL2. The effect order of transactions is the linearization order of their calls to the *clock*. Every transaction reads an initial snapshot number at $I01$. A committing transaction makes a new snapshot at $C07$. A TL2 transaction takes effect at $C07$ if it is committed and at $I01$ otherwise. The access order of read operations and writer transactions to location i is the execution order of their accesses to the $reg[i]$ register. The read method reads $reg[i]$ at $R04$ and a writer transaction writes to $reg[i]$ at $C16_i$.

Definition 13 (Marking TL2) Consider an execution history $H \in \mathbb{H}(TL2)$. Let

$$\begin{aligned} readAcc(R) &= R'R04 \\ writeAcc(T, i) &= commitOf_H(T)'C16_i \\ Eff(T) &= \begin{cases} initOf_H(T)'I01 & \text{if } T \in Aborted(H) \\ commitOf_H(T)'C07 & \text{if } T \in Committed(H) \end{cases} \end{aligned}$$

The marking \sqsubseteq for H is the reflexive closure of \sqsubset that is define as follows:

$$\begin{aligned} &\{(T, T') \mid T, T' \in Trans(H) \wedge Eff(T) \prec_{clock} Eff(T')\} \cup \\ &\{(T, R) \mid \exists i: R \in GlobalTReads(H), i = arg1(R), T \in Writers_H(i) \wedge writeAcc(T, i) \prec_H readAcc(R)\} \cup \\ &\{(R, T) \mid \exists i: R \in GlobalTReads(H), i = arg1(R), T \in Writers_H(i) \wedge readAcc(R) \lesssim_H writeAcc(T, i)\} \end{aligned}$$

We have formally proved the markability of TL2 using a novel program logic that facilitates reasoning about execution and linearization orders. To keep the focus of this paper on markability, we avoid the formal presentation of the logic and present a simplified reasoning.

In addition to the lemmas presented in the previous section, we use the rule P2X that states the program-order-preservation property. If a method call l_1 is ordered before a method call l_2 in the program, and methods l_1 and l_2 are executed, then l_1 is executed before l_2 .

Lemma 45 *TL2 preserves reads of aborted transactions (part 1).*

$$\begin{aligned} & \forall H \in \mathbb{H}(TL2): \\ & \forall R \in \text{GlobalTReads}(H): \text{Let } i = \text{arg1}_H(R), T = \text{trans}_H(R): \\ & T \in \text{Aborted}(H) \Rightarrow \text{NoWriterBetween}_{H,i}(R, \sqsubseteq, T) \end{aligned}$$

Proof Sketch.

T	T'
	$C02_i \triangleright \text{lock}[i].\text{trylock}()$
	...
	$C07 \triangleright \text{wver} = \text{clock.iaf}()$
$I01 \triangleright \text{snap} = \text{clock.read}()$...
...	
$R03 \triangleright s_1 = \text{ver}[i].\text{read}()$	
$R04 \triangleright v = \text{reg}[i].\text{read}()$	
	$C16_i \triangleright \text{reg}[i].\text{write}(v)$
	$C17_i \triangleright \text{ver}[i].\text{write}(wver)$
$R05 \triangleright l = \text{lock}[i].\text{read}()$	$C18_i \triangleright \text{lock}[i].\text{unlock}()$
$R06 \triangleright s_2 = \text{ver}[i].\text{read}()$	
$R07 \triangleright sver = \text{rver}[t].\text{read}()$	
if $(\neg(\neg l \wedge s_1 = s_2 \wedge s_2 \leq sver))$	
return \mathbb{A}	

Figure 5: Case $T \in \text{Aborted}(H) \wedge R \sqsubset T' \sqsubset T$

We consider an aborted transaction T with an unaborted global read operation R from a location i and a writer T' of i .

We assume that

T' accesses i after R

that is

$$(1) T' \sqsubset R$$

and

T' takes effect before T

that is

$$(2) T' \sqsubset T$$

We show that

TL2 aborts R .

Figure 5 depicts the two transactions.

By Definition 13 on [1], we have

$$(3) R04 \prec_H C16_i$$

By Definition 13 on [2], we have

$$(4) C07 \prec_{\text{clock}} I01$$

The method calls $R05$ and $C18_i$ are on the object $\text{lock}[i]$. We consider two cases for the linearization order of them and prove that R returns \mathbb{A} in both cases.

Case 1:

$$(5) R05 \prec_{\text{lock}[i]} C18_i$$

By P2X on the algorithm, we have

$$(6) C02_i \prec_H C07$$

$$(7) I01 \prec_H R05$$

By the Lemma XLTRANS on [6], [4] and [7], we have

$$C02_i \prec_H R05$$

thus, by the Lemma X2L, we have

$$(8) C02_i \prec_{lock[i]} R05$$

By the Lemma TRYLOCKREADM on [8] and [5], we have that

$$R05 \text{ returns } true \text{ i.e. } l = true$$

Thus,

The validation check fails and R returns \mathbb{A} .

Case 2:

$$(9) C18_i \prec_{lock[i]} R05$$

By P2X on the algorithm, we have

$$(10) C17_i \prec_H C18_i$$

$$(11) R05 \prec_H R06$$

By the Lemma XLTRANS on [10], [9] and [11], we have

$$C17_i \prec_H R06$$

Thus, by the Lemma X2L, we have

$$(12) C17_i \prec_{ver[i]} R06$$

By Lemma 54 on [12], we have

$$(13) wver \leq s_2$$

By P2X on the algorithm, we have

$$(14) R03 \prec_H R04$$

$$(15) C16_i \prec_H C17_i$$

By the Lemma XXTRANS on [14], [3] and [15], we have

$$R03 \prec_H C17_i$$

Thus, by the Lemma X2L, we have

$$(16) R03 \prec_{ver[i]} C17_i$$

By Lemma 54 on [16], we have

$$(17) s_1 < wver$$

From [13] and [17], we have

$$\neg(s_1 = s_2)$$

Thus,

The validation check fails and R returns \mathbb{A} in this case too.

□

Lemma 46 *TL2 preserves reads of aborted transactions (part 2).*

$$\forall H \in \mathbb{H}(TL2):$$

$$\forall R \in GlobalTReads(H): \text{ Let } i = arg1_H(R), T = trans_H(R):$$

$$T \in Aborted(H) \Rightarrow NoWriterBetween_{H,i}(T, \sqsubseteq, R)$$

Proof Sketch.

We consider an aborted transaction T with an unaborted global read operation R from a location i and a writer T' of i .

We assume that

T	T'
$I01 \triangleright$ $snap = clock.read()$	$C02_i \triangleright$ $lock[i].trylock()$
$I02 \triangleright$ $rver[t].write(snap)$	
	$C07 \triangleright$ $wver = clock.iaf()$
	$C16_i \triangleright$ $reg[i].write(v)$
$R04 \triangleright$ $v = reg[i].read()$	$C17_i \triangleright$ $ver[i].write(wver)$
$R05 \triangleright$ $l = lock[i].read()$	$C18_i \triangleright$ $lock[i].unlock()$
$R06 \triangleright$ $s_2 = ver[i].read()$	
$R07 \triangleright$ $sver = rver[t].read()$	
if $(\neg(\neg l \wedge s_1 = s_2 \wedge s_2 \leq sver))$ return \mathbb{A}	

Figure 6: Case $T \in Aborted(H) \wedge T \sqsubset T' \sqsubset R$

T' takes effect after T

that is

$$(1) T \sqsubset T'$$

and

T' accesses i before R

that is

$$(2) T' \sqsubset R$$

We show that

TL2 aborts R .

Figure 6 depicts the two transactions.

By Definition 13 on [1], we have

$$(3) I01 \prec_{clock} C07$$

By Definition 13 on [2], we have

$$(4) C16_i \preceq R04$$

The method calls $R05$ and $C18_i$ are on the object $lock[i]$. We consider two cases for the linearization order of them and prove that R returns \mathbb{A} in both cases.

Case 1:

$$(5) R05 \prec_{lock[i]} C18_i$$

By P2X on the algorithm, we have

$$(6) C02_i \prec_H C16_i$$

$$(7) R04 \prec_H R05$$

By the Lemma XXTRANS on [6], [4] and [7], we have

$$C02_i \prec_H R05$$

thus, by the Lemma X2L, we have

$$(8) C02_i \prec_{lock[i]} R05$$

By the Lemma TRYLOCKREADM on [8] and [5], we have that

$R05$ returns *true* i.e. $l = true$.

Thus,

The validation check fails and R returns \mathbb{A} .

Case 2:

$$(9) \ C18_i \prec_{lock[i]} R05$$

By P2X on the algorithm, we have

$$(10) \ C17_i \prec_H C18_i$$

$$(11) \ R05 \prec_H R06$$

By the Lemma XLTRANS on [10], [9] and [11], we have

$$C17_i \prec_H R06$$

Thus, by the Lemma X2L, we have

$$(12) \ C17_i \prec_{ver[i]} R06$$

By Lemma 53 on [12], we have

$$(13) \ wver \leq s_2$$

By the Lemma SCOUNTER on [3], we have

$$(14) \ snap < wver$$

The value of $sver$ is read at $R07$ from $rver$.

The thread-local register $rver$ is only assigned at $I02$ to $snap$.

Thus, we have

$$(15) \ snap = sver$$

From [13], [14] and [15], we have

$$sver > s_2$$

Thus,

The validation check fails and R returns \mathbb{A} in this case too.

□

Lemma 47 *TL2 preserves reads of aborted transactions.*

$$\forall H \in \mathbb{H}(TL2):$$

$$\forall R \in GlobalTReads(H): \text{Let } i = arg1_H(R), T = trans_H(R):$$

$$T \in Aborted(H) \Rightarrow$$

$$NoWriterBetween_{H,i}(R, \sqsubseteq, T) \wedge NoWriterBetween_{H,i}(T, \sqsubseteq, R)$$

Proof. Immediate from Lemma 45 and Lemma 46.

□

Lemma 48 *TL2 preserves reads of committed transactions (part 1).*

$$\forall H \in \mathbb{H}(TL2):$$

$$\forall R \in GlobalTReads(H): \text{Let } i = arg1_H(R), T = trans_H(R):$$

$$T \in Committed(H) \Rightarrow$$

$$NoWriterBetween_{H,i}(R, \sqsubseteq, T)$$

Proof Sketch.

We consider a committed transaction T with an unaborted global read operation R from a location i and a writer T' of i .

We assume that

T' accesses i after R

that is

$$(1) \ R \sqsubset T'$$

and

T' takes effect before T

that is

T	T'
	$C02'_i \triangleright \text{lock}[i].\text{trylock}()$
	...
$I01 \triangleright \text{snap} = \text{clock.read}()$	$C07' \triangleright \text{wver}' = \text{clock.iaf}()$
$I02 \triangleright \text{rver}[t].\text{write}(\text{snap})$...
...	
$R04 \triangleright v = \text{reg}[i].\text{read}()$	
...	$C16'_i \triangleright \text{reg}[i].\text{write}(v')$
$C07 \triangleright \text{wver} = \text{clock.iaf}()$	
...	
$C08 \triangleright \text{sver} = \text{rver}[t].\text{read}()$	
if ($\text{wver} \neq \text{sver} + 1$)	$C17'_i \triangleright \text{ver}[i].\text{write}(\text{wver}')$
$C10_i \triangleright l = \text{lock}[i].\text{read}()$	$C18'_i \triangleright \text{lock}[i].\text{unlock}()$
$C11_i \triangleright s = \text{ver}[i].\text{read}()$	
if ($\neg(\neg l \wedge s \leq \text{sver})$)	
foreach ($j \in \text{lset}$)	
$\text{lock}[j].\text{unlock}()$	
return \mathbb{A}	

Figure 7: Case $T \in \text{Committed}(H) \wedge R \sqsubset T' \sqsubset T$

(2) $T' \sqsubset T$

We show that

TL2 aborts R .

Figure 7 depicts the two transactions. We annotate the labels and variables of T' by a prime so that they do not conflict with the labels and variables of T .

By Definition 13 on [1], we have

(3) $R04 \prec_H C16_i$

By Definition 13 on [2], we have

(4) $C07' \prec_{\text{clock}} C07$

The method calls $I01$ and $C07'$ are on the object clock . We consider two cases for the linearization order of them.

Case 1:

(5) $C07' \prec_{\text{clock}} I01$

From [5] and [3],

The proof of this case reduces to the proof of Lemma 45.

Case 2:

(6) $I01 \prec_{\text{clock}} C07'$

By the Lemma SCOUNTER on [4], we have

(7) $\text{wver}' < \text{wver}$

By the Lemma SCOUNTER on [6], we have

(8) $\text{snap} < \text{wver}'$

The value of sver is read at $R07$ from rver .

The thread-local register rver is only assigned at $I02$ to snap .

Thus, we have

(9) $snap = sver$

From [8] and [9], we have

(10) $sver < wver'$

From [10] and [7], we have

(11) $wver \neq sver + 1$

Thus,

The if branch is taken.

The method calls $C10_i$ and $C18'_i$ are on the object $lock[i]$.

We consider two cases for the linearization order of them.

Case 2.1:

(12) $C10_i \prec_{lock[i]} C18'_i$

By P2X on the algorithm, we have

(13) $C02'_i \prec_H C07'$

(14) $C07 \prec_H C10_i$

By the Lemma XLTRANS on [13], [4] and [14], we have

$C02'_i \prec_H C10_i$

thus, by the Lemma X2L, we have

(15) $C02'_i \prec_{lock[i]} C10_i$

By the Lemma TRYLOCKREADM on [15] and [12], we have that

$R05$ returns *true* i.e. $l = true$

Thus,

The validation check fails and R returns \mathbb{A} .

Case 2.2:

(16) $C18'_i \prec_{lock[i]} C10_i$

By P2X on the algorithm, we have

(17) $C17'_i \prec_H C18'_i$

(18) $C10_i \prec_H C11_i$

By the Lemma XLTRANS on [17], [16] and [18], we have

$C17'_i \prec_H C11_i$

Thus, by the Lemma X2L, we have

(19) $C17'_i \prec_{ver[i]} C11_i$

By Lemma 54 on [19], we have

(20) $wver' \leq s$

From [10], [20], we have

$sver < s$

Thus,

The validation check fails and R returns \mathbb{A} in this case too.

□

Lemma 49 *TL2 preserves reads of committed transactions (part 2).*

$\forall H \in \mathbb{H}(TL2):$

$\forall R \in \text{GlobalTReads}(H):$ Let $i = \text{arg1}_H(R), T = \text{trans}_H(R):$

$T \in \text{Committed}(H) \Rightarrow$

$\text{NoWriterBetween}_{H,i}(T, \sqsubseteq, R)$

T	T'
$R04 \triangleright \quad v = \text{reg}[i].\text{read}()$	
...	
$C07 \triangleright \quad wver = \text{clock.iaf}()$	
...	$C07' \triangleright \quad wver' = \text{clock.iaf}()$
	...
	$C16'_i \triangleright \quad \text{reg}[i].\text{write}(v')$

Figure 8: Case $T \in \text{Committed}(H) \wedge T \sqsubset T' \sqsubset R$

Proof Sketch.

We consider a committed transaction T with an unaborting global read operation R from a location i and a writer T' of i . We should show that it is impossible that T' takes effect after T and T' accesses i before R .

We assume that

T' takes effect after T

that is

(1) $T \sqsubset T'$

We show that

T' accesses i after R .

that is

(2) $R \sqsubset T'$

Figure 8 depicts the two transactions. We annotate the labels and variables of T' by a prime so that they do not conflict with the labels and variables of T .

By Definition 13 on [1], we have

(3) $C07 \prec_{\text{clock}} C07'$

By Definition 13 on [2], we have to show

$R04 \prec_H C16_i$

By P2X and the algorithm, we have

(4) $C04 \prec_H C07$

(5) $C07' \prec_H C16'_i$

By the Lemma XLTRANS on [4], [3], and [5], we have

$R04 \prec_H C16_i$

□

Lemma 50 *TL2 preserves reads of committed transactions.*

$\forall H \in \mathbb{H}(\text{TL2}):$

$\forall R \in \text{GlobalTReads}(H): \text{Let } i = \text{arg1}_H(R), T = \text{trans}_H(R):$

$T \in \text{Committed}(H) \Rightarrow$

$\text{NoWriterBetween}_{H,i}(R, \sqsubseteq, T) \wedge \text{NoWriterBetween}_{H,i}(T, \sqsubseteq, R)$

Proof. Immediate from Lemma 48 and Lemma 49.

□

Lemma 51 *TL2 is read-preserving.*

$\forall H \in \mathbb{H}(\text{TL2}): \text{ReadPres}(H, \sqsubseteq)$

Proof. Immediate from Lemma 47 and Lemma 50.

□

Lemma 52 *Version registers are updated to ascending numbers.*

Let $C17_i^1$ denote the method call at line C17_i executed by a transaction T_1 and let $wver^1$ denote its argument. Similarly, let $C17_i^2$ denote the method call at line C17_i executed by a transaction T_2 and let $wver^2$ denote its argument. If $C17_i^1 \prec_{ver[i]} C17_i^2$, then $wver^1 < wver^2$.

Proof Sketch.

T_1	T_2
...	
$C02_i^1 \triangleright \text{locked}^1 = \text{lock}[i].\text{trylock}()$	
...	
$C07^1 \triangleright wver^1 = \text{clock.iaf}()$	
...	
$C17_i^1 \triangleright \text{ver}[i].\text{write}(wver^1)$	
$C18_i^1 \triangleright \text{lock}[i].\text{unlock}()$	
...	...
	$C02_i^2 \triangleright \text{locked}^2 = \text{lock}[i].\text{trylock}()$
	...
	$C07^2 \triangleright wver^2 = \text{clock.iaf}()$
	...
	$C17_i^2 \triangleright \text{ver}[i].\text{write}(wver^2)$
	$C18_i^2 \triangleright \text{lock}[i].\text{unlock}()$
	...

Figure 9: Updating Version Registers

We have that

$$(1) C17_i^1 \prec_{ver[i]} C17_i^2$$

We show that

$$wver^1 < wver^2$$

By P2X on the algorithm, we have

$$(2) C02_i^1 \prec_H C17_i^1$$

$$(3) C17_i^2 \prec_H C18_i^2$$

By the Lemma XLTRANS on [2], [1] and [3], we have

$$(4) C02_i^1 \prec_H C18_i^2$$

Thus, by the Lemma X2L, we have

$$(5) C02_i^1 \prec_{lock[i]} C18_i^2$$

From the algorithm,

$$(6) \text{ The ownership of } \text{lock}[i] \text{ is respected.}$$

By the Lemma TRYLOCK on [6] and [5], we have

$$(7) C18_i^1 \prec_{lock[i]} C02_i^2$$

By P2X on the algorithm, we have

$$(8) C07^1 \prec_H C18_i^1$$

$$(9) C02_i^2 \prec_H C07^2$$

By the Lemma XLTRANS on [8], [7], and [9], we have

$$(10) C07^1 \prec_H C07^2$$

By the Lemma X2L on [10], we have

$$(11) C07^1 \prec_{clock} C07^2$$

By the Lemma SCOUNTER on [11], we have
 $wver^1 < wver^2$

□

Lemma 53 *For every write method call W on $ver[i]$ with argument v and every read method call R on $ver[i]$ with the return value v' , if $W \prec_{ver[i]} R$ then $v \leq v'$.*

Proof Sketch.

We have

- (1) W is a write method call on $ver[i]$.
- (2) R is a read method call on $ver[i]$.
- (3) $W \prec_{ver[i]} R$.
- (4) The argument of W is v .
- (5) The return value of R is v' .

We show that

$$v \leq v'$$

Let

- (6) W' is last write on $ver[i]$ linearized before R .
- (7) The argument of W' is v'' .

By the Lemma AREG' on [6], [7], and [5], we have

$$(8) v' = v''$$

From [6], and [1], we have

$$(9) W \preceq_{ver[i]} W'$$

By the algorithm and [1], and [6], we have

$$(10) W \text{ and } W' \text{ are both at } C17.$$

By Lemma 52 on [10], [9], [4] and [7], we have

$$(11) v \leq v''$$

From [8] and [11], we have

$$v \leq v'$$

□

Lemma 54 *For every write method call W on $ver[i]$ with argument v and every read method call R on $ver[i]$ with the return value v' , if $R \prec_{ver[i]} W$ then $v' < v$.*

Proof Sketch.

We have

- (1) W is a write method call on $ver[i]$.
- (2) R is a read method call on $ver[i]$.
- (3) $R \prec_{ver[i]} W$.
- (4) The argument of W is v .
- (5) The return value of R is v' .

We show that

$$v' < v$$

Let

- (6) W' is last write on $ver[i]$ linearized before R .

(7) The argument of W' is v'' .

By the Lemma AREG' on [6], [7], and [5], we have

$$(8) \quad v' = v''$$

From [3], and [6], we have

$$(9) \quad W' \prec_{ver[i]} W$$

By the algorithm and [1], and [6], we have

$$(10) \quad W \text{ and } W' \text{ are both at } C17.$$

By Lemma 52 on [10], [9], [4] and [7], we have

$$(11) \quad v'' < v$$

From [8] and [11], we have

$$v' < v$$

□

Lemma 55 *TL2 is global-write-observant.*

$$\begin{aligned}
& \forall H \in \mathbb{H}(TL2): \\
& \forall R \in \text{GlobalTReads}(H): \exists W \in \text{GlobalTWrites}(H): \text{Let } T' = \text{trans}_H(W): \\
& \text{LastPreAccessor}_{H, \sqsubseteq}(T', R) \wedge \\
& \text{arg1}_H(R) = \text{arg1}_H(W) \wedge \text{retv}_H(R) = \text{arg2}_H(W)
\end{aligned}$$

Proof Sketch.

We consider a transaction T with an unaborting global read operation R from a location i . The read operation R is from the location i , thus,

(1) The argument of R is i .

As R is global, thus,

(2) The return value of R is the return value of $R04$.

We first show that

(3) The read method call from $\text{reg}[i]$ at $R04$ is race-free.

We assume that there is a write method call on $\text{reg}[i]$ concurrent to it and show that TL2 aborts R . Figure 10 depicts this situation.

T	T'
	$C02_i \triangleright \text{locked} = \text{lock}[i].\text{trylock}()$
	...
$R03 \triangleright s_1 = \text{ver}[i].\text{read}()$	
$R04 \triangleright v = \text{reg}[i].\text{read}()$	$C16_i \triangleright v = \text{reg}[i].\text{write}(v)$
...	$C17_i \triangleright \text{ver}[i].\text{write}(wver)$
$R05 \triangleright \text{lock}[i].\text{read}()$	$C18_i \triangleright \text{lock}[i].\text{unlock}()$
$R06 \triangleright s_2 = \text{ver}[i].\text{read}()$...
$R07 \triangleright sver = rver[t].\text{read}()$	
if $(\neg(\neg l \wedge s_1 = s_2 \wedge s_2 \leq sver))$ return \mathbb{A}	

Figure 10: $R04$ is race-free

We assume that there a race between $R04$ and $C16_i$. Thus,

(4) $R04 \sim C16_i$

The method calls $R05$ and $C18_i$ are on the object $\text{lock}[i]$.

We consider two cases for the linearization order of them and prove that R returns \mathbb{A} in both cases.

We consider two cases

Case 1:

(5) $R04 \prec_{\text{lock}[i]} C18_i$

By P2X and the algorithm, we have

(6) $C02_i \prec_H C16_i$

(7) $R04 \prec_H R05$

By the Lemma XXTRANS on [6], [4], and [7], we have

(8) $C02_i \prec_H R05$

By the Lemma X2L on [8], we have

(9) $C02_i \prec_{\text{lock}[i]} R05$

By the Lemma TRYLOCKREADM on [9] and [5], we have that

$R05$ returns *true* i.e. $l = \text{true}$

Thus,

The validation check fails and R returns \mathbb{A} .

Case 2:

(10) $C18_i \prec_{lock[i]} R04$

By P2X and the algorithm, we have

(11) $R03 \prec_H R04$

(12) $R05 \prec_H R06$

(13) $C16_i \prec_H C17_i$

(14) $C17_i \prec_H C18_i$

By the Lemma XXTRANS on [11], [4], and [13], we have

(15) $R03 \prec_H C17_i$

By Lemma 54 on [15], we have

(16) $s_1 < wver$

By the Lemma XLTRANS on [14], [10], and [12], we have

(17) $C17_i \prec_H R06$

By Lemma 53 on [17], we have

(18) $s_2 > wver$

From [15] and [17], we have

(19) $s_1 \neq s_2$

Thus,

The validation check fails and R returns \mathbb{A} .

Second, we show that

(20) The register $reg[i]$ is sequentially-written i.e. no two write methods on $reg[i]$ are concurrent.

We assume two concurrent write method calls on $reg[i]$ and show a contradiction.

Figure 11 depicts this situation.

T	T'
$C02_i \triangleright \text{locked} = lock[i].trylock()$	
...	$C02'_i \triangleright \text{locked}' = lock[i].trylock()$
	...
$C16_i \triangleright v = reg[i].write(v)$	$C16'_i \triangleright v' = reg[i].write(v')$
...	...
$C18_i \triangleright lock[i].unlock()$	
	$C18'_i \triangleright lock[i].unlock()$

Figure 11: $reg[i]$ is sequentially-written

We assume that $C16_i$ and $C16'_i$ are concurrent. Thus,

(21) $C16_i \sim C16'_i$

By P2X and the algorithm, we have

(22) $C02_i \prec_H C16_i$

(23) $C16'_i \prec_H C18'_i$

By the Lemma XXTRANS on [22], [21], and [23], we have

(24) $C02_i \prec_H C18'_i$

By the Lemma X2L on [8], we have

(25) $C02_i \prec_{lock[i]} C18'_i$

By the Lemma TRYLOCK on [25], we have that

$$(26) C18_i \prec_{lock[i]} C02'_i$$

By P2X and the algorithm, we have

$$(27) C16_i \prec_H C18_i$$

$$(28) C02'_i \prec_H C16'_i$$

By the Lemma XLTRANS on [27], [26], and [28], we have

$$(29) C16_i \prec_H C16'_i$$

That is a contradiction to [21].

By the Lemma BREG on [3], and [20], we have

$$(30) \text{ There is a write method call } w \text{ on } reg[i] \text{ such that}$$

The argument of w is equal to the return value of $R04$.

The last write method call on $reg[i]$ that is executed before $R04$ is w .

By the algorithm, we have

$$(31) \text{ The register } reg[i] \text{ is written only at } C16_i.$$

From [28] and [29], we have

There is a transaction T' such that

(We annotate the labels and variables of T' by a prime
so that they do not conflict with the labels and variables of T .)

$$(32) \text{ The argument of } C16'_i \text{ is equal to the return value of } R04.$$

$$(33) \text{ The last write method call on } reg[i] \text{ that is executed before } R04 \text{ is } C16'_i.$$

By the algorithm, we have

$$(34) \text{ The argument of } C16'_i \text{ is the value of the key } i \text{ in the map } wset[T'] \text{ in the commit.}$$

$$(35) \text{ The map } wset[T'] \text{ is updated only at } W01 \text{ in a } write \text{ of } T' \text{ such that}$$

The key is equal to the first argument of the *write*.

The value is equal to the second argument of the *write*.

From [34], and [35], we have

$$(36) \text{ There exists a write } W \text{ of } T'$$

$$(37) \text{ The first argument of } W \text{ is equal to } i.$$

$$(38) W \text{ is the last } write \text{ of } T' \text{ with the first argument equal to } i.$$

$$(39) \text{ The second argument of } W \text{ is equal to the argument of } C16'_i.$$

From [1], and [37], we have

$$(40) \text{ The first argument of } R \text{ is the first argument of } W.$$

From [2], [32], and [39], we have

$$(41) \text{ The return value of } R \text{ is the second argument of } W.$$

From [38], we have

$$(42) W \text{ is a global write.}$$

We show that

$$(43) \text{ The transaction } T' \text{ is the last pre-accessor of } R.$$

From [33], we have

$$(44) C16'_i \prec_H R04$$

By Definition 13 on [44], we have

$$(45) T' \sqsubset R$$

Now, we show that

$$(46) \text{ Every transaction } T'' \text{ other than } T' \text{ that accesses } i \text{ before } R, \text{ takes effect before } T'.$$

We assume that

$$(47) T'' \neq T'$$

$$(48) \quad T'' \sqsubset R$$

We should show that

$$T'' \sqsubset T'$$

By Definition 13 on [48], we have

(We annotate the labels and variables of T' by a double prime.)

$$(49) \quad C16''_i \prec_H R04$$

From [33], [33], and [49], we have

$$(50) \quad C16''_i \prec_H C16'_i$$

Consider Figure 12.

T''	T'
$C02''_i \triangleright \text{locked}'' = \text{lock}[i].\text{tryLock}()$	
...	$C02'_i \triangleright \text{locked}' = \text{lock}[i].\text{tryLock}()$
$C07''_i \triangleright \text{wver}'' = \text{clock}.\text{iaf}()$...
...	$C07'_i \triangleright \text{wver}' = \text{clock}.\text{iaf}()$
$C16''_i \triangleright \text{reg}[i].\text{write}(v'')$...
...	$C16'_i \triangleright \text{reg}[i].\text{write}(v')$
$C18''_i \triangleright \text{lock}[i].\text{unlock}()$...
	$C18'_i \triangleright \text{lock}[i].\text{unlock}()$

Figure 12: Effect-order of pre-accessors

By P2X and the algorithm, we have

$$(51) \quad C02''_i \prec_H C16''_i$$

$$(52) \quad C16''_i \prec_H C18'_i$$

By the Lemma XXTRANS on [51], [50], and [52], we have

$$(53) \quad C02''_i \prec_H C18'_i$$

By the Lemma X2L on [53], we have

$$(54) \quad C02''_i \prec_{\text{lock}[i]} C18'_i$$

By the Lemma TRYLOCK on [45], we have that

$$(55) \quad C18''_i \prec_{\text{lock}[i]} C02'_i$$

By P2X and the algorithm, we have

$$(56) \quad C07''_i \prec_H C18''_i$$

$$(57) \quad C02'_i \prec_H C07'_i$$

By the Lemma XLTRANS on [56], [55], and [57], we have

$$(58) \quad C07''_i \prec_H C07'_i$$

By Definition 13 on [58], we have

$$T'' \sqsubset T'.$$

The conclusion is

$$[36], [42], [40], [41], \text{ and } [43]$$

□

Lemma 56 *TL2 is local-write-observant.*

$$\forall H \in \mathbb{H}(TL2):$$

$$\forall R \in \text{LocalTReads}(H): \text{Let } T = \text{trans}_H(R), i = \text{arg1}_H(R), H' = H|T|i:$$

$$\exists W \in \text{TWrites}(H'):$$

$$W \prec_{H'} R \wedge \text{NoWriteBetween}_{H'}(W, R) \wedge$$

$$\text{retv}_{H'}(R) = \text{arg2}_{H'}(W)$$

Proof Sketch.

Let

(1) The operation R is a local read with the first argument i by the transaction T .

From [1], as R is local, we have

(2) There is a write operation before R with the first argument i by T .

From [2], let

(3) The operation W is the last write operation before R with the first argument i by the transaction T .

By the algorithm

(4) The value of a key i in $wset$ is updated only at $W01$ in a write operation with the first argument i and the value of the key i is updated to the second argument of the write operation.

From [3] and [4], we have

(5) The value of a key i in $wset$ during the execution of R is equal to the second argument of W .

Thus, by the algorithm

(6) $R01$ - $R02$ find a value for the key i in $wset$.

Thus,

(7) The return value of R is equal to the value of key i in $wset$.

From [7] and [5], we have

(8) The return value of R is equal to the second argument of W .

The conclusion is

[3] and [8]

□

Lemma 57 *TL2 is write-observant.*

$$\forall H \in \mathbb{H}(TL2): WriteObs(H, \sqsubseteq)$$

Proof. Immediate from Lemma 56 and Lemma 55.

□

Lemma 58 *TL2 is real-time-preserving.*

$$\forall H \in \mathbb{H}(TL2): RealTimePres(H, \sqsubseteq)$$

Proof Sketch.

We assume that

$$(1) T \preceq_H T'$$

We show that

$$T \sqsubseteq T'$$

By the definition of \preceq_H , from [1], we have

$$(2) \text{ All the operations of } T \text{ are executed before all the operations of } T'.$$

By the Lemma X2L, from [2], we have

$$(3) \text{ All the operations of } T \text{ on } clock \text{ are linearized before all the operations of } T' \text{ on } clock.$$

By Definition 13,

$$(4) \text{ The effect point of each transaction is one of its own operations on the } clock \text{ object.}$$

From [3] and [4], we have

$$(5) \text{ The transaction } T \text{ takes effect before the transaction } T'.$$

that is

$$T \sqsubseteq T'$$

□

Lemma 59 *The relation \sqsubseteq is a marking relation.*

$$\forall H \in \mathbb{H}(TL2): \sqsubseteq \in Marking(H)$$

Proof Sketch.

Consider Definition 13.

By the totality of the linearization order \prec_{clock} , the relation \sqsubseteq is a total on the set of transactions.

As every pair of method calls either execute in order or concurrently, every read operation of a location i is ordered either before or after every writer to i . In addition, as no method call can execute before another method call and also after after or concurrent to it, no read operation of a location i is ordered both before and after a writer to i .

□

Lemma 60 *TL2 is markable.*

$$\forall H \in \mathbb{H}(TL2): H \in FinalStateMarkable$$

Proof.

Immediate from Lemma 59, Lemma 51, Lemma 57, and Lemma 58.

□

Theorem 61 *TL2 is opaque.*

$$\forall H \in \mathbb{H}(TL2): H \in FinalStateOpaque$$

Proof.

Immediate from Lemma 60, and Theorem 18.

□

7 Marking DSTM (visible reads)

\mathcal{T} : <i>Loc</i> { <i>writer</i> : BasicRegister , <i>rset</i> : BasicSet , <i>oldVal</i> : BasicRegister , <i>newVal</i> : BasicRegister }, <i>state</i> : AtomicCASRegister [], <i>start</i> : AtomicCASRegister []	
\mathcal{D} :	
def <i>init</i> _{<i>t</i>} () <i>I01</i> ▷ <i>state</i> [<i>t</i>]. <i>write</i> (\mathbb{R}), <i>I02</i> ▷ return <i>ok</i> ,	def <i>write</i> _{<i>t</i>} (<i>i</i> , <i>v</i>) <i>W01</i> ▷ <i>r</i> ₁ = <i>start</i> [<i>i</i>]. <i>read</i> (), <i>W02</i> ▷ <i>w</i> = <i>r</i> ₁ . <i>writer</i> . <i>read</i> (), if (<i>w</i> = <i>t</i>) <i>W03</i> ▷ <i>r</i> ₁ . <i>newVal</i> . <i>write</i> (<i>v</i>), <i>W04</i> ▷ return <i>ok</i> ,
def <i>read</i> _{<i>t</i>} (<i>i</i>) <i>R01</i> ▷ <i>r</i> ₁ = <i>start</i> [<i>i</i>]. <i>read</i> (), <i>R02</i> ▷ <i>v</i> = <i>currentValue</i> _{<i>t</i>} (<i>r</i> ₁), <i>R03</i> ▷ <i>r</i> ₂ = <i>clone</i> (<i>r</i> ₁), <i>R04</i> ▷ <i>r</i> ₂ . <i>rset</i> . <i>add</i> (<i>t</i>), <i>R05</i> ▷ <i>rd</i> = <i>start</i> [<i>i</i>]. <i>cas</i> (<i>r</i> ₁ , <i>r</i> ₂), <i>R06</i> ▷ <i>s</i> = <i>state</i> [<i>t</i>]. <i>read</i> (), if ($\neg rd \vee (s = \mathbb{A})$) <i>R07</i> ▷ return \mathbb{A} else <i>R08</i> ▷ return <i>v</i> , {i05 → i06},	<i>W05</i> ▷ <i>v</i> ₂ = <i>currentValue</i> _{<i>t</i>} (<i>r</i> ₁), <i>W06</i> ▷ foreach (<i>t</i> ₂ ∈ <i>r</i> ₁ . <i>rset</i>) <i>W07</i> ▷ <i>state</i> [<i>t</i> ₂]. <i>cas</i> (\mathbb{R} , \mathbb{A}), <i>W08</i> ▷ <i>r</i> ₂ = <i>new Loc</i> (), <i>W09</i> ▷ <i>r</i> ₂ . <i>writer</i> . <i>write</i> (<i>t</i>), <i>W10</i> ▷ <i>r</i> ₂ . <i>oldVal</i> . <i>write</i> (<i>v</i> ₂), <i>W11</i> ▷ <i>r</i> ₂ . <i>newVal</i> . <i>write</i> (<i>v</i>), <i>W12</i> ▷ <i>wd</i> = <i>start</i> [<i>i</i>]. <i>cas</i> (<i>r</i> ₁ , <i>r</i> ₂), if (<i>wd</i>) <i>W13</i> ▷ return <i>ok</i> else <i>W14</i> ▷ return \mathbb{A} {i06 → i12}
def <i>commit</i> _{<i>t</i>} () <i>C01</i> ▷ <i>c</i> = <i>state</i> [<i>t</i>]. <i>cas</i> (\mathbb{R} , \mathbb{C}), if (<i>c</i>) <i>C02</i> ▷ return \mathbb{C} else <i>C03</i> ▷ return \mathbb{A} ,	
def <i>currentValue</i> _{<i>t</i>} (<i>r</i>) <i>V01</i> ▷ <i>t</i> ₂ = <i>r</i> . <i>writer</i> . <i>read</i> (), if ($\neg(t_2 = t)$) <i>V02</i> ▷ <i>state</i> [<i>t</i> ₂]. <i>cas</i> (\mathbb{R} , \mathbb{A}), <i>V03</i> ▷ <i>s</i> = <i>state</i> [<i>t</i> ₂]. <i>read</i> (), if (<i>s</i> = \mathbb{A}) <i>V04</i> ▷ return <i>r</i> . <i>oldVal</i> else <i>V05</i> ▷ return <i>r</i> . <i>newVal</i> ,	

Figure 13: *DSTMVis* DSTM (visible reads) Algorithm Specification

Notation. Let us remind the notation. Consider an execution history H .

We write $e_1 \triangleleft_H e_2$ to denote that the event e_1 comes before the event e_2 in the history H .

We use $l_1 \prec_H l_2$ to denote that l_1 is executed before l_2 . We use $l_1 \sim_H l_2$ to denote that l_1 is executed concurrently to l_2 . We use $l_1 \lesssim_H l_2$ to denote that l_1 is executed before or concurrently to l_2 .

We use $\prec_{start[i]}$ to denote the linearization order of $start[i]$.

A label $c_1'c_2$ is a call string that denotes a method call labeled c_2 that is executed in the body of the method call labeled c_1 .

We use $initOf_H(T)$ and $commitOf_H(T)$ to denote the *init* and *commit* method calls of the transaction T in the history H . We use $LastTRead_H(T)$ to denote the last read method call by the transaction T in the history H . We use $FirstTWrite_H(T, i)$ to denote the first write method call to location i by the transaction T in the history H .

Marking Relation. Now, we define the marking relation for DSTM.

Definition 14 (Marking DSTM) Consider an execution history $H \in \mathbb{H}(DSTMVis)$. Let

$$\begin{aligned} Eff(T) &= \begin{cases} commitOf_H(T)'C01 & \text{if } T \in Committed(H) \\ LastTRead_H(T)'R05 & \text{if } T \in Aborted(H) \wedge TReads(H) \neq \emptyset \\ initOf_H(T)'I01 & \text{if } T \in Aborted(H) \wedge TReads(H) = \emptyset \end{cases} \\ readAcc(R) &= R'R05 \\ writeAcc(T, i) &= FirstTWrite_H(T, i)'W12 \end{aligned}$$

The marking \sqsubseteq for H is the reflexive closure of \sqsubset that is define as follows:

$$\begin{aligned} &\{(T, T') \mid T, T' \in Trans(H) \wedge inv(Eff(T)) \triangleleft_H inv(Eff(T'))\} \cup \\ &\{(T, R) \mid \exists i: R \in GlobalTReads(H), i = arg1(R), T \in Writers_H(i) \wedge writeAcc(T, i) \prec_{start[i]} readAcc(R)\} \cup \\ &\{(R, T) \mid \exists i: R \in GlobalTReads(H), i = arg1(R), T \in Writers_H(i) \wedge readAcc(R) \prec_{start[i]} writeAcc(T, i)\} \end{aligned}$$

A committed transactions takes effect at the invocation event of $C01$, the *cas* method call in its commit method call. An aborted transaction that has a successful read method call takes effect at the invocation event of $R05$ of its last successful read method call. An aborted transaction that has no successful read method call takes effect at the invocation event of $I01$ in its initialization method call.

The access point of a read method call is at $R05$. The access point of a writer transaction to location i is at $W12$ of its first write method call to i .

8 Marking NORec

\mathcal{T} : <i>seqLock</i> : SeqLock, <i>reg</i> : BasicRegister [] <i>snap</i> : ThreadLocal BasicRegister, <i>rset</i> : ThreadLocal BasicMap, <i>wset</i> : ThreadLocal BasicMap,	
\mathcal{D} :	
def <i>init_t</i> () do I01 \triangleright (<i>s</i> , <i>l</i>) = <i>seqLock.read</i> () while (<i>l</i>), I02 \triangleright <i>snap</i> [<i>t</i>] = <i>s</i> ,	def <i>validate_t</i> () V01 \triangleright while (<i>true</i>) do V02 \triangleright (<i>s1</i> , <i>l1</i>) = <i>seqLock.read</i> (), while (<i>l1</i>) foreach ((<i>i</i> , <i>v</i>) \in <i>rset</i> [<i>t</i>]) V03 _{<i>i</i>} \triangleright <i>v'</i> = <i>reg</i> [<i>i</i>]. <i>read</i> (), if (<i>v</i> \neq <i>v'</i>), V04 _{<i>i</i>} \triangleright return false , V05 \triangleright (<i>s2</i> , <i>l2</i>) = <i>seqLock.read</i> (), if (<i>s2</i> = <i>s1</i> \wedge \neg <i>l2</i>) V06 \triangleright <i>snap</i> [<i>t</i>]. <i>write</i> (<i>s1</i>), V07 \triangleright return true , {V02 \rightarrow V03 _{<i>i</i>} , V03 _{<i>i</i>} \rightarrow V05},
def <i>read_t</i> (<i>i</i>) R01 \triangleright <i>pv</i> = <i>wset</i> [<i>t</i>]. <i>get</i> (<i>i</i>), if (<i>pv</i> \neq \perp) R02 \triangleright return pv , do R03 \triangleright <i>v</i> = <i>reg</i> [<i>i</i>]. <i>read</i> (), R04 \triangleright <i>s1</i> = <i>snap</i> [<i>t</i>]. <i>read</i> (), R05 \triangleright (<i>s2</i> , <i>l2</i>) = <i>seqLock.read</i> (), if (<i>s2</i> = <i>s1</i> \wedge \neg <i>l2</i>) R06 \triangleright break , R07 \triangleright <i>b</i> = <i>validate_t</i> (<i>i</i>), if (\neg <i>b</i>) R08 \triangleright return \mathbb{A} , while (<i>true</i>), R09 \triangleright <i>rset</i> [<i>t</i>]. <i>put</i> (<i>i</i> , <i>v</i>), R10 \triangleright return v , {R03 \rightarrow R05},	def <i>commit_t</i> (<i>i</i>) C01 \triangleright <i>e</i> = <i>wset</i> [<i>t</i>]. <i>isEmpty</i> (), if (<i>e</i>) C02 \triangleright return \mathbb{C} , do C03 \triangleright <i>s</i> = <i>snap</i> [<i>t</i>]. <i>read</i> (), C04 \triangleright <i>d</i> = <i>seqLock.compareAndLock</i> (<i>s</i>), if (<i>d</i>) C05 \triangleright break , C06 \triangleright <i>b</i> = <i>validate_t</i> (<i>i</i>), if (\neg <i>b</i>) return \mathbb{A} , while (<i>true</i>), foreach ((<i>i</i> , <i>v</i>) \in <i>wset</i> [<i>t</i>]) C07 _{<i>i</i>} \triangleright <i>reg</i> [<i>i</i>]. <i>write</i> (<i>v</i>), C08 \triangleright <i>seqLock.incAndUnlock</i> (), C09 \triangleright return \mathbb{C} {C04 \rightarrow C07 _{<i>i</i>} , C07 _{<i>i</i>} \rightarrow C08},
def <i>write_t</i> (<i>i</i> , <i>v</i>) W01 \triangleright <i>wset</i> [<i>t</i>]. <i>put</i> (<i>i</i> , <i>v</i>), W02 \triangleright return ok ,	
def <i>abort_t</i> (<i>i</i>) A01 \triangleright return \mathbb{A}	

Figure 14: NORec NORec Algorithm Specification

Notation. Let us remind the notation. Consider an execution history H .

We use $l_1 \prec_H l_2$ to denote that l_1 is executed before l_2 . We use $l_1 \sim_H l_2$ to denote that l_1 is executed concurrently to l_2 . We use $l_1 \lesssim_H l_2$ to denote that l_1 is executed before or concurrently to l_2 .

We use $\prec_{seqLock}$ to denote the linearization order of *seqLock*.

A label $c_1'c_2$ is a call string that denotes a method call labeled c_2 that is executed in the body of the method call labeled c_1 .

We use $initOf_H(T)$ and $commitOf_H(T)$ to denote the *init* and *commit* method calls of the transaction T in the history H .

Marking Relation. Now, we define the marking relation for NoRec.

Definition 15 (Marking NoRec) Consider an execution history $H \in \mathbb{H}(NORec)$. Let

$$\begin{aligned}
REff(T) &= \text{The last execution of } I01 \text{ or } V05 \\
Eff(T) &= \begin{cases} REff(T) & \text{if } T \in Aborted(H) \vee TWrites(H) = \emptyset \\ commitOf(T)'C04 & \text{if } T \in Committed(H) \wedge TWrites(H) \neq \emptyset \end{cases} \\
readAcc(T, i) &= \begin{cases} R'R03 & \text{if } REff(T) \prec_H R'R03 \\ \text{Let } REff(T) = V'V05 \text{ in } V'V03_i & \text{if } R'R03 \prec_H REff(T) \end{cases} \\
writeAcc(T, i) &= commitOf(T)'C07_i
\end{aligned}$$

The marking \sqsubseteq for H is the reflexive closure of \sqsubset that is define as follows:

$$\begin{aligned}
&\{(T, T') \mid T, T' \in Trans(H) \wedge Eff(T) \prec_{seqLock} Eff(T')\} \cup \\
&\{(T, R) \mid \exists i: R \in GlobalTReads(H), i = arg1(R), T \in Writers_H(i) \wedge writeAcc(T, i) \prec_H readAcc(T, i)\} \cup \\
&\{(R, T) \mid \exists i: R \in GlobalTReads(H), i = arg1(R), T \in Writers_H(i) \wedge readAcc(T, i) \prec_H writeAcc(T, i)\}
\end{aligned}$$

An aborted transaction or a read-only transaction takes effect at the last execution of *I01* or *V05*. This method call reads that most recent snapshot value that the transaction is still consistent for. A committed transactions that has write method calls takes effect at *C04*.

The access point of a read method call is at *R03* if the last recent snapshot is read before *R03*; otherwise, it is at *V03_i* of the latest successful validate method call. The access point of a writer transaction to location i is at *C07_i*.

9 The Cost of Read Validation

The read-preservation invariant requires the TM algorithm to check that a read location is not overwritten between the point where the location is read and the point where the transaction takes effect. This requirement motivated us to study how read-preservation can influence the time complexity of TM operations and helped us construct client scenarios that exhibit lower bounds. We present a generalization of the seminal lower bound result presented in [2]. Let us first remind some definitions from previous works on the inherent complexity of TM [1, 2, 4, 5].

An aborted transaction that did not invoke an abort operation is said to be *forcefully* aborted. We say that two transactions *conflict* if they access the same location and one of them writes to the location. A TM algorithm is (weakly) *progressive* if and only if it forcefully aborts a transaction only when it conflicts with a live transaction. More precisely, it aborts a transaction only when there is a time t at which it conflicts with another concurrent transaction that is live at time t (not committed or aborted by time t). In addition to providing progress, progressive TM algorithms are expected to retry transactions less frequently and therefore, improve performance.

A TM algorithm is *invisible-reads* if and only if no read operation mutates any base object. Mutating base objects can potentially invalidate the caches and adversely affect performance. Thus, most high-performance TM algorithms are invisible-reads. A transaction is *read-only* if and only if it does not invoke any write operations. We assume that the abort operation for a read-only transaction does not mutate any base shared object.

Two transactions *contend* on a base object o if and only if they access o and at least one of them mutates o . A TM algorithm is (strictly) *disjoint-access-parallel* if and only if two transactions contend on a base object only if they access a common memory location. Disjoint-access-parallelism can improve scalability as transactions that access disjoint memory locations access disjoint base objects.

A TM algorithm is *single-version* if and only if it stores a single value for each memory location in the base objects.

Theorem 62 *The time complexity of the commit operation of every opaque, progressive, disjoint-access-parallel and invisible-reads TM algorithm is $\Omega(|\mathcal{R}|)$ where \mathcal{R} is the read set.*

We explain the key idea here and then present the proof. Consider a TM algorithm TM that is opaque, progressive, disjoint-access-parallel and invisible-reads. Consider the following client scenario. Invoke the following methods in sequence. Wait for the response of the method call of each step before going to the next step. (1) $init_{T_1}()$, (2) $read_{T_1}(i)$ (3) $init_{T_2}()$, (4) $write_{T_2}(i, v_1)$, (5) $commit_{T_2}()$, (6) $init_{T_3}()$, (7) $read_{T_3}(j)$, (8) $abort_{T_3}()$, (9) $write_{T_1}(j, v_1)$, (10) $commit_{T_1}()$. As the TM is opaque, progressive and invisible-reads, it can be shown that it results in the history H_1 depicted in Figure 15(a). The initializing transaction T_0 (that initializes every location to v_0) and also the initializing operations of transactions are elided for brevity.

To make sure that the read location i is not overwritten, the commit operation of T_1 should access a shared object that T_2 (that is a writer of i) mutates. Assume otherwise i.e. the commit operation of T_1 does not access any shared object that T_2 mutates. Thus, T_2 is invisible to T_1 . As TM is invisible-reads, it can be shown that T_3 is invisible to other transactions. As T_2 and T_3 are invisible to T_1 , removing them from the client scenario does not affect the responses that T_1 receives. Therefore, the execution of T_1 alone results in the execution history H_2 depicted in Figure 15(b). As there is no conflicting transaction and TM is progressive, TM cannot forcefully abort the commit operation of T_1 . The commit operation should have returned \mathbb{C} but has returned \mathbb{A} that is a contradiction. Therefore, we conclude that the commit operation of T_1 accesses a shared object that T_2 mutates. The scenario can be trivially extended to an arbitrary location k in the read set \mathcal{R} by generalizing the transaction T_2 with the transaction $T_2^k = write_{T_2^k}(k, v_1) \cdot commit_{T_2^k}()$. It can be shown that for every $k \in \mathcal{R}$, the commit operation of T_1 accesses a shared object that the transaction

T_2^k mutates. The transactions $\{T_2^k \mid k \in \mathcal{R}\}$ access disjoint locations. As TM is strictly disjoint-access-parallel, these transactions access disjoint shared objects. Thus, the commit operation of T_1 accesses a separate shared object for every $k \in \mathcal{R}$. Therefore, the commit operation of T_1 accesses at least $|\mathcal{R}|$ shared objects. Therefore, the time complexity of the commit operation of T_1 is $\Omega(|\mathcal{R}|)$.

This theorem shows that designers should pick at least one of the following sources of inefficiency in the design of every opaque TM algorithm: aborting non-conflicting transactions, sharing base objects between transactions that access disjoint locations, visible reads or linear-time complexity of the commit method. As an example, TL2 shares the *clock* object between all transactions and is, therefore, not disjoint-access-parallel. In addition, it has linear-time read-validation in the commit method.

Proof.

Consider a TM algorithm TM that is opaque, progressive, disjoint-access-parallel and invisible-reads. We describe the following client scenario with three transactions T_1 , T_2 and T_3 and consider its execution with TM .

1. Invoke $init_{T_1}()$ and wait for the response.
2. Invoke $read_{T_1}(i)$ and wait for the response.
3. Invoke $init_{T_2}()$ and wait for the response.
4. Invoke $write_{T_2}(i, v_1)$ and wait for the response.
5. Invoke $commit_{T_2}()$ and wait for the response.
6. Invoke $init_{T_3}()$ and wait for the response.
7. Invoke $read_{T_3}(j)$ and wait for the response.
8. Invoke $abort_{T_3}()$ and wait for the response.
9. Invoke $write_{T_1}(j, v_1)$ and wait for the response.
10. Invoke $commit_{T_1}()$ and wait for the response.

The resulting history H_1 for this scenario is depicted in Figure 15(a). The initializing transaction T_0 (that initializes every location to v_0) and also the initializing operations of each transaction are elided for brevity.

The transaction T_1 first invokes the init and then a read operation on the location i . As TM is progressive and T_1 is not in conflict with any other transaction, TM does not forcefully abort the read operation. Therefore, it returns a value. As TM is opaque, there should be a justifying history S for the current execution history (after the read operation returns). As the initializing transaction T_0 is executed before T_1 , the real-time-order property requires T_0 to be ordered before T_1 in S . The transaction T_0 writes the initial value v_0 to every location and commits. Thus, the read operation returns v_0 .

Then, the transaction T_2 invokes the init and then a write operation to i with the value v_1 and then invokes the commit operation. As TM is invisible-reads, the read operation of T_1 is invisible to T_2 . Thus, T_2 does not observe any inconsistency and as TM is progressive, both the write and commit operations are successful.

Next, the transaction T_3 invokes the init and then a read operation on the location j and then invokes the abort operation. As there are no conflicting operations on j and TM is progressive, the read operation is not forcefully aborted. Therefore, it returns a value. As TM is opaque, there should be a justifying history S' for the current execution history (after the read operation returns). As the initializing transaction T_0 is executed before T_3 , the real-time-order property requires T_0 to be ordered before T_3 in S' . The transaction

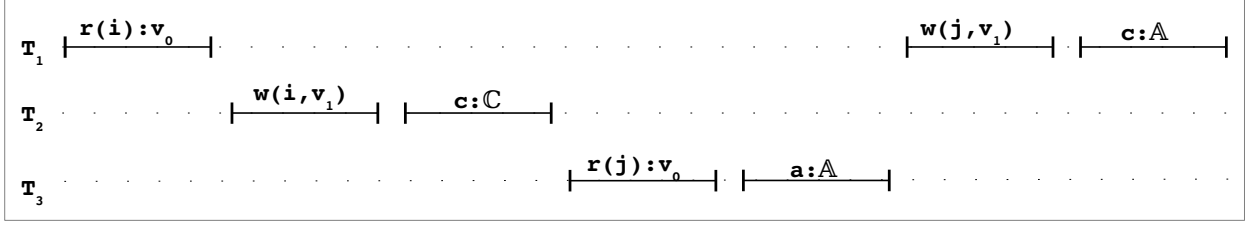
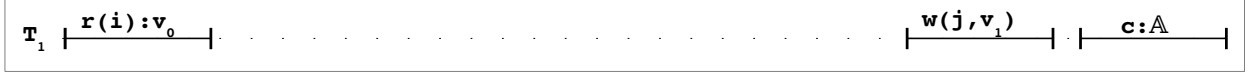
(a) H_1 (b) H_2

Figure 15: The execution histories constructed by the scenarios in the proof of Theorem 62. The letters r , w , c , and a abbreviate read, write, commit and abort operations. The initializing transaction T_0 (that initializes every location to v_0) and also the initializing operations of each transaction are elided.

T_0 is the only committed transaction that has written to j . Thus, the read operation returns the initial value v_0 .

Next, the transaction T_1 invokes a write operation on location j with the value v_1 . When this write operation is invoked, neither T_2 nor T_3 are alive and TM is progressive. Therefore, TM does not forcefully abort the write operation. Finally, T_1 invokes the commit operation. We show that the commit operation aborts i.e. returns \mathbb{A} . Let us assume otherwise, i.e. T_1 commits. As TM is opaque, there is a justifying history S'' for H_1 i.e. S'' is a sequential history that is equivalent to H_1 , is real-time-preserving and is a member of transactional sequential specification. As T_2 is executed before T_3 in H_1 , the real-time-preservation property requires T_2 to be before T_3 in S'' . Thus, there are three possible transaction orderings for S'' . We show that none of them is a justifying history.

- $S'' = H_1|T_0 \cdot H_1|T_1 \cdot H_1|T_2 \cdot H_1|T_3$

We have that $Visible(S'', T_3)|j = write_{T_0}(j, v_0) \cdot write_{T_1}(j, v_1) \cdot read_{T_3}(j):v_0 \notin SeqSpec(j)$. The read operation is expected to return the value v_1 but has returned v_0 . Thus, S'' is not a justifying history.

- $S'' = H_1|T_0 \cdot H_1|T_2 \cdot H_1|T_1 \cdot H_1|T_3$

Similar to the previous case, $Visible(S'', T_3)|j = write_{T_0}(j, v_0) \cdot write_{T_1}(i, v_1) \cdot read_{T_3}(j):v_0 \notin SeqSpec(j)$. The read operation is expected to return the value v_1 but has returned v_0 . Thus, S'' is not a justifying history.

- $S'' = H_1|T_0 \cdot H_1|T_2 \cdot H_1|T_3 \cdot H_1|T_1$

We have that $Visible(S'', T_1)|i = write_{T_0}(i, v_0) \cdot write_{T_2}(i, v_1) \cdot read_{T_1}(i):v_0 \notin SeqSpec(i)$. The read operation is expected to return the value v_1 but has returned v_0 . Thus, S'' is not a justifying history.

Thus, we arrive at a contradiction. Therefore, we conclude that the commit operation of T_1 returns \mathbb{A} .

Now, we argue that the commit operation of T_1 should access a shared object that T_2 mutates. Assume otherwise i.e. the commit operation of T_1 does not access any shared object that T_2 mutates. Thus, T_2 is invisible to T_1 . As TM is invisible-reads, the read operation of T_3 does not mutate any shared objects. Furthermore, T_3 is a read-only transaction. Thus, its abort operation does not mutate any shared objects. Therefore, T_3 is invisible to other transactions. As T_2 and T_3 are invisible to T_1 , removing them from the client scenario does not affect the responses that T_1 receives. Therefore, the execution of T_1 alone results

in the execution history H_2 depicted in Figure 15(b). As there is no conflicting transaction and TM is progressive, TM cannot forcefully abort the commit operation of T_1 . The commit operation should have returned \mathbb{C} but has returned \mathbb{A} that is a contradiction. Therefore, we conclude that the commit operation of T_1 accesses a shared object that T_2 mutates.

In the above client scenario, the read set of T_1 was the singleton set i . The scenario can be trivially extended to an arbitrary location k in a read set \mathcal{R} .

1. Invoke $init_{T_1}()$ and wait for the response.
2. For each $i \in \mathcal{R}$:
 - 2.1. Invoke $read_{T_1}(i)$ and wait for the response.
3. Invoke $init_{T_2^k}()$ and wait for the response.
4. Invoke $write_{T_2^k}(k, v_1)$ and wait for the response.
5. Invoke $commit_{T_2^k}()$ and wait for the response.
6. Invoke $init_{T_3}()$ and wait for the response.
7. Invoke $read_{T_3}(j)$ and wait for the response.
8. Invoke $abort_{T_3}()$ and wait for the response.
9. Invoke $write_{T_1}(j, v_1)$ and wait for the response.
10. Invoke $commit_{T_1}()$ and wait for the response.

A similar reasoning concludes that for every $k \in \mathcal{R}$, the commit operation of T_1 accesses a shared object that T_2^k mutates.

The transactions T_2^k (for $k \in \mathcal{R}$) access disjoint locations. As TM is strictly disjoint-access-parallel, the transactions T_2^k (for $k \in \mathcal{R}$) access disjoint shared objects. Thus, the commit operation of T_1 accesses a separate shared object for every $k \in \mathcal{R}$. Therefore, the commit operation of T_1 accesses at least $|\mathcal{R}|$ shared objects. Therefore, the time complexity of the commit operation of T_1 is $\Omega(|\mathcal{R}|)$. □

Theorem 63 *The time complexity of the commit operation of every opaque, progressive, and invisible-reads TM algorithm that stores information about a constant number of locations in each shared object is $\Omega(|\mathcal{R}|)$ where \mathcal{R} is the read set.*

The proof of this theorem uses the same client scenario as the proof of Theorem 62. The main difference is the final step of reasoning. As information about a constant number c of locations can be obtained from each shared object, the commit operation of T_1 has to read at least $|\mathcal{R}|/c$ shared objects.

Proof.

The proof of this theorem flows similar to the proof of Theorem 62 to the point that we have that

- (1) For every $k \in \mathcal{R}$, the commit operation of T_1 reads a shared object that T_2^k mutates.

We have that

- (2) Each transaction T_2^k accesses a separate location k .

From the premises we have that

(3) Information about a constant number c of locations is stored in each shared object.

From 2 and 3, we have that

(4) At most c of the set of writer transactions $T_2^k, k \in \mathcal{R}$ can write to the same shared object.

From 1 and 4, we have

The commit operation of T_1 has to read at least $|\mathcal{R}|/c$ shared objects.

Therefore,

The time complexity of the commit operation of T_1 is $\Omega(|\mathcal{R}|)$. □

We restate Theorem 3 of [2] below. Our Theorem 63 generalizes this theorem by dropping the single-version requirement. Note that the assumption about limited capacity of shared objects is stated before the theorem in [2] and explicitly in the theorem here.

Theorem 64 (*Theorem 3 of [2]*) *The time complexity of every opaque, progressive, single-version and invisible-reads TM algorithm that stores information about a constant number of locations in each shared object is $\Omega(|I|)$ (where I is the set of locations).*

References

- [1] Hagit Attiya, Eshcar Hillel, and Alessia Milani. Inherent limitations on disjoint-access parallel implementations of transactional memory. *Theory of Computing Systems*, 49(4), 2011.
- [2] R. Guerraoui and M. Kapalka. On the correctness of transactional memory. In *PPOPP*, 2008.
- [3] Rachid Guerraoui, Thomas A. Henzinger, and Vasu Singh. Model checking transactional memories. *Distributed Computing*, 2010.
- [4] Rachid Guerraoui and Michal Kapalka. On obstruction-free transactions. In *SPAA*, 2008.
- [5] Dmitri Perelman, Rui Fan, and Idit Keidar. On maintaining multiple versions in stm. In *PODC*, 2010.