

On the Semantics of Snapshot Isolation

Azalea Raad¹, Ori Lahav², and Viktor Vafeiadis¹

¹ MPI-SWS, Germany

² Tel Aviv University, Israel

{azalea,viktor}@mpi-sws.org

orilahav@tau.ac.il

Abstract. Snapshot isolation (SI) is a standard transactional consistency model used in databases, distributed systems and software transactional memory (STM). Its semantics is formally defined both declaratively as an acyclicity axiom, and operationally as a concurrent algorithm with memory bearing timestamps.

We develop two simpler equivalent operational definitions of SI as lock-based reference implementations that do not use timestamps. Our first locking implementation is prescient in that requires *a priori* knowledge of the data accessed by a transaction and carries out transactional writes eagerly (in-place). Our second implementation is non-prescient and performs transactional writes lazily by recording them in a local log and propagating them to memory at commit time. Whilst our first implementation is simpler and may be better suited for developing a program logic for SI transactions, our second implementation is more practical due to its non-prescience. We show that both implementations are sound and complete against the declarative SI specification and thus yield equivalent operational definitions for SI.

We further consider, for the first time *formally*, the use of SI in a context with racy non-transactional accesses, as can arise in STM implementations of SI. We introduce *robust snapshot isolation* (RSI), an adaptation of SI with similar semantics and guarantees in this mixed setting. We present a declarative specification of RSI as an acyclicity axiom and analogously develop two operational models as lock-based reference implementations (one eager, one lazy). We show that these operational models are both sound and complete against the declarative RSI model.

1 Introduction

Transactions are the *de facto* synchronisation mechanism in databases and geo-replicated distributed systems, and are thus gaining adoption in the shared-memory setting via *software transactional memory* (STM) [33,20]. In contrast to other synchronisation mechanisms, transactions readily provide atomicity, isolation, and consistency guarantees for sequences of operations, allowing programmers to focus on the high-level design of their systems.

However, providing these guarantees comes at a significant cost. As such, various transactional consistency models in the literature trade off consistency

guarantees for better performance. At nearly the one end of the spectrum, we have *serialisability* [28], which requires transactions to appear to have been executed in some total order. Serialisability provides strong guarantees, but is widely considered too expensive to implement. The main problem is that two conflicting transactions (e.g. one reading from and one updating the same datum) cannot both execute and commit in parallel.

Consequently, most major databases, both centralised (e.g. Oracle and MS SQL Server) and distributed [15,32,29], have opted for a slightly weaker model called *snapshot isolation* (SI) [7] as their default consistency model. SI has much better performance than serialisability by allowing conflicting transactions to execute concurrently and commit successfully as long as they do not have a write-write conflict. This in effect allows reads of SI transactions to read from an earlier memory snapshot than the one affected by their writes, and permits the *write skew anomaly* [11] depicted in Fig. 1. Besides this anomaly, however, SI is essentially the same as serialisability: Cerone et al. [11] provide a widely applicable condition under which SI and serialisability coincide for a given set of transactions. For these reasons, SI has also started gaining adoption in the generic programming language setting via STM implementations [1,16,8,26,25] that provide SI semantics for their transactions.

The formal study of SI, however, has so far not accounted for the more general STM setting in which both transactions and uninstrumented non-transactional code can access the same memory locations. While there exist two equivalent definitions of SI—one declarative in terms of an acyclicity constraint [10,11] and one operational in terms of an optimistic multi-version concurrency control algorithm [7]—neither definition supports *mixed-mode* (i.e. both transactional and non-transactional) accesses to the same locations. Extending the definitions to do so is difficult for two reasons: (1) the operational definition attaches a timestamp to every memory location, which heavily relies on the absence of non-transactional accesses; and (2) there are subtle interactions between the transactional implementation and the weak memory model underlying the non-transactional accesses.

In this article, we address these limitations of SI. We develop two simple lock-based reference implementations for SI that do not use timestamps. Our first implementation is *prescient* [19] in that it requires *a priori* knowledge of the data accessed by a transaction, and performs transactional writes *eagerly* (in-place). Our second implementation is non-prescient and carries out transactional writes *lazily* by first recording them in a local log and subsequently propagating them to memory at commit time. Our first implementation is simpler and may be better suited for understanding and developing a program logic for SI transactions, whilst our second implementation is more practical due to its non-prescience. We show that both implementations are sound and complete against the declarative SI specification and thus yield equivalent operational definitions for SI.

We then extend both our eager and lazy implementations to make them robust under uninstrumented non-transactional accesses, and characterise declaratively

the semantics we obtain. We call this extended model *robust snapshot isolation* (RSI) and show that it gives reasonable semantics with mixed-mode accesses.

To provide SI semantics, instead of timestamps, our implementations use *multiple-readers-single-writer* (MRSW) locks. They acquire locks in reader mode to take a snapshot of the memory locations accessed by a transaction and then promote the relevant locks to writer mode to enforce an ordering on transactions with write-write conflicts. As we discuss in §4, the equivalence of the RSI implementation and its declarative characterisation depends heavily upon the axiomatisation of MRSW locks: here, we opted for the weakest possible axiomatisation that does not order any concurrent reader lock operations and present an MRSW lock implementation that achieves this.

Outline In §2 we present an overview of our contributions by describing our reference implementations for both SI and RSI. In §3 we define the declarative framework for specifying STM programs. In §4 we present the declarative SI specification against which we demonstrate the soundness and completeness of our SI implementations. In §5 we formulate a declarative specification for RSI and demonstrate the soundness and completeness of our RSI implementations. We discuss related and future work in §6.³

2 Background and Main Ideas

As noted earlier, the key challenge in specifying STM transactions lies in accounting for the interactions between mixed-mode accesses to the same data. One simple approach is to treat each non-transactional access as a singleton mini-transaction and to provide *strong isolation* [27,9], i.e. full isolation between transactional and non-transactional code. This, however, requires *instrumenting* non-transactional accesses to adhere to same access policies as transactional ones (e.g. acquiring the necessary locks), which incurs a substantial performance penalty for non-transactional code. A more practical approach is to enforce isolation only amongst transactional accesses, an approach known as *weak isolation* [27,9], adopted by the relaxed transactions of C++ [2].

As our focus is on STMs with SI guarantees, instrumenting non-transactional accesses is not feasible. In particular, as we expect many more non-transactional accesses than transactional ones, we do not want to incur any performance degradation on non-transactional code when executed in parallel with transactional code. As such, we opt for an STM with SI guarantees under *weak isolation*. Under weak isolation, however, transactions with explicit abort instructions are problematic as their intermediate state may be observed by non-transactional code. As such, weakly isolated STMs (e.g. C++ relaxed transactions [2]) often forbid explicit aborts altogether. Throughout our development we thus make two simplifying assumptions: (1) transactions are not nested; and (2) there are no explicit abort instructions, following the example of weakly isolated relaxed

³ A full version of this article is available at [31].

T1: $[a := x; //0 \parallel x := a + 1;$ (LU) Lost Update SI: ✗	T2: $[b := x; //0 \parallel x := b + 1;$	T1: $[a := x; //0 \parallel y := 1;$	T2: $[b := y; //0 \parallel x := 1;$	T1: $[y := 1;$ T3: $[a := x; //0 \parallel$ T2: $[b := y; //0 \parallel x := 1;$ (WS2) Write Skew Variant SI: ✓
T1: $[x := 1; \parallel y := 1;$ T3: $[a := y; //2 \parallel$ (LU2) Lost Update Variant SI: ✗	T2: $[b := x; //0 \parallel y := 2;$	$x := 1;$ T1: $[a := z; \parallel b := y; //0 \parallel$	$y := 1;$ T2: $[c := z; \parallel d := x; //0$	$x := 1; \parallel y := 1; \parallel$ T2: $[a := y; //1 \parallel b := x; //0$ (MPT) Message Passing RSI: ✗

Fig. 1: Litmus tests illustrating transaction anomalies and their admissibility under SI and RSI. In all tests, initially, $x = y = z = 0$. The $//v$ annotation next to a read records the value read.

transactions of C++. As we describe later in §2.3, it is straightforward to lift the latter restriction (2) for our lazy implementations.

For non-transactional accesses, we naturally have to pick some consistency model. For simplicity and uniformity, we pick the release/acquire (RA) subset of the C++ memory model [6,23], a well-behaved platform-independent memory model, whose compilation to x86 requires no memory fences.

Snapshot Isolation (SI) The initial model of SI in [7] is described informally in terms of a multi-version concurrent algorithm as follows. A transaction T proceeds by taking a *snapshot* S of the shared objects. The execution of T is then carried out locally: read operations query S and write operations update S . Once T completes its execution, it attempts to *commit* its changes and succeeds *only if* it is not *write-conflicted*. Transaction T is write-conflicted if another *committed* transaction T' has written to a location also written to by T , since T recorded its snapshot. If T fails the conflict check it aborts and may restart; otherwise, it commits its changes, and its changes become visible to all other transactions that take a snapshot thereafter.

To realise this, the shared state is represented as a series of *multi-versioned* objects: each object is associated with a history of several versions at different *timestamps*. In order to obtain a snapshot, a transaction T chooses a *start-timestamp* t_0 , and reads data from the committed state as of t_0 , ignoring updates after t_0 . That is, updates committed after t_0 are invisible to T . In order to commit, T chooses a *commit-timestamp* t_c larger than any existing start- or commit-timestamp. Transaction T is deemed write-conflicted if another transaction T' has written to a location also written to by T *and* the commit-timestamp of T' is in the execution interval of T ($[t_0, t_c]$).

2.1 Towards an SI Reference Implementation without Timestamps

While the SI description above is suitable for understanding SI, it is not useful for integrating the SI model in a language such as C/C++ or Java. From a programmer’s perspective, in such languages the various threads directly access the *uninstrumented* (single-versioned) shared memory; they do not access their own instrumented snapshot at a particular timestamp, which is loosely related to the snapshots of other threads. Ideally, what we would therefore like is an equivalent description of SI in terms of accesses to uninstrumented shared memory and a synchronisation mechanism such as locks.

In what follows, we present our first lock-based reference implementation for SI that does not rely on timestamps. To do this, we assume that the locations accessed by a transaction can be statically determined. Specifically, we assume that each transaction T is supplied with its *read set*, RS , and *write set*, WS , containing those locations read and written by T , respectively (a static over-approximation of these sets suffices for soundness.). As such, our first reference implementation is *prescient* [19] in that it requires *a priori* knowledge of the locations accessed by the transaction. Later in §2.3 we lift this assumption and develop an SI reference implementation that is *non-prescient* and similarly does not rely on timestamps.

Conceptually, a candidate implementation of transaction T would (1) obtain a snapshot of the locations read by T ; (2) lock those locations written by T ; (3) execute T *locally*; and (4) unlock the locations written. The snapshot is obtained via `snapshot(RS)` in Fig. 3 where the values of locations in RS are recorded in a local array \mathbf{s} . The local execution of T is carried out by executing $\langle T \rangle$ in Fig. 3, which is obtained from T by (i) modifying read operations to read locally from the snapshot in \mathbf{s} , and (ii) updating the snapshot after each write operation. Note that the snapshot must be obtained *atomically* to reflect the memory state at a particular instance (*cf.* start-timestamp). An obvious way to ensure the snapshot atomicity is to lock the locations in the read set, obtain a snapshot, and unlock the read set. However, as we must allow for two transactions *reading* from the same location to execute in parallel, we opt for *multiple-readers-single-writer* (MRSW) locks.

Let us now try to make this general pattern more precise. As a first attempt, consider the implementation in Fig. 2a written in a simple while language, which releases all the reader locks at the end of the snapshot phase before acquiring any writer locks. This implementation is unsound as it admits the lost update (LU) anomaly in Fig. 1 disallowed under SI [11]. To understand this, consider a scheduling where T_2 runs between lines 3 and 4 of T_1 in Fig. 2a, which would result in T_1 having read a stale value. The problem is that the writer locks on WS are acquired too late, allowing two conflicting transactions to run concurrently. To address this, writer locks must be acquired early enough to pre-empt the concurrent execution of write-write-conflicting transactions. Note that locks have to be acquired early even for locations only written by a transaction to avoid exhibiting a variant of the lost update anomaly (LU2).

<pre> 1. for (x ∈ RS) lock_r x 2. snapshot(RS); 3. for (x ∈ RS) unlock_r x 4. for (x ∈ WS) lock_w x 5. (T); 6. for (x ∈ WS) unlock_w x </pre> <p style="text-align: center;">(a)</p>	<pre> 1. for (x ∈ WS) lock_w x; 2. for (x ∈ RS \ WS) lock_r x 3. snapshot(RS); 4. for (x ∈ RS \ WS) unlock_r x 5. (T); 6. for (x ∈ WS) unlock_w x </pre> <p style="text-align: center;">(b)</p>	<pre> 1. for (x ∈ RS ∪ WS) lock_r x 2. snapshot(RS); 3. for (x ∈ RS ∪ WS) { 4. if (x ∈ WS) promote x 5. else unlock_r x; } 6. (T); 7. for (x ∈ WS) unlock_w x </pre> <p style="text-align: center;">(c)</p>
<p>Sound: ✗ allows (LU), (LU2)</p>	<p>Sound: ✓ Complete: ✗ disallows (WS)</p>	<p>Sound: ✓ Complete: ✗ disallows (WS2)</p>

Fig. 2: Candidate SI implementations of transaction T given read/write sets RS,WS

As such, our second candidate implementation in Fig. 2b brings forward the acquisition of writer locks. Whilst this implementation is sound (and disallows lost update), it nevertheless disallows behaviours deemed valid under SI such as the write skew anomaly (WS) in Fig. 1, and is thus incomplete. The problem is that such early acquisition of writer locks not only pre-empts concurrent execution of *write-write-conflicting* transactions, but also those of *read-write-conflicting* transactions (e.g. WS) due to the exclusivity of writer locks.

To remedy this, in our third candidate implementation in Fig. 2c we first acquire weaker reader locks on all locations in RS or WS, and later *promote* the reader locks on WS to exclusive writer ones, while releasing the reader locks on RS. The promotion of a reader lock signals its intent for exclusive ownership and awaits the release of the lock by other readers before claiming it exclusively as a writer. To avoid deadlocks, we further assume that $RS \cup WS$ is ordered so that locks are promoted in the same order by all threads.

Although this implementation is “more complete” than the previous one, it is still incomplete as it disallows certain behaviour admitted by SI. In particular, consider a variant of the write skew anomaly (WS2) depicted in Fig. 1, which is admitted under SI, but not admitted by this implementation.

To understand why this is admitted by SI, recall the operational SI model using timestamps. Let the domain of timestamps be that of natural numbers \mathbb{N} . The behaviour of (WS2) can be achieved by assigning the following execution intervals for T1: $[t_0^{T1}=2, t_c^{T1}=2]$; T2: $[t_0^{T2}=1, t_c^{T2}=4]$; and T3: $[t_0^{T3}=3, t_c^{T3}=3]$. To see why the implementation in Fig. 2c does not admit the behaviour in (WS2), let us assume without loss of generality that x is ordered before y . Upon executing lines 3-5, a) T1 promotes y ; b) T2 promotes x and then c) releases the reader lock on y ; and d) T3 releases the reader lock on x . To admit the behaviour in (WS2), the release of y in (c) must occur before the promotion of y in (a) since otherwise T2 cannot read 0 for y . Similarly, the release of x in (d) must occur before its promotion in (b). On the other hand, since T3 is executed by the same thread after T1, we know that (a) occurs before (d). This however leads to circular execution: (b)→(c)→(a)→(d)→(b), which cannot be realised.

<pre> 0. LS := ∅; 1. for (x ∈ RS ∪ WS) lock_r x 2. snapshot(RS); 3. for (x ∈ RS \ WS) unlock_r x 4. for (x ∈ WS) { 5. if (can-promote x) LS.add(x) 6. else { 7. for (x ∈ LS) unlock_w x 8. for (x ∈ WS \ LS) unlock_r x 9. goto line 0 } 10. } 11. (T); 12. for (x ∈ WS) unlock_w x </pre>	<pre> snapshot(RS) \triangleq for (x ∈ RS) s_x := x snapshot_{RSI}(RS) \triangleq start: for (x ∈ RS) s_x := x for (x ∈ RS) { if (s_x != x) goto start } (a:=x) \triangleq a:=s_x (x:=a) \triangleq x:=a; s_x:=a (S₁; S₂) \triangleq (S₁); (S₂) (while(e) S) \triangleq while(e) (S) ... and so on ... </pre>
--	---

Fig. 3: SI implementation of transaction T given RS, WS; the code in blue ensures deadlock avoidance. The RSI implementation (§5) is obtained by replacing snapshot on line 2 with snapshot_{RSI}.

To overcome this, in our final candidate execution in Fig. 3 (ignoring the code in blue), after obtaining a snapshot, we *first* release the reader locks on RS, and *then* promote the reader locks on WS, rather than simultaneously in one pass. As we demonstrate in §4, the implementation in Fig. 3 is both *sound and complete* against its declarative SI specification.

Avoiding Deadlocks As two distinct reader locks on x may simultaneously attempt to promote their locks, promotion is done on a ‘first-come-first-served’ basis to avoid *deadlocks*. A call to `can-promote x` by reader *r* thus returns a boolean denoting either (i) successful promotion (true); or (ii) failed promotion as another reader *r'* is currently promoting a lock on x (false). In the latter case, *r* must release its reader lock on x to ensure the successful promotion of x by *r'* and thus avoid deadlocks. To this end, our implementation in Fig. 3 includes a deadlock avoidance mechanism (code in blue) as follows. We record a list LS of those locks on the write set that have been successfully promoted so far. When promoting a lock on x succeeds (line 5), the LS is extended with x. On the other hand, when promoting x fails (line 6), all those locks promoted so far (i.e. in LS) as well as those yet to be promoted (i.e. in WS \ LS) are released and the transaction is restarted.

Remark 1. Note that the deadlock avoidance code in blue does not influence the correctness of the implementation in Fig. 3, and is merely included to make the reference implementation more realistic. In particular, the implementation without the deadlock avoidance code is both sound and complete against the SI specification, provided that the conditional `can-promote` call on line 5 is replaced by the blocking `promote` call.

Avoiding Over-Synchronisation due to MRSW Locks Consider the store buffering program (SBT) shown in Fig. 1. If, for a moment, we ignore transactional accesses, our underlying memory model (RA)—as well as all other weak memory models—allows the annotated weak behaviour. Intuitively, placing the two transactions that only *read z* in (SBT) should still allow the weak behaviour since the two transactions do not need to synchronise in any way. Nevertheless, most MRSW lock implementations forbid this outcome because they use a single global counter to track the number of readers that have acquired the lock, which inadvertently also synchronises the readers with one another. As a result, the two read-only transactions act as memory fences forbidding the weak outcome of (SBT). To avoid such synchronisation, in the technical appendix [31] we provide a different MRSW implementation using a separate location for each thread so that reader lock acquisitions do not synchronise.

To keep the presentation simple, we henceforth assume an abstract specification of a MRSW lock library providing operations for acquiring/releasing reader/writer locks, as well as promoting reader locks to writer ones. We require that (1) calls to writer locks (to acquire, release or promote) *synchronise* with all other calls to the lock library; and (2) writer locks provide *mutual exclusion* while held. We formalise these notions in §4. These requirements do not restrict synchronisation between *two read* lock calls: two read lock calls may or may not synchronise. Synchronisation between read lock calls is relevant only for the completeness of our RSI implementation (handling mixed-mode code); for that result, we further require that (3) read lock calls not synchronise.

2.2 Handling Racy Mixed-Mode Accesses

Let us consider what happens when data accessed by a transaction is modified concurrently by an uninstrumented atomic non-transactional write. Since such writes do not acquire any locks, the snapshots taken may include values written by non-transactional accesses. The result of the snapshot then depends on the order in which the variables are read. Consider the (MPT) example in Fig. 1. In our implementation, if in the snapshot phase *y* is read before *x*, then the annotated weak behaviour is not possible because the underlying model (RA) disallows this weak “message passing” behaviour. If, however, *x* is read before *y*, then the weak behaviour is possible. In essence, this means that the SI implementation described so far is of little use when there are races between transactional and non-transactional code. Technically, our SI implementation violates *monotonicity* with respect to wrapping code inside a transaction. The weak behaviour of the (MPT) example is disallowed by RA if we remove the transaction block T2, and yet it is exhibited by our SI implementation with the transaction block.

To get monotonicity under RA, it suffices for the snapshots to read the variables in the same order they are accessed by the transactions. Since a static calculation of this order is not always possible, following [30], we achieve this by reading each variable twice. In more detail, our `snapshotRSI` implementation in Fig. 3 takes *two* snapshots of the locations read by the transaction, and checks that they both return the same values for each location. This ensures

that every location is read both before and after every other location in the transaction, and hence all the high-level happens-before orderings in executions of the transactional program are also respected by its implementation. As we demonstrate in §5, our RSI implementation is both *sound and complete* against our proposed declarative semantics for RSI. There is however one caveat: since equality of values is used to determine whether the two snapshots agree, we will miss cases where different non-transactional writes to a location write the same value. In our formal development (see §5), we thus assume that if multiple non-transactional writes write the same value to the same location, they cannot race with the same transaction. Note that this assumption cannot be lifted without instrumenting non-transactional writes, and thus impeding performance substantially. That is, to lift this restriction we must instead replace every non-transactional *write* $x := v$ with `lock_w x; x := v; unlock_w x`.

2.3 Non-Prescient Reference Implementations without Timestamps

Recall that the SI and RSI implementations in §2.1 are prescient in that they require knowledge of the read and write sets of transactions beforehand. In what follows we present alternative SI and RSI implementations that are *non-prescient*.

Non-Prescient SI Reference Implementation In Fig. 4 we present a *lazy* lock-based reference implementation for SI. This implementation is *non-prescient* and does not require *a priori* knowledge of the read set \mathbf{RS} and the write set \mathbf{WS} . Rather, the \mathbf{RS} and \mathbf{WS} are computed on the fly as the execution of the transaction unfolds. As with the SI implementation in Fig. 3, this implementation does not rely on timestamps and uses MRSW locks to synchronise concurrent accesses to shared data. As before, the implementation consults a local *snapshot* at \mathbf{s} for read operations. However, unlike the eager implementation in Fig. 3 where transactional writes are performed *in-place*, the implementation in Fig. 4 is *lazy* in that it logs the writes in the local array \mathbf{s} and propagates them to memory at commit time, as we describe shortly.

Ignoring the code in blue, the implementation in Fig. 4 proceeds with initialising \mathbf{RS} and \mathbf{WS} with \emptyset (line 1); it then populates the local snapshot array at \mathbf{s} with initial value \perp for each location \mathbf{x} (line 2). It then executes $\langle \mathbf{T} \rangle$ which is obtained from \mathbf{T} as follows. For each *read* operation $\mathbf{a} := \mathbf{x}$ in \mathbf{T} , first the value of $\mathbf{s}[\mathbf{x}]$ is inspected to ensure it contains a snapshot of \mathbf{x} . If this is not the case (i.e. $\mathbf{x} \notin \mathbf{RS} \cup \mathbf{WS}$), a reader lock on \mathbf{x} is acquired, a snapshot of \mathbf{x} is recorded in $\mathbf{s}[\mathbf{x}]$, and the read set \mathbf{RS} is extended with \mathbf{x} . The snapshot value in $\mathbf{s}[\mathbf{x}]$ is subsequently returned in \mathbf{a} . Analogously, for each write operation $\mathbf{x} := \mathbf{a}$, the \mathbf{WS} is extended with \mathbf{x} , and the written value is lazily logged in $\mathbf{s}[\mathbf{x}]$. Recall from our candidate executions in Fig. 2 that to ensure implementation correctness, for each written location \mathbf{x} , the implementation must first acquire a reader lock on \mathbf{x} , and subsequently promote it to a writer lock. As such, for each write operation in \mathbf{T} , the implementation first checks if a reader lock for \mathbf{x} has been acquired (i.e. $\mathbf{x} \in \mathbf{RS} \cup \mathbf{WS}$) and obtains one if this is not the case.

<pre> 0. LS:=∅; 1. RS:=∅; WS:=∅; 2. for (x∈Locs) s[x]:=⊥ 3. ⟨T⟩; 4. for (x∈RS\WS) unlock_r x 5. for (x∈WS) { 6. if (can-promote x) LS.add(x) 7. else { 8. for (x∈LS) unlock_w x 9. for (x∈WS\LS) unlock_r x 10. goto line 0 } } 11. for (x∈WS) x:=s[x] 12. for (x∈WS) unlock_w x </pre>	<pre> ⟨a:=x⟩ ≜ if (x∉RS∪WS) { lock_r x; RS.add(x); s[x]:=x; } a:=s[x]; ⟨x:=a⟩ ≜ if (x∉RS∪WS) lock_r x; WS.add(x); s[x]:=a; ⟨S₁;S₂⟩ ≜ ⟨S₁⟩;⟨S₂⟩ ⟨while(e) S⟩ ≜ while(e) ⟨S⟩ ... and so on... </pre>
--	--

Fig. 4: Non-prescient SI implementation of transaction T with RS and WS computed on the fly; the code in blue ensures deadlock avoidance.

Once the execution of $\langle T \rangle$ is completed, the implementation proceeds to *commit* the transaction. To this end, the reader locks on RS are released (line 4), reader locks on WS are promoted to writer ones (line 6), the writes logged in \mathbf{s} are propagated to memory (line 11), and finally the writer locks on WS are released (line 12). As we demonstrate later in §4, the implementation in Fig. 4 is both sound and complete against the declarative SI specification.

Note that the implementation in Fig. 4 is optimistic in that it logs the writes performed by the transaction in the local array \mathbf{s} and propagates them to memory at commit time, rather than performing the writes *in-place* as with its pessimistic counterpart in Fig. 3. As before, the code in blue ensures deadlock avoidance and is identical to its counterpart in Fig. 3. As before, this deadlock avoidance code does not influence the correctness of the implementation and is merely included to make the reference implementation more practical.

Non-Prescient RSI Reference Implementation In Fig. 5 we present a *lazy* lock-based reference implementation for RSI. As with its SI counterpart, this implementation is non-prescient and computes the RS and WS on the fly. As before, the implementation does not rely on timestamps and uses MRSW locks to synchronise concurrent accesses to shared data. Similarly, the implementation consults the local *snapshot* at \mathbf{s} for read operations, whilst logging write operations lazily in a *write sequence* at \mathbf{wseq} , as we describe shortly.

Recall from the RSI implementation in §2.1 that to ensure snapshot validity, each location is read twice to preclude intermediate non-transactional writes. As such, when writing to a location x , the initial value read (recorded in \mathbf{s}) must not be *overwritten* by the transaction to allow for subsequent validation of the snapshot. To this end, for each location x , the snapshot array \mathbf{s} contains a *pair* of values, (r, c) , where r denotes the snapshot value (initial value read), and c denotes the current value which may have overwritten the snapshot value.

Recall that under weak isolation, the intermediate values written by a transaction may be observed by non-transactional reads. For instance, given the

<pre> 0. LS:=∅; 1. RS:=∅; WS:=∅; wseq:= []; 2. for (x∈Locs) s[x]:= (⊥, ⊥) 3. (T); 4. for (x∈RS) {(r,-):=s[x]; 5. if (x!=r) { //read x again 6. for (x∈RS∪WS) unlock_r x 7. goto line 0 } } 8. for (x∈RS\WS) unlock_r x 9. for (x∈WS) { 10. if (can-promote x) LS.add(x) 11. else { 12. for (x∈LS) unlock_w x 13. for (x∈WS\LS) unlock_r x 14. goto line 0 } } 15. for ((x,v)∈wseq) x:=v 16. for (x∈WS) unlock_w x </pre>	<pre> (a:=x) ≜ if (x∉RS∪WS) { lock_r x; RS.add(x); r:=x; s[x]:= (r,r); } (-,c):=s[x]; a:=c; (x:=a) ≜ if (x∉RS∪WS) lock_r x WS.add(x); (r,-):=s[x]; s[x]:= (r,a); wseq:=wseq++[(x,a)]; (S₁;S₂) ≜ (S₁);(S₂) (while(e) S) ≜ while(e) (S) ... and so on... </pre>
--	---

Fig. 5: Non-prescient RSI implementation of transaction T with RS and WS computed on the fly; the code in blue ensures deadlock avoidance.

T: $[x := 1; x := 2 \mid a := x]$ program, the non-transactional read $a := x$, may read either 1 or 2 for x . As such, at commit time, it is not sufficient solely to propagate the last written value (in program order) to each location (e.g. to propagate only the $x := 2$ write in the example above). Rather, to ensure implementation completeness, one must propagate all written values to memory, in the order they appear in the transaction body. To this end, we track the values written by the transaction as a (FIFO) *write sequence* at location `wseq`, containing items of the form (x, v) , denoting the location written (x) and the associated value (v).

Ignoring the code in blue, the implementation in Fig. 5 initialises RS and WS with \emptyset , initialises `wseq` as an empty sequence $[]$ (line 1), and populates the local snapshot array `s` with initial value (\perp, \perp) for each location x (line 2). It then executes (T) , obtained from T in an analogous manner to that in Fig. 4. For every read $a:=x$ in (T) , the current value recorded for x in `s` (namely c when $s[x]$ holds $(-, c)$) is returned in `a`. Dually, for every write $x:=a$ in (T) , the current value recorded for x in `s` is updated to `a`, and the write is logged in the write sequence `wseq` by appending (x, a) to it.

Upon completion of (T) , the snapshot in `s` is *validated* (lines 4-7). Each location x in RS is thus read again and its value is compared against the snapshot value in `s[x]`. If validation fails (line 5), the locks acquired are released (line 6) and the transaction is restarted (line 7).

If validation succeeds, the transaction is committed: the reader locks on RS are released (line 8), the reader locks on WS are promoted (line 10), the writes in `wseq` are propagated to memory in FIFO order (line 15), and finally the writer locks on WS are released (line 16).

As we show in §5, the implementation in Fig. 5 is both sound and complete against our proposed declarative specification for RSI. As before, the code in blue ensures deadlock avoidance; it does not influence the implementation correctness and is merely included to make the implementation more practical.

Supporting Explicit Abort Instructions It is straightforward to extend the lazy implementations in Fig. 4 and Fig. 5 to handle transactions containing explicit abort instructions. More concretely, as the effects (writes) of a transaction are logged locally and are not propagated to memory until commit time, upon reaching an `abort` in $\langle T \rangle$ no roll-back is necessary, and one can simply release the locks acquired so far and return. That is, one can extend $\langle \cdot \rangle$ in Fig. 4 and Fig. 5, and define $\langle \text{abort} \rangle \triangleq \text{for } (x \in \text{RS} \cup \text{WS}) \text{ unlock_r } x; \text{ return.}$

3 A Declarative Framework for STM

We present the notational conventions used in the remainder of this article, and describe a general framework for declarative concurrency models. Later in this article, we present SI, its extension with non-transactional accesses, and their lock-based implementations as instances of this general definition.

Notation Given a relation r on a set A , we write $r^?$, r^+ and r^* for the reflexive, transitive and reflexive-transitive closure of r , respectively. We write r^{-1} for the inverse of r ; $r|_A$ for $r \cap (A \times A)$; $[A]$ for the identity relation on A , i.e. $\{(a, a) \mid a \in A\}$; $\text{irreflexive}(r)$ for $\nexists a. (a, a) \in r$; and $\text{acyclic}(r)$ for $\text{irreflexive}(r^+)$. Given two relations r_1 and r_2 , we write $r_1; r_2$ for their (left) relational composition, i.e. $\{(a, b) \mid \exists c. (a, c) \in r_1 \wedge (c, b) \in r_2\}$. Lastly, when r is a strict partial order, we write $r|_{\text{imm}}$ for the *immediate* edges in r : $\{(a, b) \in r \mid \nexists c. (a, c) \in r \wedge (c, b) \in r\}$.

Assume finite sets of *locations* LOC ; *values* VAL ; *thread identifiers* TID , and *transaction identifiers* TXID . We use x, y, z to range over locations, v over values, τ over thread identifiers, and ξ over transaction identifiers.

Definition 1 (Events). *An event is a tuple $\langle n, \tau, \xi, l \rangle$, where $n \in \mathbb{N}$ is an event identifier, $\tau \in \text{TID} \uplus \{0\}$ is a thread identifier (0 is used for initialisation events), $\xi \in \text{TXID} \uplus \{0\}$ is a transaction identifier (0 is used for non-transactional events), and l is an event label that takes one of the following forms:*

- A memory access label: $\text{R}(x, v)$ for reads; $\text{W}(x, v)$ for writes; and $\text{U}(x, v_r, v_w)$ for updates.
- A lock label: $\text{RL}(x)$ for reader lock acquisition; $\text{RU}(x)$ for reader lock release; $\text{WL}(x)$ for writer lock acquisition; $\text{WU}(x)$ for writer lock release; and $\text{PL}(x)$ for reader to writer lock promotion.

We typically use a, b , and e to range over events. The functions tid , tx , lab , typ , loc , val_r and val_w respectively project the thread identifier, transaction identifier, label, type (in $\{\text{R}, \text{W}, \text{U}, \text{RL}, \text{RU}, \text{WL}, \text{WU}, \text{PL}\}$), location, and read/written values of an event, where applicable. We assume only reads and writes are used in transactions ($\text{tx}(a) \neq 0 \implies \text{typ}(a) \in \{\text{R}, \text{W}\}$).

Given a relation r on events, we write r_{loc} for $\{(a, b) \in r \mid \text{loc}(a) = \text{loc}(b)\}$. Analogously, given a set A of events, we write A_x for $\{a \in A \mid \text{loc}(a) = x\}$.

Definition 2 (Execution graphs). *An execution graph, G , is a tuple of the form $(E, \text{po}, \text{rf}, \text{mo}, \text{lo})$, where:*

- E is a set of events, assumed to contain a set E_0 of initialisation events, consisting of a write event with label $\mathbb{W}(x, 0)$ for every $x \in \text{Loc}$. The sets of read events in E is denoted by $\mathcal{R} \triangleq \{e \in E \mid \text{typ}(e) \in \{\text{R}, \text{U}\}\}$; write events by $\mathcal{W} \triangleq \{e \in E \mid \text{typ}(e) \in \{\text{W}, \text{U}\}\}$; update events by $\mathcal{U} \triangleq \mathcal{R} \cap \mathcal{W}$; and lock events by $\mathcal{L} \triangleq \{e \in E \mid \text{typ}(e) \in \{\text{RL}, \text{RU}, \text{WL}, \text{WU}, \text{PL}\}\}$. The sets of reader lock acquisition and release events, \mathcal{RL} and \mathcal{RU} , writer lock acquisition and release events, \mathcal{WL} and \mathcal{WU} , and lock promotion events \mathcal{PL} are defined analogously. The set of transactional events in E is denoted by \mathcal{T} ($\mathcal{T} \triangleq \{e \in E \mid \text{tx}(e) \neq 0\}$); and the set of non-transactional events is denoted by \mathcal{NT} ($\mathcal{NT} \triangleq E \setminus \mathcal{T}$).
- $\text{po} \subseteq E \times E$ denotes the ‘program-order’ relation, defined as a disjoint union of strict total orders, each ordering the events of one thread, together with $E_0 \times (E \setminus E_0)$ that places the initialisation events before any other event. We assume that events belonging to the same transaction are ordered by po , and that any other event po -between them also belongs to the same transaction.
- $\text{rf} \subseteq \mathcal{W} \times \mathcal{R}$ denotes the ‘reads-from’ relation, defined between write and read events of the same location with matching read and written values; it is total and functional on reads, i.e. every read is related to exactly one write.
- $\text{mo} \subseteq \mathcal{W} \times \mathcal{W}$ denotes the ‘modification-order’ relation, defined as a disjoint union of strict total orders, each ordering the write events on one location.
- $\text{lo} \subseteq \mathcal{L} \times \mathcal{L}$ denotes the ‘lock-order’ relation, defined as a disjoint union of strict orders, each of which (partially) ordering the lock events to one location.

In the context of an execution graph $G=(E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ —we often use “ G .” as a prefix to make this explicit—the ‘same-transaction’ relation, $\text{st} \in \mathcal{T} \times \mathcal{T}$, is the equivalence relation given by $\text{st} \triangleq \{(a, b) \in \mathcal{T} \times \mathcal{T} \mid \text{tx}(a) = \text{tx}(b)\}$. Given a relation $r \subseteq E \times E$, we write $r_{\mathcal{T}}$ for lifting r to transaction classes: $r_{\mathcal{T}} \triangleq \text{st}; (r \setminus \text{st}); \text{st}$. For instance, when $(w, r) \in \text{rf}$, w is a transaction ξ_1 event and r is a transaction ξ_2 event, then all all events in ξ_1 are $\text{rf}_{\mathcal{T}}$ -related to all events in ξ_2 . Analogously, we write $r_{\mathcal{I}}$ to restrict r to its *intra-transactional* edges (within a transaction): $r_{\mathcal{I}} \triangleq r \cap \text{st}$; and write $r_{\mathcal{E}}$ to restrict r to its *extra-transactional* edges (outside a transaction): $r_{\mathcal{E}} \triangleq r \setminus \text{st}$. Lastly, the ‘reads-before’ relation is defined by $\text{rb} \triangleq (\text{rf}^{-1}; \text{mo}) \setminus [E]$. Intuitively, rb relates a read r to all writes w that are mo -after the write r reads from; i.e. when $(w', r) \in \text{rf}$ and $(w', w) \in \text{mo}$, then $(r, w) \in \text{rb}$. In the transactional literature, this is known as the *anti-dependency* relation [4,3].

Execution graphs of a given program represent traces of shared memory accesses generated by the program. The set of execution graphs associated with programs written in our while language can be straightforwardly defined by induction over the structure of programs as in e.g. [35]. Each execution of a program P has a particular program *outcome*, prescribing the final values of local variables in each thread. In this initial stage, the execution outcomes are almost unrestricted as there are very few constraints on the rf , mo and lo relations. Such restrictions and thus the permitted outcomes of a program are determined by defining the set of *consistent* executions, which is defined separately for each model we consider. Given a program P and a model M , the set $\text{outcomes}_M(P)$ collects the outcomes of every M -consistent execution of P .

4 Snapshot Isolation (SI)

We present a declarative specification of SI and demonstrate that the SI implementations presented in Fig. 3 and Fig. 4 are both sound and complete with respect to the SI specification.

In [11] Cerone and Gotsman developed a declarative specification for SI using dependency graphs [4,3]. Below we adapt their specification to the notation of §3. As with [11], throughout this section, we take *SI execution graphs* to be those in which $E = \mathcal{T} \subseteq (\mathcal{R} \cup \mathcal{W}) \setminus \mathcal{U}$. That is, the SI model handles transactional code only, consisting solely of read and write events (excluding updates).

Definition 3 (SI consistency [11]). *An SI execution $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ is SI-consistent if the following conditions hold:*

- $\text{rf}_I \cup \text{mo}_I \cup \text{rb}_I \subseteq \text{po}$ (INT)
- $\text{acyclic}((\text{po}_T \cup \text{rf}_T \cup \text{mo}_T); \text{rb}_T?)$ (EXT)

Informally, (INT) ensures the consistency of each transaction internally, while (EXT) provides the synchronisation guarantees among transactions. In particular, we note that the two conditions together ensure that if two read events in the same transaction read from the same location x , and no write to x is po-between them, then they must read from the same write (known as ‘internal read consistency’).

Next, we provide an alternative equivalent formulation of SI-consistency which will serve as the basis of our extension with non-transactional accesses in §5.

Proposition 1. *An SI execution $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ is SI-consistent if and only if INT holds and the ‘SI-happens-before’ relation $\text{si-hb} \triangleq (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T \cup \text{si-rb})^+$ is irreflexive, where $\text{si-rb} \triangleq [\mathcal{R}_E]; \text{rb}_T; [\mathcal{W}]$ and $\mathcal{R}_E \triangleq \{r \mid \exists w. (w, r) \in \text{rf}_E\}$.*

Proof. The full proof is given in the technical appendix [31].

Intuitively, SI-happens-before orders events of different transactions; this order is due to either the program order (po_T), or synchronisation enforced by the implementation ($\text{rf}_T \cup \text{mo}_T \cup \text{si-rb}$). By contrast, events of the same transaction are unordered, as the implementation may well execute them in a different order (in particular, by taking a snapshot, it executes external reads before the writes).

In more detail, the rf_T corresponds to transactional synchronisation due to *causality*, i.e. when one transaction T_2 observes an effect of an earlier transaction T_1 . The inclusion of rf_T ensures that T_2 cannot read from T_1 without observing its *entire* effect. This in turn ensures that transactions exhibit ‘all-or-nothing’ behaviour: they cannot mix-and-match the values they read. For instance, if T_1 writes to both x and y , transaction T_2 may not read x from T_1 but read y from an earlier (in ‘happens-before’ order) transaction T_0 .

The mo_T corresponds to transactional synchronisation due to *write-write conflicts*. Its inclusion enforces write-conflict-freedom of SI transactions: if T_1 and T_2 both write to x via events w_1 and w_2 such that $(w_1, w_2) \in \text{mo}$, then T_1 must commit before T_2 , and thus its entire effect must be visible to T_2 .

To understand **si-rb**, first note that \mathcal{R}_E denotes the *external* transactional reads (i.e. those reading a value written by another transaction). That is, the \mathcal{R}_E are the read events that get their values from the transactional snapshot phases. By contrast, internal reads (those reading a value written by the same transaction) happen only after the snapshot is taken. Now let there be an **rb_T** edge between two transactions, T_1 and T_2 . This means there exist a read event r of T_1 and a write event w of T_2 such that $(r, w) \in \mathbf{rb}$; i.e. there exists w' such that $(w', r) \in \mathbf{rf}$ and $(w', w) \in \mathbf{mo}$. If r reads internally (i.e. w' is an event in T_1), then T_1 and T_2 are conflicting transactions and as accounted by **mo_T** described above, all events of T_1 happen before those of T_2 . Now, let us consider the case when r reads externally (w' is not in T_1). From the timestamped model of SI, there exists a start-timestamp $t_0^{T_1}$ as of which the T_1 snapshot (all its external reads including r) is recorded. Similarly, there exists a commit-timestamp $t_c^{T_2}$ as of which the updates of T_2 (including w) are committed. Moreover, since $(r, w) \in \mathbf{rb}$ we know $t_0^{T_1} < t_c^{T_2}$ (otherwise r must read the value written by w and not w'). That is, we know all events in the snapshot of T_1 (i.e. all external reads in T_1) happen before all writes of T_2 .⁴

We use the declarative framework in §3 to formalise the semantics of our implementation. Here, our programs include only non-transactional code, and thus *implementation execution graphs* are taken as those in which $\mathcal{T} = \emptyset$. Furthermore, we assume that locks in implementation programs are used in a *well-formed* manner: the sequence of lock events for *each location*, in each thread (following po), should match (a prefix of) the regular expression $(\mathbf{RL} \cdot \mathbf{RU} \mid \mathbf{WL} \cdot \mathbf{WU} \mid \mathbf{RL} \cdot \mathbf{PL} \cdot \mathbf{WU})^*$. For instance, a thread never releases a lock, without having acquired it earlier in the program. As a consistency predicate on execution graphs, we use the C11 release/acquire consistency augmented with certain constraints on lock events.

Definition 4. *An implementation execution graph $G = (E, \text{po}, \mathbf{rf}, \mathbf{mo}, \mathbf{lo})$ is RA-consistent if the following hold, where $\mathbf{hb} \triangleq (\text{po} \cup \mathbf{rf} \cup \mathbf{lo})^+$ denotes the ‘RA-happens-before’ relation:*

- $\forall x. \forall a \in \mathcal{WL}_x \cup \mathcal{WU}_x \cup \mathcal{PL}_x, b \in \mathcal{L}_x. a = b \vee (a, b) \in \mathbf{lo} \vee (b, a) \in \mathbf{lo}$ (WSYNC)
- $[\mathcal{WL} \cup \mathcal{PL}]; (\mathbf{lo} \setminus \text{po}); [\mathcal{L}] \subseteq \text{po}; [\mathcal{WU}]; \mathbf{lo}$ (WEX)
- $[\mathcal{RL}]; (\mathbf{lo} \setminus \text{po}); [\mathcal{WL} \cup \mathcal{PL}] \subseteq \text{po}; [\mathcal{RU} \cup \mathcal{PL}]; \mathbf{lo}$ (RSHARE)
- acyclic($\mathbf{hb}_{\text{loc}} \cup \mathbf{mo} \cup \mathbf{rb}$) (ACYC)

The (WSYNC) states that write lock calls (to acquire, release or promote) *synchronise* with all other calls to the same lock.

The next two constraints ensure the ‘single-writer-multiple-readers’ paradigm. In particular, (WEX) states that write locks provide *mutual exclusion* while held: any lock event l of thread τ **lo**-after a write lock acquisition or promotion event l' of another thread τ' , is **lo**-after a subsequent write lock release event u of τ' (i.e. $(l', u) \in \text{po}$ and $(u, l) \in \mathbf{lo}$). As such, the lock cannot be acquired (in read or write mode) by another thread until it has been released by its current owner.

The (RSHARE) analogously states that once a thread acquires a lock in read mode, the lock cannot be acquired in write mode by other threads until it has

⁴ By taking **rb_T** instead of **si-rb** in Prop. 1 one obtains a characterisation of *serialisability*.

either been released, or promoted to a writer lock (and subsequently released) by its owner. Note that this does not preclude other threads from simultaneously acquiring the lock in read mode. In the technical appendix [31] we present two MRSW lock implementations that satisfy the conditions outlined above.

The last constraint (**ACYC**) is that of C11 RA consistency [23], with the **hb** relation extended with **lo**.

Remark 2. Our choice of implementing the SI STMs on top of the RA fragment is purely for presentational convenience. Indeed, it is easy to observe that execution graphs of $\langle P \rangle$ are data race free, and thus, **ACYC** could be replaced by any condition that implies $\forall x. ([\mathcal{W}_x]; (\text{po} \cup \text{lo})^+; [\mathcal{W}_x]; (\text{po} \cup \text{lo})^+; [\mathcal{R}_x]) \cap \text{rf} = \emptyset$ and that is implied by $\text{acyclic}(\text{po} \cup \text{rf} \cup \text{lo} \cup \text{mo} \cup \text{rb})$. In particular, the C11 non-atomic accesses or sequentially consistent accesses may be used.

We next show that our SI implementations in Fig. 3 and Fig. 4 are *sound and complete* with respect to the declarative specification given above. The proofs are non-trivial and the full proofs are given in the technical appendix [31].

Theorem 1 (Soundness and completeness). *Let P be a transactional program; let $\langle P \rangle_E$ denote its eager implementation as given in Fig. 3 and $\langle P \rangle_L$ denote its lazy implementation as given in Fig. 4. Then:*

$$\text{outcomes}_{\text{SI}}(P) = \text{outcomes}_{\text{RA}}(\langle P \rangle_E) = \text{outcomes}_{\text{RA}}(\langle P \rangle_L)$$

Proof. The full proofs for both implementations is given in the technical appendix [31].

Stronger MRSW Locks As noted in §2, for both (prescient and non-prescient) SI implementations our soundness and completeness proofs show that the same result holds for a stronger lock specification, in which reader locks synchronise as well. Formally, this specification is obtained by adding the following to Def. 4:

$$- \forall x. \forall a, b \in \mathcal{RL}_x \cup \mathcal{RU}_x. a = b \vee (a, b) \in \text{lo} \vee (b, a) \in \text{lo} \quad (\text{RSYNC})$$

Soundness of this stronger specification ($\text{outcomes}_{\text{RA}}(\langle P \rangle_x) \subseteq \text{outcomes}_{\text{SI}}(P)$ for $x \in \{E, L\}$) follows immediately from Thm. 1. Completeness ($\text{outcomes}_{\text{SI}}(P) \subseteq \text{outcomes}_{\text{RA}}(\langle P \rangle_x)$ for $x \in \{E, L\}$), however, is more subtle, as we need to additionally satisfy (**RSYNC**) when constructing **lo**. While we can do so for SI, it is essential for the completeness of our RSI implementations that reader locks not synchronise, as shown by (**SBT**) in §2.

In the technical appendix [31] we present two MRSW lock implementations *sound* against the **lo** conditions in Def. 4. Additionally, the first implementation is *complete* against the conditions of Def. 4 augmented with (**RSYNC**), whilst the second is complete against the conditions of Def. 4 alone.




Fig. 6: RSI-inconsistent executions due to (a) rsi-po; (b) $[\mathcal{N}\mathcal{T}]$; rf; st; (c) $(\text{mo}; \text{rf})_{\mathcal{T}}$

5 Robust Snapshot Isolation (RSI)

We explore the semantics of SI STMs in the presence of non-transactional code with *weak isolation* guarantees (see §2). We refer to this model as *robust snapshot isolation* (RSI), due to its ability to provide SI guarantees between transactions even in the presence of non-transactional code. We propose the first declarative specification of RSI programs and develop two lock-based reference implementations that are both *sound and complete* against our proposed specification.

A Declarative Specification of RSI STMs We formulate a declarative specification of RSI semantics by adapting the SI semantics in Prop. 1 to account for non-transactional accesses. To specify the abstract behaviour of RSI programs, *RSI execution graphs* are taken to be those in which $\mathcal{L} = \emptyset$. Moreover, as with SI graphs, RSI execution graphs are those in which $\mathcal{T} \subseteq (\mathcal{R} \cup \mathcal{W}) \setminus \mathcal{U}$. That is, RSI transactions comprise solely read and write events, excluding updates.

Definition 5 (RSI consistency). *An execution $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ is RSI-consistent iff INT holds and acyclic($\text{rsi-hb}_{\text{loc}} \cup \text{mo} \cup \text{rb}$), where $\text{rsi-hb} \triangleq (\text{rsi-po} \cup \text{rsi-rf} \cup \text{mo}_{\mathcal{T}} \cup \text{si-rb})^+$ is the ‘RSI-happens-before’ relation, with $\text{rsi-po} \triangleq (\text{po} \setminus \text{po}_{\mathcal{T}}) \cup [\mathcal{W}]$; $\text{po}_{\mathcal{T}}$; $[\mathcal{W}]$ and $\text{rsi-rf} \triangleq (\text{rf}; [\mathcal{N}\mathcal{T}]) \cup ([\mathcal{N}\mathcal{T}]; \text{rf}; \text{st}) \cup \text{rf}_{\mathcal{T}} \cup (\text{mo}; \text{rf})_{\mathcal{T}}$.*

As with SI and RA, we characterise the set of executions admitted by RSI as graphs that lack cycles of certain shapes. To account for non-transactional accesses, similar to RA, we require $\text{rsi-hb}_{\text{loc}} \cup \text{mo} \cup \text{rb}$ to be acyclic (recall that $\text{rsi-hb}_{\text{loc}} \triangleq \{(a, b) \in \text{rsi-hb} \mid \text{loc}(a) = \text{loc}(b)\}$). The RSI-happens-before relation rsi-hb includes both the synchronisation edges enforced by the transactional implementation (as in si-hb), and those due to non-transactional accesses (as in hb of the RA consistency). The rsi-hb relation itself is rather similar to si-hb . In particular, the $\text{mo}_{\mathcal{T}}$ and si-rb subparts can be justified as in si-hb ; the difference between the two lies in rsi-po and rsi-rf .

To justify rsi-po , recall from §4 that si-hb includes $\text{po}_{\mathcal{T}}$. The rsi-po is indeed a strengthening of $\text{po}_{\mathcal{T}}$ to account for non-transactional events: it additionally includes (i) po to and from non-transactional events; and (ii) po between two *write* events in a transaction. We believe (i) comes as no surprise to the reader;

for (ii), consider the execution graph in Fig. 6a, where transaction T is denoted by the dashed box labelled T, comprising the write events w_1 and w_2 . Removing the T block (with w_1 and w_2 as non-transactional writes), this execution is deemed inconsistent, as this weak “message passing” behaviour is disallowed in the RA model. We argue that the analogous transactional behaviour in Fig. 6a must be similarly disallowed to maintain monotonicity with respect to wrapping non-transactional code in a transaction (see Thm. 3). As in SI, we cannot include the entire po in rsi-hb because the write-read order in transactions is not preserved by the implementation.

Similarly, rsi-rf is a strengthening of rf_T to account for non-transactional events: in the absence of non-transactional events rsi-rf reduces to rf_T ∪ (mo; rf)_T which is contained in si-hb. The rf; [NT] part is required to preserve the ‘happens-before’ relation for non-transactional code. That is, as rf is included in the hb relation of underlying memory model (RA), it is also included in rsi-hb.

The [NT]; rf; st part asserts that in an execution where a read event r of transaction T reads from a non-transactional write w , the snapshot of T reads from w and so all events of T happen after w . Thus, in Fig. 6b, r' cannot read from the overwritten initialisation write to y .

For the (mo; rf)_T part, consider the execution graph in Fig. 6c where there is a write event w of transaction T₁ and a read event r of transaction T₂ such that $(w, r) \in \text{mo; rf}$. Then, transaction T₂ must acquire the read lock of loc(w) after T₁ releases the writer lock, which in turn means that every event of T₁ happens before every event of T₂.

Remark 3. Recall that our choice of modelling SI and RSI STMs in the RA fragment is purely for presentational convenience (see Remark 2). Had we chosen a different model, the RSI consistency definition (Def. 5) would largely remain unchanged, with the exception of rsi-rf \triangleq (sw; [NT]) ∪ ([NT]; sw; st) ∪ rf_T ∪ (mo; rf)_T, where in the highlighted changes the rf relation is replaced with sw, denoting the ‘synchronises-with’ relation. As in the RA model sw \triangleq rf, we have inlined this in Def. 5.

SI and RSI Consistency We next demonstrate that in the absence of non-transactional code, the definitions of SI-consistency (Prop. 1) and RSI-consistency (Def. 5) coincide. That is, for all executions G , if $G.\mathcal{NT} = \emptyset$, then G is SI-consistent if and only if G is RSI-consistent.

Theorem 2. *For all executions G , if $G.\mathcal{NT} = \emptyset$, then:*

$$G \text{ is SI-consistent} \iff G \text{ is RSI-consistent}$$

Proof. The full proof is given in the technical appendix [31].

Note that the above theorem implies that for all transactional programs P , if P contains no non-transactional accesses, then $\text{outcomes}_{\text{SI}}(P) = \text{outcomes}_{\text{RSI}}(P)$.

RSI Monotonicity We next prove the *monotonicity* of RSI when wrapping non-transactional events into a transaction. That is, wrapping a block of non-transactional code inside a new transaction does not introduce additional behaviours. More concretely, given a program P , when a block of non-transactional code in P is wrapped inside a new transaction to obtain a new program P_T , then $\text{outcomes}_{\text{RSI}}(P_T) \subseteq \text{outcomes}_{\text{RSI}}(P)$. This is captured in the theorem below, with its full proof given in the technical appendix [31].

Theorem 3 (Monotonicity). *Let P_T and P be RSI programs such that P_T is obtained from P by wrapping a block of non-transactional code inside a new transaction. Then:*

$$\text{outcomes}_{\text{RSI}}(P_T) \subseteq \text{outcomes}_{\text{RSI}}(P)$$

Proof. The full proof is given in the technical appendix [31].

Lastly, we show that our RSI implementations in §2 (Fig. 3 and Fig. 5) are sound and complete with respect to Def. 5. This is captured in the theorem below. The soundness and completeness proofs are non-trivial; the full proofs are given in the technical appendix [31].

Theorem 4 (Soundness and completeness). *Let P be a program that possibly mixes transactional and non-transactional code. Let $(P)_E$ denote its eager RSI implementation as given in Fig. 3 and $(P)_L$ denote its lazy RSI implementation as given in Fig. 5. If for every location x and value v , every RSI-consistent execution of P contains either (i) at most one non-transactional write of v to x ; or (ii) all non-transactional writes of v to x are happens-before-ordered with respect to all transactions accessing x , then:*

$$\text{outcomes}_{\text{RSI}}(P) = \text{outcomes}_{\text{RA}}((P)_E) = \text{outcomes}_{\text{RA}}((P)_L)$$

Proof. The full proofs for both implementations are given in the technical appendix [31].

6 Related and Future Work

Much work has been done in formalising the semantics of weakly consistent *database transactions* [7,34,3,4,10,12,11,13,14,18], both operationally and declaratively. On the operational side, Berenson et al. [7] gave an operational model of SI as a multi-version concurrent algorithm. Later, Sovran et al. [34] described and operationally defined the *parallel snapshot isolation* model (PSI), as a close relative of SI with weaker guarantees.

On the declarative side, Adya et al. [3,4] introduced *dependency graphs* (similar to execution graphs of our framework in §3) for specifying transactional semantics and formalised several ANSI isolation levels. Cerone et al. [10,12] introduced *abstract executions* and formalised several isolation levels including SI and PSI. Later in [11], they used dependency graphs of Adya to develop equivalent SI

and PSI semantics; recently in [13], they provided a set of algebraic laws for connecting these two declarative styles.

To facilitate client-side reasoning about the behaviour of database transactions, Gotsman et al. [18] developed a proof rule for proving invariants of client applications under a number of consistency models.

Recently, Kaki et al. [21] developed a program logic to reason about transactions under ANSI SQL isolation levels (including SI). To do this, they formulated an operational model of such programs (parametric in the isolation level). They then proved the soundness of their *logic* with respect to their proposed operational model. However, the authors did not establish the *soundness* or *completeness* of their *operational model* against existing formal semantics, e.g. [11]. The lack of the completeness result means that their proposed operational model may exclude behaviours deemed valid by the corresponding declarative models. This is a particular limitation as possibly many valid behaviours cannot be shown correct using the logic and is thus detrimental to its usability.

By contrast, transactional semantics in the STM setting with mixed (both transactional and non-transactional) accesses is under-explored on both operational and declarative sides. Recently, Dongol et al. [17] applied execution graphs [5] to specify *serialisable* STM programs under weak memory models. Raad et al. [30] formalised the semantics of PSI STMs declaratively (using execution graphs) and operationally (as lock-based reference implementations). Neither work, however, handles the semantics of SI STMs with weak isolation guarantees.

Finally, Khyzha et al. [22] formalise the sufficient conditions on STMs and their programs that together ensure strong isolation. That is, non-transactional accesses can be viewed as singleton transactions (transactions containing single instructions). However, their conditions require *serialisability* for fully transactional programs, and as such, RSI transactions do not meet their conditions. Nevertheless, we conjecture that a DRF guarantee for strong atomicity, similar to [22], may be established for RSI. That is, if all executions of a fully transactional program have no races between singleton and non-singleton transactions, then it is safe to replace all singleton transactions by non-transactional accesses.

In the future, we plan to build on the work presented here by developing reasoning techniques that would allow us to verify properties of STM programs. This can be achieved by either extending existing program logics for weak memory, or developing new ones for currently unsupported models. In particular, we can reason about the SI models presented here by developing custom proof rules in the existing program logics for RA such as [24,35].

Acknowledgements We thank the VMCAI reviewers for their constructive feedback. The first author was supported in part by a European Research Council (ERC) Consolidator Grant for the project “RustBelt”, under the European Union’s Horizon 2020 Framework Programme (grant agreement number 683289). The second author was supported by the Israel Science Foundation (grant number 5166651), and by Len Blavatnik and the Blavatnik Family foundation.

References

1. The Clojure Language: Refs and Transactions, <http://clojure.org/refs>
2. Technical specification for C++ extensions for transactional memory (2015), <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2015/n4514.pdf>
3. Adya, A.: Weak consistency: A generalized theory and optimistic implementations for distributed transactions. Ph.D. thesis, MIT (1999)
4. Adya, A., Liskov, B., O’Neil, P.: Generalized isolation level definitions. In: Proceedings of the 16th International Conference on Data Engineering. pp. 67–78 (2000)
5. Alglave, J., Maranget, L., Tautschnig, M.: Herding cats: Modelling, simulation, testing, and data mining for weak memory. *ACM Trans. Program. Lang. Syst.* **36**(2), 7:1–7:74 (2014)
6. Batty, M., Owens, S., Sarkar, S., Sewell, P., Weber, T.: Mathematizing C++ concurrency. In: Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. pp. 55–66 (2011)
7. Berenson, H., Bernstein, P., Gray, J., Melton, J., O’Neil, E., O’Neil, P.: A critique of ANSI SQL isolation levels. In: Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data. pp. 1–10 (1995)
8. Bieniusa, A., Fuhrmann, T.: Consistency in hindsight: A fully decentralized STM algorithm. In: Proceedings of the 2010 IEEE International Symposium on Parallel and Distributed Processing, IPDPS 2010. pp. 1 – 12 (2010)
9. Blundell, C., C. Lewis, E., M. K. Martin, M.: Deconstructing transactions: The subtleties of atomicity. In: 4th Annual Workshop on Duplicating, Deconstructing, and Debunking (2005)
10. Cerone, A., Bernardi, G., Gotsman, A.: A framework for transactional consistency models with atomic visibility. In: Proceedings of the 26th International Conference on Concurrency Theory. pp. 58–71 (2015)
11. Cerone, A., Gotsman, A.: Analysing snapshot isolation. In: Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing. pp. 55–64 (2016)
12. Cerone, A., Gotsman, A., Yang, H.: Transaction chopping for parallel snapshot isolation. In: Proceedings of the 29th International Symposium on Distributed Computing - Volume 9363. pp. 388–404 (2015)
13. Cerone, A., Gotsman, A., Yang, H.: Algebraic laws for weak consistency. In: CONCUR (2017)
14. Crooks, N., Pu, Y., Alvisi, L., Clement, A.: Seeing is believing: A client-centric specification of database isolation. In: Proceedings of the ACM Symposium on Principles of Distributed Computing. pp. 73–82. PODC ’17, ACM, New York, NY, USA (2017). <https://doi.org/10.1145/3087801.3087802>, <http://doi.acm.org/10.1145/3087801.3087802>
15. Daudjee, K., Salem, K.: Lazy database replication with snapshot isolation. In: Proceedings of the 32Nd International Conference on Very Large Data Bases. pp. 715–726 (2006)
16. Dias, R.J., Distefano, D., Seco, J.a.C., Lourenço, J.a.M.: Verification of snapshot isolation in transactional memory java programs. In: Proceedings of the 26th European Conference on Object-Oriented Programming. pp. 640–664. ECOOP’12, Springer-Verlag, Berlin, Heidelberg (2012). <https://doi.org/10.1007/978-3-642-31057-7-28>, <http://dx.doi.org/10.1007/978-3-642-31057-7-28>
17. Dongol, B., Jagadeesan, R., Riely, J.: Transactions in relaxed memory architectures. *Proc. ACM Program. Lang.* **2**(POPL), 18:1–18:29 (Dec 2017). <https://doi.org/10.1145/3158106>, <http://doi.acm.org/10.1145/3158106>

18. Gotsman, A., Yang, H., Ferreira, C., Najafzadeh, M., Shapiro, M.: 'cause i'm strong enough: Reasoning about consistency choices in distributed systems. In: Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. pp. 371–384. POPL '16, ACM, New York, NY, USA (2016). <https://doi.org/10.1145/2837614.2837625>, <http://doi.acm.org/10.1145/2837614.2837625>
19. Harris, T., Larus, J., Rajwar, R.: Transactional Memory, 2Nd Edition. Morgan and Claypool Publishers, 2nd edn. (2010)
20. Herlihy, M., Moss, J.E.B.: Transactional memory: Architectural support for lock-free data structures. In: Proceedings of the 20th Annual International Symposium on Computer Architecture. pp. 289–300 (1993)
21. Kaki, G., Nagar, K., Najafzadeh, M., Jagannathan, S.: Alone together: Compositional reasoning and inference for weak isolation. Proc. ACM Program. Lang. **2**(POPL), 27:1–27:34 (Dec 2017). <https://doi.org/10.1145/3158115>, <http://doi.acm.org/10.1145/3158115>
22. Khyzha, A., Attiya, H., Gotsman, A., Rinetzky, N.: Safe privatization in transactional memory. In: Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming. pp. 233–245 (2018)
23. Lahav, O., Giannarakis, N., Vafeiadis, V.: Taming release-acquire consistency. In: Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. pp. 649–662 (2016)
24. Lahav, O., Vafeiadis, V.: Owicki-Gries reasoning for weak memory models. In: Proceedings, Part II, of the 42Nd International Colloquium on Automata, Languages, and Programming - Volume 9135. pp. 311–323 (2015)
25. Litz, H., Cheriton, D., Firoozshahian, A., Azizi, O., Stevenson, J.P.: SI-TM: Reducing transactional memory abort rates through snapshot isolation. SIGPLAN Not. pp. 383–398 (2014)
26. Litz, H., Dias, R.J., Cheriton, D.R.: Efficient correction of anomalies in snapshot isolation transactions. ACM Trans. Archit. Code Optim. **11**(4), 65:1–65:24 (Jan 2015). <https://doi.org/10.1145/2693260>
27. Martin, M., Blundell, C., Lewis, E.: Subtleties of transactional memory atomicity semantics. IEEE Comput. Archit. Lett. **5**(2), 17–17 (2006)
28. Papadimitriou, C.H.: The serializability of concurrent database updates. J. ACM **26**(4), 631–653 (Oct 1979). <https://doi.org/10.1145/322154.322158>, <http://doi.acm.org/10.1145/322154.322158>
29. Peng, D., Dabek, F.: Large-scale incremental processing using distributed transactions and notifications. In: Proceedings of the 9th USENIX Conference on Operating Systems Design and Implementation. pp. 251–264 (2010)
30. Raad, A., Lahav, O., Vafeiadis, V.: On parallel snapshot isolation and release/acquire consistency. In: Proceedings of the 27th European Symposium on Programming (2018), to appear
31. Raad, A., Lahav, O., Vafeiadis, V.: The technical appendix for this paper is available at <https://arxiv.org/abs/1805.06196> (2018)
32. Serrano, D., Patino-Martinez, M., Jimenez-Peris, R., Kemme, B.: Boosting database replication scalability through partial replication and 1-copy-snapshot-isolation. In: Proceedings of the 13th Pacific Rim International Symposium on Dependable Computing. pp. 290–297 (2007)
33. Shavit, N., Touitou, D.: Software transactional memory. In: Proceedings of the Fourteenth Annual ACM Symposium on Principles of Distributed Computing. pp. 204–213 (1995)

34. Sovran, Y., Power, R., Aguilera, M.K., Li, J.: Transactional storage for geo-replicated systems. In: Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles. pp. 385–400 (2011)
35. Vafeiadis, V., Narayan, C.: Relaxed separation logic: A program logic for C11 concurrency. In: Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications. pp. 867–884 (2013)

A MRSW Lock Implementations

We consider two different MRSW library implementations, both satisfying the lock synchronisation guarantees required by [Def. 4](#). As we demonstrate shortly, our first implementation ([Appendix A.1](#)) offers additional synchronisation guarantees by ensuring that any two calls to the library (including those of read locks) synchronise. That is, our first implementation satisfies the (**RSYNC**) axiom on [page 16](#). By contrast, the synchronisation guarantees of our second implementation ([Appendix A.2](#)) are exactly those required by [Def. 4](#), where a library call to a *write* lock synchronises with all other lock library calls. As we discussed earlier in [§2](#), whilst *both* library implementations can be used in our *SI* implementation, *only the weaker* (second) implementation can be used in our *RSI* implementation.

A.1 Fully Synchronising MRSW Lock Implementation

Our first MRSW lock library is implemented in the RA fragment of C11 [\[23\]](#), and is given in [Fig. 7](#). In this implementation, the lock associated with each location x resides at location $x+1$, written $x1$. The state of a lock $x1$ is represented by an integer value. A lock $x1$ may hold either:

- i) value 0, denoting that the lock is free (not held in read or write mode); or
- ii) value 1, denoting that the lock is held (exclusively) in write mode; or
- iii) an even value $2n$ with $n > 0$, denoting that the lock is held in (shared) read mode by n readers; or
- iv) an odd value $2n+1$ with $n > 0$, denoting that the lock is currently being promoted, awaiting the release of n readers.

As such, the implementation of `lock_w x` simply spins until it can atomically update (via `CAS`) the value of $x1$ from zero (free) to one (acquired in write mode). The `CAS(x1, v, v')` denotes the atomic ‘compare-and-set’ operation, where *either* $x1$ currently holds value v in which case it is atomically updated to v' and true is returned; *or* $x1$ currently holds a value other than v in which case it is left unchanged and false is returned. Dually, the implementation of `unlock_w x` simply releases the write lock by atomically assigning $x1$ to zero.

The implementation of `can-promote x` is more involved. As multiple readers may attempt to promote their reader locks simultaneously, promotion is granted on a ‘first-come-first-served’ bases. As such, the implementation of `can-promote x` first reads the value of $x1$. If $x1$ holds an odd value, then another reader is currently promoting $x1$ and thus promotion of $x1$ fails by returning false. On the other hand, if $x1$ holds an even value, then its value is atomically decremented (to an odd value) to signal the intention to promote. The implementation then proceeds by spinning until all other readers have released their locks on x (i.e. $x1 == 1$), at which point true is returned to denote the successful acquisition of x in write mode. Note that once a reader has signalled its intention to promote x (by decrementing $x1$ to an odd value), any other such attempt to promote the lock on x , as well as calls to acquire it in read mode will fail thereafter until such time that x is released by its current promoter.

```

lock_r x  $\triangleq$ 
  start: a := x1;
        if (is-odd a) goto start;
        if (!CAS(x1, a, a+2))
          goto start;

unlock_r x  $\triangleq$  FAA(x1, -2);

lock_w x  $\triangleq$  while (!CAS(x1, 0, 1)) skip;

can-promote x  $\triangleq$ 
  start: a := x1;
        if (is-odd a) return false;
        if (!CAS(x1, a, a-1))
          goto start;
        while (x1 != 1) skip;
        return true;

unlock_w x  $\triangleq$  x1 := 0;

```

Fig. 7: Fully synchronising MRSW lock implementation in the RA fragment of C11

The implementation of `lock_r x` is similar. It first checks whether `x1` is odd (held in write mode or being promoted). If so then the implementation spins until `x1` is even (free or held in read mode), at which point its value is incremented by two (to increase the number of readers by one) using the atomic ‘fetch-and-add’ (FAA) operation, and `x` is successfully acquired in read mode. Dually, the implementation of `unlock_r x` atomically decrements the value of `x1` by two to decrease the reader count by one.

Synchronisation Guarantees In what follows we demonstrate that our implementation satisfies the (**WSYNC**), (**WEX**) and (**RSHARE**) conditions in Def. 4, as well as the stronger (**RSYNC**) condition discussed on page 16.

Observe that a successful acquisition of a writer lock is done via an atomic update operation (when the CAS succeeds). That is, a call to `lock_w x` returns only when the CAS is successful, i.e. when no other thread holds a lock on `x`. Similarly, a call to `unlock_w x` returns after an atomic write to `x1` assigning it to zero. Moreover, once `x` is acquired in write mode by a thread τ , no other thread can acquire it (all other calls to `lock_r x` and `lock_w x` spin until `x` is released by τ). As such, any RA-consistent execution graph of a program P containing a call by thread τ to `lock_w x` followed by its subsequent release via `unlock_w x` includes a trace of the following form, where wl and wu are events of thread τ , wl denotes the update event associated with the successful CAS acquiring the writer lock, and wu denotes the write event associated with its release:

$$\xrightarrow{rf} wl : U(x1, 0, 1) \xrightarrow{po_{x1}|_{imm}} wu : W(x1, 0) \quad (1)$$

Note that if the use of locks in P is well-formed (i.e. the sequence of lock events in each thread following po matches (a prefix of) the regular expression $(RL \cdot RU \mid WL \cdot WU \mid RL \cdot PL \cdot WU)^*$ – see page 15), no other write event on `x1` can happen between wl and wu . This is because i) well-formedness of lock traces ensures that the lock is released by the acquiring thread τ itself, and ii) the lock

on x cannot be acquired after wl and before wu : calls to `lock_w x` fail because $x1$ is non-zero; and calls to `can-promote x` and `lock_r x` fail because $x1 = 1$ holds an odd value. This in turn ensures the mutual exclusion property of writer locks, as required by the (WEX) condition in Def. 4.

Similarly, a call to `can-promote x` returns after the CAS succeeds and subsequently the condition of the while loop amounts to false and no other thread owns a lock on x (in read or write mode). As such, any RA-consistent execution graph of a program P containing a call by thread τ to `can-promote x` followed by its subsequent release via `unlock_w x` includes a trace of the following form, where sp , pl and wu are events of thread τ , sp denotes the update event associated with the successful CAS signalling promotion, pl denotes the atomic read event in the final iteration of the while loop denoting successful promotion, and wu denotes the write event associated with its release:

$$\xrightarrow{\text{rf}} sp : \text{U}(x1, 2i+2, 2i+1) \xrightarrow{\text{rf}} b_1 \xrightarrow{\text{rf}} \dots \xrightarrow{\text{rf}} b_i \xrightarrow{\text{rf}} pl : \text{R}(x1, 1) \xrightarrow{\text{po}_{x1} | \text{imm}} wu : \text{W}(x1, 0) \quad (2)$$

for some i denoting the reader count on x , with $b_1 \dots b_m \in \mathcal{RU}_x$, $rl \xrightarrow{\text{po}_{x1} | \text{imm}} ru$. In other words, once τ has signalled its intention to promote x , no other lock on x can be acquired – calls to `lock_w x` fail as the value of $x1$ is non-zero; calls to `lock_r x` fail as $x1$ holds an odd value. However, existing reader locks must be released before x can be successfully promoted. As with writer locks above, if the use of locks in P is well-formed, no other event on $x1$ can happen between pl and wu . This once again ensures the mutual exclusion property of promoted locks, as required by the (WEX) condition in Def. 4.

Analogously, a call to `lock_r x` returns when the CAS succeeds; and a call to `unlock_r x` returns after the atomic FAA operation. Moreover, once x is acquired in reader mode by a thread τ , no other thread can acquire it in write mode (all other calls to `lock_w x` spin) until it has been released by τ (and potentially other readers). As such, any RA-consistent execution graph of a program P containing a call to `lock_r x` followed by its subsequent release via `unlock_r x` includes a trace of the following form, where rl and ru are events of thread τ , rl denotes the update event associated with the successful CAS acquiring the reader lock, and ru denotes the update event associated with its release via FAA:

$$\xrightarrow{\text{rf}} rl : \text{U}(x1, 2i, 2i+2) \xrightarrow{\text{rf}} \underbrace{a_1 \xrightarrow{\text{rf}} \dots \xrightarrow{\text{rf}} a_n}_{\text{zero or more}} \xrightarrow{\text{rf}} \underbrace{sp : \text{U}(x1, 2k+2, 2k+1)}_{\text{zero or one}} \xrightarrow{\text{rf}} \underbrace{b_1 \xrightarrow{\text{rf}} \dots \xrightarrow{\text{rf}} b_m}_{\text{zero or more}} \xrightarrow{\text{rf}} ru : \text{U}(x1, j+2, j) \quad (3)$$

for some $i, j, k > 0$, where $a_1 \dots a_n \in \mathcal{RL}_x \cup \mathcal{RU}_x$, $b_1 \dots b_m \in \mathcal{RU}_x$, $rl \xrightarrow{\text{po}_{x1} | \text{imm}} ru$. In other words, other threads may also acquire x in read mode (denoted by events of $a_1 \dots a_i$) and their calls may be interleaved between the read lock and unlock of τ . Moreover, another thread may signal to promote its reader lock on x (denoted by sp) in between the acquisition and release of the reader lock on x by τ . Lastly, once the lock has been signalled for promotion, no other reader lock on x can be acquired, though existing reader locks may be released (denoted by

$b_1 \cdots b_m$). Note that if the use of locks in P is well-formed, no thread can acquire x in write mode (as x holds a non-zero value) or successfully promote it before x is released by τ via ru . More concretely, threads may signal their intention to promote (see sp above). However, they cannot successfully promote it before τ has released it in ru as $x1$ holds a value greater than 1 (see pl in (2)). This ensures the (RSHARE) condition in Def. 4.

Observe that given an RA-consistent execution graph involving calls to the above MRSW lock library above, for each location x and its lock at $x1$, the trace of events on $x1$ comprises po and rf edges, as demonstrated by the traces in (1), (2) and (3). In other words, any two events on $x1$ are related by $(po \cup rf)_{loc}^+$, i.e. $(po \cup rf)_{loc}^+$ is total for each $x1$. Moreover, from the RA-consistency of our execution we know that $(po \cup rf)_{loc}^+$ is acyclic. As such, since $(po \cup rf)_{loc}^+$ is transitively closed, we know that it is a strict total order. For each location x , we thus define lo_x as the strict total order given by $lo_x \triangleq (po_{x1} \cup rf_{x1})^+$. As such, our implementation satisfies both the (WSYNC) condition in Def. 4 and the stronger (RSYNC) condition discussed on page 16.

A.2 Write Synchronising MRSW Lock Implementation

Our second MRSW lock implementation is similarly implemented in the RA fragment of C11 [23] and is given in Fig. 8. In this implementation, each lock x is represented as an *ordered* map at location $x+1$, written $x1$. The map at $x1$ contains one entry per thread as follows. For each thread with identifier τ , the $x1[\tau]$ map entry records the current locking privileges of τ on x . More concretely, when $x1[\tau] = 0$, then τ does not hold the x lock; when $x1[\tau] = 2$, then τ holds x in read mode; and when $x1[\tau] = 1$; then *some* thread (either τ or another thread) either holds x in write mode, or it is in the process of acquiring x in write mode. The x lock is held in write mode only when all entries in $x1$ are mapped to one. As we describe shortly, for thread τ to acquire x in write mode, it must inspect each entry in $x1$ (in order), wait for it be free (zero) and then set it to one. As we discuss shortly, this in-order inspection of entries allows us to avoid deadlocks. In our implementation, we assume that the thread identifier can be obtained by calling `getTID`. We identify the top-most thread by $\tau = 0$; as such, the entry of top-most thread in each map is ordered before all other threads.

We proceed with a more detailed explanation of our implementation after introducing our map notation. We write $\mathbf{1}$ to denote a map where all entries have value 1; similarly, we write $\mathbf{0}$ to denote a map where all entries have value 0. Lastly, we write $S \subseteq x1$, to denote that the values held in map $x1$ are a superset of S . The lock map $x1$ associated with location x can be in one of the following states:

- $x1 = \mathbf{0}$ when x is free;
- $x1 = \mathbf{1}$ when x is held in write mode;
- $\{2\} \subseteq x1$ when x is held in read mode (by those threads τ where $x1[\tau] = 2$).

When thread τ calls `lock_r` x , it simply spins until the lock is free ($x1 = \mathbf{0}$ and thus $x1[\tau] = 0$), at which point it acquires it in read mode by setting $x1[\tau]$ to two. Dually, when τ calls `unlock_r` x it simply sets $x1[\tau]$ to zero.

```

lock_r x  $\triangleq$ 
  t:=getTID;
  while (!CAS(xl[t], 0, 2)) skip;

unlock_r x  $\triangleq$  t:=getTID; xl[t]:=0

lock_w x  $\triangleq$ 
  for (t  $\in$  dom(xl))
    while (!CAS(xl[t], 0, 1)) skip;

unlock_w x  $\triangleq$ 
  for (t  $\in$  dom(xl)) xl[t]:=0;

can-promote x  $\triangleq$ 
  t:=getTID;
  if (t==0) xl[t]:=1;
  else {
    retry: a:=xl[0];
    if (a==1) return false;
    if (!CAS(xl[0], 0, 1))
      goto retry;
  } xl[t]:=1;
  for (i  $\in$  dom(xl) && i  $\notin$  {0,t})
    while (!CAS(xl[i], 0, 1)) skip;
  return true;

```

Fig. 8: Write synchronising MRSW lock implementation in the RA fragment of C11

Analogously, when τ calls `lock_w x`, it traverses the `xl` map in order, spinning on each entry until it is free (0) and subsequently acquiring it (by setting it to 1). Conversely, when τ calls `unlock_w x`, it releases `x` by traversing `xl` in order and setting each entry to one.

To understand the implementation of lock promotion, first consider the case where `can-promote x` is called by $\tau \neq 0$, i.e. a thread other than the top-most thread. The implementation of `can-promote x` then inspects the *first* entry in the map (`xl[0]`), i.e. that of the top-most thread. If `xl[0] = 1`, then `x` is currently being acquired by another thread; the promotion thus fails and `false` is returned. If on the other hand `xl[0] \neq 1` (i.e. `xl[0] = 0` or `xl[0] = 2`), the implementation spins until it is zero and atomically updates it to one, signalling its intention to promote `x`. This pre-empts the promotion of `x` by other threads: any such attempt would fail as now `xl[0] = 1`. The implementation then sets its own entry (`xl[τ]`) to one, traverses the map in order, and spins on each entry until they too can be set to one. At this point the lock is successfully promoted and `true` is returned. Note that it is safe for τ to update its own entry `xl[τ]` to one: at this point in execution no thread holds the writer lock on `x`, no thread can promote its lock on `x`, and those threads with a reader lock on `x` never access the `xl[τ]` entry – the read lock calls of another thread τ' solely accesses `xl[τ']`.

Let us now consider the case when the top-most thread with $\tau = 0$ calls `can-promote x`. Since prior to a `can-promote x` call τ owns a reader lock on `x`, i.e. `xl[τ] = 2`, no other thread can promote its `x` lock. As such, τ successfully sets `xl[τ]` to one, signalling its intention to promote `x`. In other words, the promotion is skewed in favour of the top-most thread: if a thread races against the top-most thread to promote `x`, the top-most thread always wins. With the exception of the top-most thread, promotion is done on a ‘first-come-first-served’

basis. The rest of the implementation is then carried out as before: the map $\mathbf{x1}$ is traversed in turn and each entry is set to one.

Synchronisation Guarantees In what follows we demonstrate that our implementation satisfies the (**WSYNC**), (**WEX**) and (**RSHARE**) conditions in [Def. 4](#), while it does not satisfy the stronger (**RSYNC**) condition discussed on [page 16](#).

Observe that a successful acquisition of a writer lock is done via several atomic update operations (via successful **CAS** operations in the for loop). That is, a call to `lock_w x` returns only once the **CAS** on all $\mathbf{x1}$ entries succeeds, i.e. when no other thread holds a lock on \mathbf{x} . Similarly, a call to `unlock_w x` returns after atomic writes on each entry in $\mathbf{x1}$, assigning them to zero. Moreover, once \mathbf{x} is acquired in write mode by a thread τ , no other thread can acquire it: all other calls to `lock_r x` spin until the relevant $\mathbf{x1}$ entry is set to zero by τ ; all calls to `lock_w x` spin until all entries in $\mathbf{x1}$ are set to zero by τ . As such, any RA-consistent execution graph of a program P containing a call by thread τ to `lock_w x` or `can-promote x`, followed by its subsequent release via `unlock_w x` includes a trace of the following form, where the domain of the $\mathbf{x1}$ map is $0 \cdots n$ for some n , the $wl_0, \dots, wl_n, wu_0, \dots, wu_n$ are events of thread τ , each wl_i denotes the update event associated with the successful **CAS** (in `lock_w x`) on the $\mathbf{x1}[i]$ entry, and each wu_i denotes the write event setting the $\mathbf{x1}[i]$ entry to zero (in `unlock_w x`):

$$\begin{array}{c}
\stackrel{rf}{\rightarrow} wl_0 : \mathbf{U}(\mathbf{x1}[0], 0, 1) \\
\downarrow \text{po} \\
\vdots \\
\downarrow \text{po} \\
\stackrel{rf}{\rightarrow} wl_n : \mathbf{U}(\mathbf{x1}[n], 0, 1) \xrightarrow{\text{po}} wu_0 : \mathbf{W}(\mathbf{x1}[0], 0) \xrightarrow{\text{po}} \dots \xrightarrow{\text{po}} wu_n : \mathbf{W}(\mathbf{x1}[0], 0)
\end{array} \tag{4}$$

As before, note that if the use of locks in P is well-formed (see [page 15](#)), no other write event on $\mathbf{x1}[i]$ can happen between wl_i and wu_i . This is because no thread can acquire the write lock on $\mathbf{x1}$ after wl_0 and before wu_n ; and thread τ_i cannot acquire the read lock on $\mathbf{x1}$ after wl_i and before wu_i . This in turn ensures the mutual exclusion property of writer locks, as required by the (**WEX**) condition in [Def. 4](#).

Analogously, a call to `lock_r x` returns when the **CAS** succeeds; and a call to `unlock_r x` returns after the atomic write operation setting $\mathbf{x1}[\tau]$ to zero. As such, any RA-consistent execution graph of a program P containing a call to `lock_r x` by thread τ , followed by its subsequent release via `unlock_r x` includes a trace of the following form, where rl and ru are events of thread τ , rl denotes the update event associated with the successful **CAS** acquiring the reader lock, and ru denotes the write event associated with its release by assigning it to zero:

$$\stackrel{rf}{\rightarrow} rl : \mathbf{U}(\mathbf{x1}[\tau], 0, 2) \xrightarrow{\text{po}} ru : \mathbf{W}(\mathbf{x1}[\tau], 0) \tag{5}$$

Note that each read lock call on \mathbf{x} by thread τ , accesses $\mathbf{x1}[\tau]$ alone and no other entry in $\mathbf{x1}$. As such, two threads may simultaneously acquire the reader

lock on \mathbf{x} , ensuring the (**RSHARE**) condition in [Def. 4](#). Moreover, as read lock calls by two distinct threads τ, τ' access *disjoint* memory locations ($\mathbf{x1}[\tau]$ and $\mathbf{x1}[\tau']$), they never synchronise. That is, the MRSW lock implementation in [Fig. 8](#) does not satisfy the (**RSYNC**) axiom on page 16.

Lastly, we demonstrate that our implementation satisfies the (**WSYNC**) condition in [Def. 4](#), when accessed by an arbitrary (finite) number of threads n . Observe that given an RA-consistent execution graph involving calls to the MRSW lock library above, for each location \mathbf{x} , its lock at $\mathbf{x1}$, and each thread τ , the trace of events on $\mathbf{x1}[\tau]$ comprises **po** and **rf** edges, as demonstrated by the traces in [\(4\)](#) and [\(5\)](#). In other words, any two event on $\mathbf{x1}[\tau]$ are related by $(\mathbf{po} \cup \mathbf{rf})_{loc}^+$, i.e. $(\mathbf{po} \cup \mathbf{rf})_{loc}^+$ is total for each $\mathbf{x1}[\tau]$. We thus use this total order to determine the **lo** (synchronisation order) between two lock events, where at least one of them is a write lock event. Let us now pick two distinct lock events, $w, l \in \mathcal{L}_{\mathbf{x}}$, where at least one of them w is a write lock event, i.e. $w \in \mathcal{WL}_{\mathbf{x}} \cup \mathcal{PL}_{\mathbf{x}} \cup \mathcal{WU}_{\mathbf{x}}$. Either 1) w, l are events of the same thread; or 2) w and l are events of distinct threads τ, τ' , respectively. In the first case the two events are related by **po** one way or another, and as $\mathbf{po} \subseteq \mathbf{hb}$, we know that the two events synchronise. In the second case, there are two additional cases to consider: either i) $l \in \mathcal{WL}_{\mathbf{x}} \cup \mathcal{PL}_{\mathbf{x}} \cup \mathcal{WU}_{\mathbf{x}}$; or $l \in \mathcal{RL}_{\mathbf{x}} \cup \mathcal{RU}_{\mathbf{x}}$. In case (i), as we discussed above we know that each call amounts to a trace akin to that in [\(4\)](#). That is, we know the trace of τ contains $wl : \mathbf{U}(\mathbf{x1}[0], 0, 1) \xrightarrow{\mathbf{po}} wu : \mathbf{W}(\mathbf{x1}[0], 0)$, and the trace of τ' contains $wl' : \mathbf{U}(\mathbf{x1}[0], 0, 1) \xrightarrow{\mathbf{po}} wu' : \mathbf{W}(\mathbf{x1}[0], 0)$. Moreover, as we discussed above, we know that no other write event on $\mathbf{x1}[0]$ can happen between wl and wu , and between wl' and wu' . As $(\mathbf{po} \cup \mathbf{rf})_{loc}^+$ is total for each $\mathbf{x1}[\tau]$, we then know that we either have $wl \xrightarrow{\mathbf{po}} wu \xrightarrow{(\mathbf{po} \cup \mathbf{rf})_{loc}^+} wl' \xrightarrow{\mathbf{po}} wu'$, or we have $wl' \xrightarrow{\mathbf{po}} wu' \xrightarrow{(\mathbf{po} \cup \mathbf{rf})_{loc}^+} wl \xrightarrow{\mathbf{po}} wu$. As such, in both cases we know that the two events synchronise.

In case (ii), we know that the write call results in a trace akin to that in [\(4\)](#), while the read call results in a trace similar to that in [\(5\)](#). That is, we know the trace of τ contains $wl : \mathbf{U}(\mathbf{x1}[\tau'], 0, 1) \xrightarrow{\mathbf{po}} wu : \mathbf{W}(\mathbf{x1}[\tau'], 0)$, and the trace of τ' contains $rl : \mathbf{U}(\mathbf{x1}[\tau'], 0, 2) \xrightarrow{\mathbf{po}} ru : \mathbf{W}(\mathbf{x1}[\tau'], 0)$. Moreover, as we discussed above, we know that no other write event on $\mathbf{x1}[\tau']$ can happen between wl and wu . As $(\mathbf{po} \cup \mathbf{rf})_{loc}^+$ is total for $\mathbf{x1}[\tau']$, we then know that we either have $wl \xrightarrow{\mathbf{po}} wu \xrightarrow{(\mathbf{po} \cup \mathbf{rf})_{loc}^+} rl \xrightarrow{\mathbf{po}} ru$, or we have $rl \xrightarrow{\mathbf{po}} ru \xrightarrow{(\mathbf{po} \cup \mathbf{rf})_{loc}^+} wl \xrightarrow{\mathbf{po}} wu$. As such, in both cases we know that the two events synchronise.

B Auxiliary Results

Proposition 2 (SI-consistency). *An execution graph G is SI-consistent if and only if INT holds and the ‘SI-happens-before’ relation $\text{si-hb} \triangleq (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T \cup \text{si-rb})^+$ is irreflexive, where $\text{si-rb} \triangleq [\mathcal{R}_E]; \text{rb}_T; [\mathcal{W}]$ and $\mathcal{R}_E \triangleq \text{codom}(\text{rf}_E)$.*

Proof. Pick an arbitrary execution graph G . We are then required to show:

$$\text{acyclic}((\text{po}_T \cup \text{rf}_T \cup \text{mo}_T); \text{rb}_T^?) \iff \text{irreflexive}(\text{si-hb})$$

The \Rightarrow direction

We proceed by contradiction. Assume that $\text{acyclic}((\text{po}_T \cup \text{rf}_T \cup \text{mo}_T); \text{rb}_T^?)$ and $\neg \text{irreflexive}(\text{si-hb})$ both hold. We then know there exists e such that $(e, e) \in \text{si-hb} = (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T \cup \text{si-rb})^+$. Note that $\text{si-rb}; \text{si-rb} = \emptyset$ because $\mathcal{R}_E \cap \mathcal{W} = \emptyset$. Hence, the si-hb cycle cannot have adjacent si-rb edges, which means that we have a cycle in $((\text{po}_T \cup \text{rf}_T \cup \text{mo}_T); \text{si-rb}^?)^+$, which contradicts our first assumption as $\text{si-rb} \subseteq \text{rb}_T$.

The \Leftarrow direction

We proceed by contradiction. Let us assume $\text{irreflexive}(\text{si-hb})$ and $\neg \text{acyclic}((\text{po}_T \cup \text{rf}_T \cup \text{mo}_T); \text{rb}_T^?)$. We then know there exists e such that $(e, e) \in ((\text{po}_T \cup \text{rf}_T \cup \text{mo}_T); \text{rb}_T^?)^+$, i.e. there exists e such that $(e, e) \in (\text{rb}_T^?; (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T))^+$. There are now two cases to consider: either 1) $(e, e) \in (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T)^+$; or 2) $\exists a, d. (e, a) \in (\text{rb}_T^?; (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T))^*$, $(a, d) \in \text{rb}_T; (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T)$ and $(d, e) \in (\text{rb}_T^?; (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T))^*$, i.e. $\exists a. (a, a) \in \text{rb}_T; (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T); (\text{rb}_T^?; (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T))^*$. In case (1) from the definition of si-hb we then have $(e, e) \in \text{si-hb}$, contradicting the assumption that $\text{irreflexive}(\text{si-hb})$ holds.

In case (2) we then know there exists b, c such that $[b]_{\text{st}} \neq [c]_{\text{st}}$, $(a, b) \in \text{rb}_T$, $(b, c) \in \text{po}_T \cup \text{rf}_T \cup \text{mo}_T$ and $(c, a) \in (\text{rb}_T^?; (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T))^*$. From the definitions of po_T , rf_T and mo_T we then have $[b]_{\text{st}} \times [c]_{\text{st}} \subseteq \text{po}_T \cup \text{rf}_T \cup \text{mo}_T$ and thus from the definition of si-hb we have $[b]_{\text{st}} \times [c]_{\text{st}} \subseteq \text{si-hb}$. As such, from Lemma 1 below we have $[b]_{\text{st}} \times [a]_{\text{st}} \subseteq \text{si-hb}$. On the other hand, from $(a, b) \in \text{rb}_T$ we know there exist r, w, w' such that $r \in [a]_{\text{st}}$, $w \in [b]_{\text{st}}$, $[a]_{\text{st}} \neq [b]_{\text{st}}$, $(r, w) \in \text{rb}$, $(w', r) \in \text{rf}$ and $(w', w) \in \text{mo}$. There are now two cases to consider: i) $w' \in [a]_{\text{st}}$; or ii) $w' \notin [a]_{\text{st}}$.

In case (2.i) we then have $[a]_{\text{st}} \times [b]_{\text{st}} \subseteq \text{mo}_T \subseteq \text{si-hb}$. As such, we have $(a, b) \in \text{si-hb}$. Since we have also established $[b]_{\text{st}} \times [a]_{\text{st}} \subseteq \text{si-hb}$, we have $(b, a) \in \text{si-hb}$. By transitivity, we have $(a, a) \in \text{si-hb}$, contradicting the assumption that $\text{irreflexive}(\text{si-hb})$ holds.

In case (2.ii) we then know $r \in \mathcal{R}_E$. As such we have $(r, w) \in \text{si-rb} \subseteq \text{si-hb}$. Since we have also established $[b]_{\text{st}} \times [a]_{\text{st}} \subseteq \text{si-hb}$, we have $(w, r) \in \text{si-hb}$. By transitivity, we have $(r, r) \in \text{si-hb}$, contradicting the assumption that $\text{irreflexive}(\text{si-hb})$ holds.

Theorem 5. For all executions G , if $G.\mathcal{NT} = \emptyset$, then:

$$G \text{ is SI-consistent} \iff G \text{ is RSI-consistent}$$

Proof (the \Leftarrow direction). Pick an arbitrary G such that $G.\mathcal{NT} = \emptyset$ and G is RSI-consistent. Let us proceed by contradiction and assume G is not SI-consistent. That is, there exists a such that $(a, a) \in G.\text{si-hb}$. From the definition of **si-hb** we then have $G.\text{si-hb} \subseteq G.\text{rsi-hb}$. As such, we have $(a, a) \in G.\text{rsi-hb}$, contradicting the assumption that G is RSI-consistent.

Proof (the \Rightarrow direction). Pick an arbitrary G such that $G.\mathcal{NT} = \emptyset$ and G is SI-consistent. As $G.\mathcal{NT} = \emptyset$, we then have $G.\text{rsi-hb} = (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T \cup \text{si-rb} \cup A)^+$ with $A \triangleq G.([\mathcal{W}]; \text{po}_I; [\mathcal{W}])$. That is, $G.\text{rsi-hb} = (G.\text{si-hb} \cup A)^+$. In what follows we demonstrate:

$$A; \text{si-hb} \subseteq \text{si-hb} \quad \text{and} \quad \text{si-hb}; A \subseteq \text{si-hb} \tag{6}$$

As such, since **si-hb** is transitively closed, we have $G.\text{rsi-hb} = G.\text{si-hb} \cup A$.

Moreover, we have $G.\text{mo} \subseteq G.\text{mo}_T \cup G.\text{mo}_I$. As such, since from (INT) we have $G.\text{mo}_I \subseteq G.\text{po}_I$, we have $G.\text{mo}_I \subseteq A$ and thus $G.\text{mo} \subseteq G.\text{rsi-hb}_{loc}$.

Let $B \triangleq G.\text{rb} \cap G.\text{po}_I$. We also have $G.\text{rb} \subseteq G.\text{si-rb} \cup G.\text{rb}_I$. As such, since from (INT) we have $G.\text{rb}_I \subseteq G.\text{po}_I$, we have $G.\text{rb}_I \subseteq B$ and thus $G.\text{rb} \subseteq G.\text{rsi-hb}_{loc} \cup B$. Consequently, we have $G.(\text{rsi-hb}_{loc} \cup \text{mo} \cup \text{rb}) = G.(\text{rsi-hb}_{loc} \cup B) = G.(\text{si-hb}_{loc} \cup A_{loc} \cup B)$. Therefore, $G.(\text{rsi-hb}_{loc} \cup \text{mo} \cup \text{rb})^+ = G.(\text{si-hb}_{loc} \cup A_{loc} \cup B)^+$.

Observe that $A; B = \emptyset$. In what follows we demonstrate that:

$$B; A \subseteq B \tag{7}$$

$$\text{si-hb}; B \subseteq \text{si-hb} \quad \text{and} \quad B; \text{si-hb} \subseteq \text{si-hb} \tag{8}$$

Note that since G is SI-consistent we know that A is irreflexive. As such, to show that $G.(\text{rsi-hb}_{loc} \cup \text{mo} \cup \text{rb})^+ = G.(\text{si-hb}_{loc} \cup A_{loc} \cup B)^+$ is irreflexive, from (6), (7), (8), and since B is irreflexive, it suffices to show that $G.\text{si-hb}$ is irreflexive, which follows immediately from the SI-consistency of G .

TS. (6)

To show $A; \text{si-hb} \subseteq \text{si-hb}$, pick an arbitrary $(a, b) \in A; \text{si-hb}$. We then know there exists c such that $a, c \in \mathcal{W}$, $(a, c) \in A$ and $(c, b) \in \text{si-hb}$. From the definition of **si-hb** and since $c \in \mathcal{W}$, we know there exists d such that $(c, d) \in \text{po}_T \cup \text{rf}_T \cup \text{mo}_T$ and $(d, b) \in \text{si-hb}^*$. As $(a, c) \in A \subseteq \text{st}$, we thus know that $(a, d) \in \text{po}_T \cup \text{rf}_T \cup \text{mo}_T \subseteq \text{si-hb}$. Consequently, since $(d, b) \in \text{si-hb}^*$ and **si-hb** is transitively closed, we have $(a, b) \in \text{si-hb}$, as required.

To show $\text{si-hb}; A \subseteq \text{si-hb}$, pick an arbitrary $(a, b) \in \text{si-hb}; A$. We then know there exists c such that $b, c \in \mathcal{W}$, $(a, c) \in \text{si-hb}$ and $(c, b) \in A$. From the definition of **si-hb**, we know there exists d such that $(a, d) \in \text{si-hb}^*$ and $(d, c) \in \text{po}_T \cup \text{rf}_T \cup \text{mo}_T \cup \text{si-rb}$. As $(c, b) \in A \subseteq \text{st}$ and since $b, c \in \mathcal{W}$ we thus know that $(d, b) \in \text{po}_T \cup \text{rf}_T \cup \text{mo}_T \cup \text{si-rb} \subseteq \text{si-hb}$. Consequently, since $(a, d) \in \text{si-hb}^*$ and **si-hb** is transitively closed, we have $(a, b) \in \text{si-hb}$, as required.

TS. (7)

Pick an arbitrary a, b such that $(a, b) \in B; A$. We then know there exists c such that $b, c \in \mathcal{W}$, $(a, c) \in B$, $(c, b) \in A$ and $\text{tx}(a) = \text{tx}(b) = \text{tx}(c)$. As $b, c \in \mathcal{W}$, we know either $(b, c) \in \text{mo}_I$ or $(c, b) \in \text{mo}_I$. However, as G is SI-consistent, from (INT) we know that $\text{mo}_I \subseteq \text{po}_I$ and thus we have $(c, b) \in \text{mo}_I \cap \text{po}_I$. As such, since $(a, c) \in \text{rb} \cap \text{po}_I$, $(c, b) \in \text{mo}_I \cap \text{po}_I$ and po_I is transitively closed, from the definition of rb we have $(a, b) \in \text{rb} \cap \text{po}_I = B$, as required.

TS. (8)

To show $B; \text{si-hb} \subseteq \text{si-hb}$, pick an arbitrary $(a, b) \in B; \text{si-hb}$. We then know there exists c such that $a \in \mathcal{R}$, $c \in \mathcal{W}$, $(a, c) \in B$ and $(c, b) \in \text{si-hb}$. From the definition of si-hb and since $c \in \mathcal{W}$, we know there exists d such that $(c, d) \in \text{po}_T \cup \text{rf}_T \cup \text{mo}_T$ and $(d, b) \in \text{si-hb}^*$. As $(a, c) \in B \subseteq \text{st}$, we thus know that $(a, d) \in \text{po}_T \cup \text{rf}_T \cup \text{mo}_T \subseteq \text{si-hb}$. Consequently, since $(d, b) \in \text{si-hb}^*$ and si-hb is transitively closed, we have $(a, b) \in \text{si-hb}$, as required.

To show $\text{si-hb}; B \subseteq \text{si-hb}$, pick an arbitrary $(a, b) \in \text{si-hb}; B$. We then know there exists c such that $c \in \mathcal{R}$, $b \in \mathcal{W}$, $(a, c) \in \text{si-hb}$ and $(c, b) \in B$. From the definition of si-hb , and since $c \in \mathcal{R}$, we know there exists d such that $(a, d) \in \text{si-hb}^*$ and $(d, c) \in \text{po}_T \cup \text{rf}_T \cup \text{mo}_T$. As $(c, b) \in B \subseteq \text{st}$, we thus know that $(d, b) \in \text{po}_T \cup \text{rf}_T \cup \text{mo}_T \subseteq \text{si-hb}$. Consequently, since $(a, d) \in \text{si-hb}^*$ and si-hb is transitively closed, we have $(a, b) \in \text{si-hb}$, as required.

Theorem 6 (Monotonicity). *Let G_T be an RSI execution graph obtained from an RSI execution graph G by wrapping some non-transactional events inside a new transaction. If G_T is RSI-consistent, then so is G .*

Proof. First, we show that $[\mathcal{W}]; G.\text{rsi-hb} \subseteq G_T.\text{po}_I \cup G_T.\text{rsi-hb}$. Let A denote the events of the new transactions in G_T . Let $\langle a, b \rangle \in [\mathcal{W}]; G.\text{rsi-hb}$, and consider a

$$(\text{po} \setminus \text{po}_I) \cup [\mathcal{W}]; \text{po}_I; [\mathcal{W}] \cup \text{rf}; [\mathcal{NT}] \cup [\mathcal{NT}]; \text{rf}; \text{st} \cup \text{rf}_T \cup (\text{mo}; \text{rf})_T \cup \text{mo}_T \cup \text{si-rb}$$

path from a to b of minimal length (all relations are in G). Note that the only possible edges on this path that do not appear in a corresponding relation in G_T are $[A]; \text{rf}; [A]$ and $[A]; \text{po}; [A] \setminus (\mathcal{W} \times \mathcal{W})$ edges. Suppose first that $[A]; \text{rf}; [A]$ is used on this path. Immediately before or after $[A]; \text{rf}; [A]$, we can only have $\text{po} \setminus \text{po}_I$, but then the two edges can be replaced by a shorter path that uses only $\text{po} \setminus \text{po}_I$. Thus, in this case it follows that $\langle a, b \rangle \in [A]; \text{rf}; [A] \subseteq G_T.\text{po}_I$. Next, suppose that $[A]; \text{po}; [A] \setminus (\mathcal{W} \times \mathcal{W})$ is used on this path (the minimality ensures such edge is used only once). First, if this edge is the first edge on the path, then (again) it is the only edge on the path (immediately after $[A]; \text{po}; [A]$, we can only have $\text{rf}; [\mathcal{NT} \setminus A]$ or $[\mathcal{NT}]; \text{rf}; \text{st}$, but both start with a write). Hence, in this case we have $\langle a, b \rangle \in [A]; \text{po}; [A] \subseteq G_T.\text{po}_I$. Second, consider the case that the $[A]; \text{po}; [A]$ edge is not the first on the path. Immediately before $[A]; \text{po}; [A]$, we can only have a $[E \setminus A]; \text{rf}$ edge. Then, in G_T , the two edges can be joined into either $[\mathcal{NT}]; G_T.\text{rf}; G_T.\text{st}$ or $G_T.\text{rf}_T$. Hence, we obtain $\langle a, b \rangle \in G_T.\text{rsi-hb}$.

Now, suppose that G is not RSI-consistent. If **INT** does not hold, then G_T is also not RSI-consistent and we are done. Otherwise, $G.\text{rsi-hb}_{loc} \cup G.\text{mo} \cup G.\text{rb}$ is cyclic. Since $G.\text{mo}$ is total on writes to each location, it follows that $\langle a, a \rangle \in G.\text{rsi-hb}; (G.\text{mo} \cup G.\text{rb})$ for some $a \in G.\mathcal{W}$. Since $G.\text{mo} = G_T.\text{mo}$, $G.\text{rb} = G_T.\text{rb}$, our claim above entails that $\langle a, a \rangle \in (G_T.\text{po}_I \cup G_T.\text{rsi-hb}); (G_T.\text{mo} \cup G_T.\text{rb})$. Now, $\langle a, a \rangle \in G_T.\text{po}_I; (G_T.\text{mo} \cup G_T.\text{rb})$ implies that **INT** does not hold for G_T , while $\langle a, a \rangle \in G_T.\text{rsi-hb}; (G_T.\text{mo} \cup G_T.\text{rb})$ implies that $G_T.\text{rsi-hb}_{loc} \cup G_T.\text{mo} \cup G_T.\text{rb}$ is cyclic. In both cases, G_T is not RSI-consistent.

Lemma 1. For all a, b, c :

$$[a]_{\text{st}} \neq [b]_{\text{st}} \wedge [a]_{\text{st}} \times [b]_{\text{st}} \subseteq \text{si-hb} \wedge (b, c) \in (\text{rb}_T^?; (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T))^* \Rightarrow [a]_{\text{st}} \times [c]_{\text{st}} \subseteq \text{si-hb}$$

Proof. Note that $(\text{rb}_T^?; (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T))^* = \bigcup_{n \in \mathbb{N}} S_n$, where $S_0 \triangleq \text{id}$ and for all $n \in \mathbb{N}$ we have $S_{n+1} \triangleq (\text{rb}_T^?; (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T)); S_n$. We thus demonstrate instead:

$$\forall a, b, c. \forall n \in \mathbb{N}. [a]_{\text{st}} \neq [b]_{\text{st}} \wedge [a]_{\text{st}} \times [b]_{\text{st}} \subseteq \text{si-hb} \wedge (b, c) \in S_n \Rightarrow [a]_{\text{st}} \times [c]_{\text{st}} \subseteq \text{si-hb}$$

We proceed by induction on n .

Base case $n = 0$

Follows immediately from the assumptions of the lemma and the definition of S_0 .

Inductive case $n = m+1$

$$\forall a, b, c. \forall i \leq m.$$

$$[a]_{\text{st}} \neq [b]_{\text{st}} \wedge [a]_{\text{st}} \times [b]_{\text{st}} \subseteq \text{si-hb} \wedge (b, c) \in S_i \Rightarrow [a]_{\text{st}} \times [c]_{\text{st}} \subseteq \text{si-hb} \quad (\text{I.H.})$$

Pick arbitrary a, b, c such that $[a]_{\text{st}} \neq [b]_{\text{st}}$, $[a]_{\text{st}} \times [b]_{\text{st}} \subseteq \text{si-hb}$, and $(b, c) \in S_n$. From the definition of S_n we then know there exists d such that $(b, d) \in \text{rb}_T^?; (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T)$ and $(d, c) \in S_m$. There are now two cases to consider: 1) $(b, d) \in \text{po}_T \cup \text{rf}_T \cup \text{mo}_T$; or 2) $(b, d) \in \text{rb}_T; (\text{po}_T \cup \text{rf}_T \cup \text{mo}_T)$.

In case (1) from the definitions of po_T , rf_T and mo_T we have $[b]_{\text{st}} \times [d]_{\text{st}} \subseteq \text{po}_T \cup \text{rf}_T \cup \text{mo}_T$ and thus from the definition of si-hb we have $[b]_{\text{st}} \times [d]_{\text{st}} \subseteq \text{si-hb}$. Since we also have $[a]_{\text{st}} \times [b]_{\text{st}} \subseteq \text{si-hb}$ and si-hb is transitively closed, we have $[a]_{\text{st}} \times [d]_{\text{st}} \subseteq \text{si-hb}$. As $(d, c) \in S_m$, from (I.H.) we then have $[a]_{\text{st}} \times [c]_{\text{st}} \subseteq \text{si-hb}$, as required.

In case (2) we then know there exists e such that $(b, e) \in \text{rb}_T$ and $(e, d) \in \text{po}_T \cup \text{rf}_T \cup \text{mo}_T$. From the definitions of po_T , rf_T and mo_T we then have $[e]_{\text{st}} \times [d]_{\text{st}} \subseteq \text{po}_T \cup \text{rf}_T \cup \text{mo}_T$ and thus from the definition of si-hb we have $[e]_{\text{st}} \times [d]_{\text{st}} \subseteq \text{si-hb}$. On the other hand, from $(b, e) \in \text{rb}_T$ we know there exist r, w, w' such that $r \in [b]_{\text{st}}$, $w \in [e]_{\text{st}}$, $[b]_{\text{st}} \neq [e]_{\text{st}}$, $[d]_{\text{st}} \neq [e]_{\text{st}}$, $(r, w) \in \text{rb}$, $(w', r) \in \text{rf}$ and $(w', w) \in \text{mo}$. There are now two cases to consider: i) $w' \in [b]_{\text{st}}$; or ii) $w' \notin [b]_{\text{st}}$.

In case (2.i) we then have $[b]_{\text{st}} \times [e]_{\text{st}} \subseteq \text{mo}_T \subseteq \text{si-hb}$. Since we also have $[a]_{\text{st}} \times [b]_{\text{st}} \subseteq \text{si-hb}$ and $[e]_{\text{st}} \times [d]_{\text{st}} \subseteq \text{si-hb}$, from the transitivity of si-hb we have $[a]_{\text{st}} \times [d]_{\text{st}} \subseteq \text{si-hb}$. As $(d, c) \in S_m$, from (I.H.) we then have $[a]_{\text{st}} \times [c]_{\text{st}} \subseteq \text{si-hb}$, as required.

In case (2.ii) we then know $r \in \mathcal{R}_E$. As such we have $(r, w) \in \text{si-rb} \subseteq \text{si-hb}$. As $[a]_{\text{st}} \times [b]_{\text{st}} \subseteq \text{si-hb}$ and $r \in [b]_{\text{st}}$ we then have $[a]_{\text{st}} \times \{w\} \subseteq \text{si-hb}$. Similarly, as $[e]_{\text{st}} \times [d]_{\text{st}} \subseteq \text{si-hb}$ and $w \in [e]_{\text{st}}$ we then have $\{w\} \times [d]_{\text{st}} \subseteq \text{si-hb}$. As such, from the transitivity of si-hb we have $[a]_{\text{st}} \times [d]_{\text{st}} \subseteq \text{si-hb}$. As $(d, c) \in S_m$, from (I.H.) we then have $[a]_{\text{st}} \times [c]_{\text{st}} \subseteq \text{si-hb}$, as required.

Lemma 2 (Lock ordering). Given an RA-consistent execution graph $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ of the SI or RSI implementations in Fig. 3, for all $w, w' \in$

$\mathcal{WL}_x, pl, pl' \in \mathcal{PL}_x, wu, wu' \in \mathcal{WU}_x, rl, rl' \in \mathcal{RL}_x,$ and $ru \in \mathcal{RU}_x$:

$$\begin{aligned} (wl, wu), (wl', wu') \in \text{po}_x|_{\text{imm}} &\Rightarrow (wu, wl') \in \text{lo}_x \vee (wu', wl) \in \text{lo}_x \\ (wl, wu), (rl, pl, wu') \in \text{po}_x|_{\text{imm}} &\Rightarrow (wu, rl) \in \text{lo}_x \vee (wu', wl) \in \text{lo}_x \\ (rl, pl, wu), (rl', pl', wu') \in \text{po}_x|_{\text{imm}} &\Rightarrow (wu, rl') \in \text{lo}_x \vee (wu', rl) \in \text{lo}_x \end{aligned} \quad (\text{WWSYNC})$$

$$\begin{aligned} (wl, wu), (rl, ru) \in \text{po}_x|_{\text{imm}} &\Rightarrow (wu, rl) \in \text{lo}_x \vee (ru, wl) \in \text{lo}_x \\ (rl, pl, wu), (rl', ru) \in \text{po}_x|_{\text{imm}} &\Rightarrow (wu, rl') \in \text{lo}_x \vee (ru, pl) \in \text{lo}_x \end{aligned} \quad (\text{RWSYNC})$$

where given a relation r we write $(a, b, c) \in r$ as a shorthand for $(a, b), (b, c) \in r$.

Proof (Proof (WWSYNC)). Pick an arbitrary RA-consistent execution graph $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ of the SI or RSI implementations in Fig. 3, and pick arbitrary $wl, wl' \in \mathcal{WL}_x, pl, pl' \in \mathcal{PL}_x, wu, wu' \in \mathcal{WU}_x$ and $rl, rl' \in \mathcal{RL}_x$. We are then required to show:

$$(wl, wu), (wl', wu') \in \text{po}_x|_{\text{imm}} \Rightarrow (wu, wl') \in \text{lo}_x \vee (wu', wl) \in \text{lo}_x \quad (9)$$

$$(wl, wu), (rl, pl, wu') \in \text{po}_x|_{\text{imm}} \Rightarrow (wu, rl) \in \text{lo}_x \vee (wu', wl) \in \text{lo}_x \quad (10)$$

$$(rl, pl, wu), (rl', pl', wu') \in \text{po}_x|_{\text{imm}} \Rightarrow (wu, rl') \in \text{lo}_x \vee (wu', rl) \in \text{lo}_x \quad (11)$$

RTS. (9)

We proceed by contradiction. Let $(wl, wu), (wl', wu') \in \text{po}_x|_{\text{imm}}$ and $(wu, wl') \notin \text{lo}_x \wedge (wu', wl) \notin \text{lo}_x$. Since lo_x is totally ordered w.r.t. write lock events, we then have $(wl', wu) \in \text{lo}_x$ and $(wl, wu') \in \text{lo}_x$.

From the mutual exclusion (WEX) afforded by lock events, and since $(wl', wu) \in \text{lo}_x$ we know there exists $wu'' \in \mathcal{WU}_x$ such that $wl' \xrightarrow{\text{po}} wu''$ and $wu'' \xrightarrow{\text{lo}} wu$. Moreover, since $wl' \xrightarrow{\text{po}_x|_{\text{imm}}} wu'$ we have $wu' \xrightarrow{\text{po}^*} wu''$. Once again, since lo is total w.r.t. write lock events and agrees with po (otherwise we would have a cycle contradicting the assumption that G is RA-consistent), we have $wu' \xrightarrow{\text{lo}^*} wu''$. We then have $wu' \xrightarrow{\text{lo}^*} wu'' \xrightarrow{\text{lo}} wu$ and thus since lo is an order (i.e. is transitive), we have $wu' \xrightarrow{\text{lo}} wu$.

Following a similar argument symmetrically, we get $wu \xrightarrow{\text{lo}} wu'$. We then have $wu' \xrightarrow{\text{lo}} wu \xrightarrow{\text{lo}} wu'$, contradicting the assumption that lo is a strict order.

RTS. (10)

We proceed by contradiction. Let $(wl, wu), (rl, pl, wu') \in \text{po}_x|_{\text{imm}}$ and $(wu, rl) \notin \text{lo}_x \wedge (wu', wl) \notin \text{lo}_x$. Since lo_x is totally ordered w.r.t. write lock events, we then have $(rl, wu) \in \text{lo}_x$ and $(wl, wu') \in \text{lo}_x$.

From (RSHARE) and since $(rl, wu) \in \text{lo}_x$ we know there exists $l \in \mathcal{RU}_x \cup \mathcal{PL}_x$ such that $rl \xrightarrow{\text{po}} l$ and $l \xrightarrow{\text{lo}} wu$. Moreover, since $rl \xrightarrow{\text{po}_x|_{\text{imm}}} pl$ we have $pl \xrightarrow{\text{po}^*} l$. Once again, since lo is total w.r.t. write lock events and agrees with po (otherwise we would have a cycle contradicting the assumption that G is RA-consistent), we have $pl \xrightarrow{\text{lo}^*} l$. We then have $pl \xrightarrow{\text{lo}^*} l \xrightarrow{\text{lo}} wu$ and thus since lo is an order (i.e. is transitive), we have $pl \xrightarrow{\text{lo}} wu$. Similarly, from the mutual exclusion (WEX) of

write locks and since $pl \xrightarrow{\text{lo}} wu$, we know there exists $wu'' \in \mathcal{WU}_x$ such that $pl \xrightarrow{\text{po}} wu''$ and $wu'' \xrightarrow{\text{lo}} wu$. Moreover, since $pl \xrightarrow{\text{po}_x|_{\text{imm}}} wu'$ we have $wu' \xrightarrow{\text{po}^*} wu''$. Again, since lo is total w.r.t. write lock events and agrees with po (otherwise we would have a cycle contradicting the assumption that G is RA-consistent), we have $wu' \xrightarrow{\text{lo}^*} wu''$. We then have $wu' \xrightarrow{\text{lo}^*} wu'' \xrightarrow{\text{lo}} wu$ and thus since lo is an order (i.e. is transitive), we have $wu' \xrightarrow{\text{lo}} wu$.

Analogously, from (WEX) and since $(wl, wu') \in \text{lo}$, we know there exists $wu''' \in \mathcal{WU}_x$ such that $wl \xrightarrow{\text{po}} wu'''$ and $wu''' \xrightarrow{\text{lo}} wu'$. Moreover, since $wl \xrightarrow{\text{po}_x|_{\text{imm}}} wu$ we have $wu \xrightarrow{\text{po}^*} wu'''$. Once again, since lo is total w.r.t. write lock events and agrees with po (otherwise we would have a cycle contradicting the assumption that G is RA-consistent), we have $wu \xrightarrow{\text{lo}^*} wu'''$. We then have $wu \xrightarrow{\text{lo}^*} wu''' \xrightarrow{\text{lo}} wu'$ and thus since lo is an order (i.e. is transitive), we have $wu \xrightarrow{\text{lo}} wu'$.

We then have $wu' \xrightarrow{\text{lo}} wu \xrightarrow{\text{lo}} wu'$, contradicting the assumption that lo is a strict order.

RTS. (11)

We proceed by contradiction. Let $(rl, pl, wu), (rl', pl', wu') \in \text{po}_x|_{\text{imm}}$ and $(wu, rl') \notin \text{lo}_x \wedge (wu', rl) \notin \text{lo}_x$. Since lo_x is totally ordered w.r.t. write lock events, we then have $(rl, wu') \in \text{lo}_x$ and $(rl', wu) \in \text{lo}_x$.

From (RSHARE) and since $(rl, wu') \in \text{lo}_x$ we know there exists $l \in \mathcal{RU}_x \cup \mathcal{PL}_x$ such that $rl \xrightarrow{\text{po}} l$ and $l \xrightarrow{\text{lo}} wu'$. Moreover, since $rl \xrightarrow{\text{po}_x|_{\text{imm}}} pl$ we have $pl \xrightarrow{\text{po}^*} l$. Once again, since lo is total w.r.t. write lock events and agrees with po (otherwise we would have a cycle contradicting the assumption that G is RA-consistent), we have $pl \xrightarrow{\text{lo}^*} l$. We then have $pl \xrightarrow{\text{lo}^*} l \xrightarrow{\text{lo}} wu'$ and thus since lo is an order (i.e. is transitive), we have $pl \xrightarrow{\text{lo}} wu'$. Similarly, from (WEX) and since $pl \xrightarrow{\text{lo}} wu'$, we know there exists $wu'' \in \mathcal{WU}_x$ such that $pl \xrightarrow{\text{po}} wu''$ and $wu'' \xrightarrow{\text{lo}} wu'$. Moreover, since $pl \xrightarrow{\text{po}_x|_{\text{imm}}} wu$ we have $wu \xrightarrow{\text{po}^*} wu''$. Again, since lo is total w.r.t. write lock events and agrees with po (otherwise we would have a cycle contradicting the assumption that G is RA-consistent), we have $wu \xrightarrow{\text{lo}^*} wu''$. We then have $wu \xrightarrow{\text{lo}^*} wu'' \xrightarrow{\text{lo}} wu'$ and thus since lo is an order (i.e. is transitive), we have $wu \xrightarrow{\text{lo}} wu'$.

Following a similar argument symmetrically, we get $wu' \xrightarrow{\text{lo}} wu$. We then have $wu \xrightarrow{\text{lo}} wu' \xrightarrow{\text{lo}} wu$, contradicting the assumption that lo is a strict order.

Proof (Proof (RWSYNC)). Pick an arbitrary RA-consistent execution graph $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ of the SI or RSI implementations in Fig. 3, and pick arbitrary $wl \in \mathcal{WL}_x, pl \in \mathcal{PL}_x, wu \in \mathcal{WU}_x, rl, rl' \in \mathcal{RL}_x$ and $ru \in \mathcal{RU}_x$. We are then required to show:

$$(wl, wu), (rl, ru) \in \text{po}_x|_{\text{imm}} \Rightarrow (wu, rl) \in \text{lo}_x \vee (ru, wl) \in \text{lo}_x \quad (12)$$

$$(rl, pl, wu), (rl', ru) \in \text{po}_x|_{\text{imm}} \Rightarrow (wu, rl') \in \text{lo}_x \vee (ru, pl) \in \text{lo}_x \quad (13)$$

RTS. (12)

We proceed by contradiction. Let $(wl, wu), (rl, ru) \in \text{po}_x|_{\text{imm}}$ and $(wu, rl) \notin \text{lo}_x \wedge (ru, wl) \notin \text{lo}_x$. Since lo_x is totally ordered w.r.t. write lock events, we then have $(rl, wu) \in \text{lo}_x$ and $(wl, ru) \in \text{lo}_x$.

From (WEX) and since $(wl, ru) \in \text{lo}_x$ we know there exists $wu' \in \mathcal{WU}_x$ such that $wl \xrightarrow{\text{po}} wu'$ and $wu' \xrightarrow{\text{lo}} ru$. Moreover, since $wl \xrightarrow{\text{po}_x|_{\text{imm}}} wu$ we have $wu \xrightarrow{\text{po}^*} wu'$. Once again, since lo is total w.r.t. write lock events and agrees with po (otherwise we would have a cycle contradicting the assumption that G is RA-consistent), we have $wu \xrightarrow{\text{lo}^*} wu'$. We then have $wu \xrightarrow{\text{lo}^*} wu' \xrightarrow{\text{lo}} ru$ and thus since lo is an order (i.e. is transitive), we have $wu \xrightarrow{\text{lo}} ru$.

Analogously, from (RSHARE) and since $(rl, wu) \in \text{lo}_x$ we know there exists $l \in \mathcal{RU}_x \cup \mathcal{PL}_x$ such that $rl \xrightarrow{\text{po}} l$ and $l \xrightarrow{\text{lo}} wu$. Moreover, since $rl \xrightarrow{\text{po}_x|_{\text{imm}}} ru$ we have $ru \xrightarrow{\text{po}^*} l$. We then have $l \xrightarrow{\text{lo}} wu \xrightarrow{\text{lo}} ru \xrightarrow{\text{po}^*} l$, contradicting the assumption that G is RA-consistent.

RTS. (13)

We proceed by contradiction. Let $(rl, pl, wu), (rl', ru) \in \text{po}_x|_{\text{imm}}$ and $(wu, rl') \notin \text{lo}_x \wedge (ru, pl) \notin \text{lo}_x$. Since lo_x is totally ordered w.r.t. write lock events, we then have $(rl', wu) \in \text{lo}_x$ and $(pl, ru) \in \text{lo}_x$.

From (WEX) and since $(pl, ru) \in \text{lo}_x$ we know there exists $wu' \in \mathcal{WU}_x$ such that $pl \xrightarrow{\text{po}} wu'$ and $wu' \xrightarrow{\text{lo}} ru$. Moreover, since $pl \xrightarrow{\text{po}_x|_{\text{imm}}} wu$ we have $wu \xrightarrow{\text{po}^*} wu'$. Once again, since lo is total w.r.t. write lock events and agrees with po (otherwise we would have a cycle contradicting the assumption that G is RA-consistent), we have $wu \xrightarrow{\text{lo}^*} wu'$. We then have $wu \xrightarrow{\text{lo}^*} wu' \xrightarrow{\text{lo}} ru$ and thus since lo is an order (i.e. is transitive), we have $wu \xrightarrow{\text{lo}} ru$.

Analogously, from (RSHARE) and since $(rl', wu) \in \text{lo}_x$ we know there exists $l \in \mathcal{RU}_x \cup \mathcal{PL}_x$ such that $rl \xrightarrow{\text{po}} l$ and $l \xrightarrow{\text{lo}} wu$. Moreover, since $rl \xrightarrow{\text{po}_x|_{\text{imm}}} ru$ we have $ru \xrightarrow{\text{po}^*} l$. We then have $l \xrightarrow{\text{lo}} wu \xrightarrow{\text{lo}} ru \xrightarrow{\text{po}^*} l$, contradicting the assumption that G is RA-consistent.

C Soundness and Completeness of the Eager SI Implementation

Notation Given an execution graph $G' = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ we write \mathcal{T}/st for the set of equivalence classes of \mathcal{T} induced by **st**; $[a]_{\text{st}}$ for the equivalence class that contains a ; and \mathcal{T}_ξ for the equivalence class of transaction $\xi \in \text{TXID}$: $\mathcal{T}_\xi \triangleq \{a \mid \text{tx}(a) = \xi\}$. We write $\text{si-consistent}(G')$ to denote that G' is SI-consistent; and write $\text{RA-consistent}(G)$ to denote that G is RA-consistent.

Given an execution graph G of the SI implementation in Fig. 3, let us assign a transaction identifier to each transaction executed by the program; and given a transaction ξ , let RS_ξ and WS_ξ denote its read and write sets, respectively. Observe that given a transaction ξ of the SI implementation in Fig. 3 with $\text{RS}_\xi \cup \text{WS}_\xi = \{\mathbf{x}_1, \dots, \mathbf{x}_i\}$, the trace of ξ , written θ_ξ , is of the following form:

$$\theta_\xi = FS^* \xrightarrow{\text{po|imm}} Rs \xrightarrow{\text{po|imm}} RUs \xrightarrow{\text{po|imm}} PLs \xrightarrow{\text{po|imm}} Ts \xrightarrow{\text{po|imm}} Us$$

where FS^* denotes the sequence of events attempting but failing to obtain a snapshot, and

- Rs denotes the sequence of events acquiring the reader locks (on all locations accessed) and capturing a snapshot of the read set, and is of the form $rl_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} rl_{\mathbf{x}_i} \xrightarrow{\text{po|imm}} S_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} S_{\mathbf{x}_i}$, where for all $n \in \{1 \dots i\}$:

$$rl_{\mathbf{x}_n} = \text{RL}(\mathbf{x}1_n) \quad S_{\mathbf{x}_n} = \begin{cases} rs_{\mathbf{x}_n} \xrightarrow{\text{po|imm}} ws_{\mathbf{x}_n} & \text{if } \mathbf{x} \in \text{RS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

with $rs_{\mathbf{x}_n} \triangleq \text{R}(\mathbf{x}_n, v_n)$ and $ws_{\mathbf{x}_n} \triangleq \text{W}(\mathbf{s}[\mathbf{x}_n], v_n)$, for some v_n .

- RUs denotes the sequence of events releasing the reader locks (when the given location is in the read set only) and is of the form $ru_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} ru_{\mathbf{x}_i}$, where for all $n \in \{1 \dots i\}$:

$$ru_{\mathbf{x}_n} = \begin{cases} \text{RU}(\mathbf{x}1_n) & \text{if } \mathbf{x}_n \notin \text{WS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- PLs denotes the sequence of events promoting the reader locks to writer ones (when the given location is in the write set), and is of the form $pl_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} pl_{\mathbf{x}_i}$, where for all $n \in \{1 \dots i\}$:

$$pl_{\mathbf{x}_n} = \begin{cases} \text{PL}(\mathbf{x}1_n) & \text{if } \mathbf{x}_n \in \text{WS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- Ts denotes the sequence of events corresponding to the execution of (\mathbb{T}) in Fig. 3 and is of the form $t_1 \xrightarrow{\text{po}_{\text{imm}}} \dots \xrightarrow{\text{po}_{\text{imm}}} t_k$, where for all $m \in \{1 \dots k\}$:

$$t_m = \begin{cases} \mathbf{R}(\mathbf{s}[\mathbf{x}_n], v_n) & \text{if } O_m = \mathbf{R}(\mathbf{x}_n, v_n) \\ \mathbf{W}(\mathbf{x}_n, v_n) \xrightarrow{\text{po}_{\text{imm}}} \mathbf{W}(\mathbf{s}[\mathbf{x}_n], v_n) & \text{if } O_m = \mathbf{W}(\mathbf{x}_n, v_n) \end{cases}$$

where O_m denotes the m th event in the trace of the original \mathbb{T} ;

- Us denotes the sequence of events releasing the locks on the write set. That is, the events in Us correspond to the execution of the last line of the implementation in Fig. 3, and is of the form $wu_{x_1} \xrightarrow{\text{po}_{\text{imm}}} \dots \xrightarrow{\text{po}_{\text{imm}}} wu_{x_i}$, where for all $n \in \{1 \dots i\}$:

$$wu_{x_n} = \begin{cases} \mathbf{WU}(y_{1_n}) & \text{if } \mathbf{x}_n \in \mathbf{WS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

Given a transaction trace θ_ξ , we write e.g. $\xi.Ls$ to refer to its constituent Ls sub-trace and write $Ls.E$ for the set of events related by po in Ls . Similarly, we write $\xi.E$ for the set of events related by po in θ_ξ . Note that $G.E = \bigcup_{\xi \in \text{Tx}} \xi.E$.

C.1 Implementation Soundness

In order to establish the soundness of our implementation, it suffices to show that given an RA-consistent execution graph of the implementation G , we can construct a corresponding SI-consistent execution graph G' with the same outcome.

Given a transaction $\xi \in \text{Tx}$ with $\mathbf{RS}_\xi \cup \mathbf{WS}_\xi = \{\mathbf{x}_1 \dots \mathbf{x}_i\}$ and trace $\theta_\xi = F_s^* \xrightarrow{\text{po}_{\text{imm}}} R_s \xrightarrow{\text{po}_{\text{imm}}} RU_s \xrightarrow{\text{po}_{\text{imm}}} PL_s \xrightarrow{\text{po}_{\text{imm}}} Ts \xrightarrow{\text{po}_{\text{imm}}} Us$, with $Ts = t_1 \xrightarrow{\text{po}_{\text{imm}}} \dots \xrightarrow{\text{po}_{\text{imm}}} t_k$, we construct the corresponding implementation trace θ'_ξ as follows:

$$\theta'_\xi \triangleq t'_1 \xrightarrow{\text{po}} \dots \xrightarrow{\text{po}} t'_k$$

where for all $m \in \{1 \dots k\}$:

$$\begin{aligned} t'_m &= \mathbf{R}(\mathbf{x}_n, rb_n) \text{ when } t_m = \mathbf{R}(\mathbf{s}[\mathbf{x}_n], rb_n) \\ t'_m &= \mathbf{W}(\mathbf{x}_n, rb_n) \text{ when } t_m = \mathbf{W}(\mathbf{x}_n, rb_n) \xrightarrow{\text{po}_{\text{imm}}} \mathbf{W}(\mathbf{s}[\mathbf{x}_n], rb_n) \end{aligned}$$

such that in the first case the identifier of t'_m is that of t_m ; and in the second case the identifier of t'_m is that of the first event in t_m . We then define:

$$\mathbf{RF}_\xi \triangleq \left\{ (w, t'_j) \left| \begin{array}{l} t'_j \in Ts' \wedge \exists \mathbf{x}, v. t'_j = \mathbf{R}(\mathbf{x}, v) \wedge w = \mathbf{W}(\mathbf{x}, v) \\ \wedge (w \in \xi.E \Rightarrow w \xrightarrow{\text{po}} t'_j \wedge \\ \quad (\forall e \in \xi.E. w \xrightarrow{\text{po}} e \xrightarrow{\text{po}} t'_j \Rightarrow (\text{loc}(e) \neq \mathbf{x} \vee e \notin \mathcal{W}))) \\ \wedge (w \notin \xi.E \Rightarrow (\forall e \in \xi.E. (e \xrightarrow{\text{po}} t'_j \Rightarrow (\text{loc}(e) \neq \mathbf{x} \vee e \notin \mathcal{W}))) \\ \wedge (w, \xi.rs_{\mathbf{x}}) \in G.\mathbf{rf}) \end{array} \right. \right\}$$

We are now in a position to demonstrate the soundness of our implementation. Given an RA-consistent execution graph G of the implementation, we construct an SI execution graph G' as follows and demonstrate that $\text{si-consistent}(G')$ holds.

– $G'.E = \bigcup_{\xi \in \text{TX}} \theta'_\xi.E$, with the $\text{tx}(\cdot)$ function defined as:

$$\text{tx}(a) \triangleq \xi \quad \text{where} \quad a \in \theta'_\xi$$

- $G'.\text{po} = G.\text{po}|_{G'.E}$
- $G'.\text{rf} = \bigcup_{\xi \in \text{TX}} \text{RF}_\xi$
- $G'.\text{mo} = G.\text{mo}|_{G'.E}$
- $G'.\text{lo} = \emptyset$

Observe that the events of each θ'_ξ trace coincides with those of the equivalence class \mathcal{T}_ξ of G' . That is, $\theta'_\xi.E = \mathcal{T}_\xi$.

Lemma 3. *Given an RA-consistent execution graph G of the implementation and its corresponding SI execution graph G' constructed as above, for all $a, b, \xi_a, \xi_b, \mathbf{x}$:*

$$\xi_a \neq \xi_b \wedge a \in \xi_a.E \wedge b \in \xi_b.E \wedge \text{loc}(a) = \text{loc}(b) = \mathbf{x} \Rightarrow$$

$$((a, b) \in G'.\text{rf} \Rightarrow \xi_a.wu_{\mathbf{x}} \xrightarrow{G.\text{hb}} \xi_b.rl_{\mathbf{x}}) \tag{14}$$

$$\wedge ((a, b) \in G'.\text{mo} \Rightarrow \xi_a.wu_{\mathbf{x}} \xrightarrow{G.\text{hb}} \xi_b.rl_{\mathbf{x}}) \tag{15}$$

$$\wedge ((a, b) \in G'.\text{mo}; \text{rf} \Rightarrow \xi_a.wu_{\mathbf{x}} \xrightarrow{G.\text{hb}} \xi_b.rl_{\mathbf{x}}) \tag{16}$$

$$\wedge ((a, b) \in G'.\text{rb} \Rightarrow (\mathbf{x} \in \text{WS}_{\xi_a} \wedge \xi_a.wu_{\mathbf{x}} \xrightarrow{G.\text{hb}} \xi_b.rl_{\mathbf{x}}) \vee (\mathbf{x} \notin \text{WS}_{\xi_a} \wedge \xi_a.ru_{\mathbf{x}} \xrightarrow{G.\text{hb}} \xi_b.pl_{\mathbf{x}})) \tag{17}$$

Proof. Pick an arbitrary RA-consistent execution graph G of the implementation and its corresponding SI execution graph G' constructed as above. Pick an arbitrary $a, b, \xi_a, \xi_b, \mathbf{x}$ such that $\xi_a \neq \xi_b$, $a \in \xi_a.E$, $a \in \xi_a.E$, and $\text{loc}(a) = \text{loc}(b) = \mathbf{x}$.

RTS. (14)

Assume $(a, b) \in G'.\text{rf}$. From the definition of $G'.\text{rf}$ we then know $(a, \xi_a.rs_{\mathbf{x}}) \in G.\text{rf}$. On the other hand, from [Lemma 2](#) we know that either i) $\mathbf{x} \in \text{WS}_{\xi_b}$ and $\xi_b.wu_{\mathbf{x}} \xrightarrow{G.\text{hb}} \xi_b.rl_{\mathbf{x}}$; or ii) $\mathbf{x} \notin \text{WS}_{\xi_b}$ and $\xi_b.ru_{\mathbf{x}} \xrightarrow{G.\text{hb}} \xi_b.pl_{\mathbf{x}}$; or iii) $\xi_a.wu_{\mathbf{x}} \xrightarrow{G.\text{hb}} \xi_b.rl_{\mathbf{x}}$. In case (i) we then have $a \xrightarrow{G.\text{rf}} \xi_a.rs_{\mathbf{x}} \xrightarrow{G.\text{po}} \xi_b.wu_{\mathbf{x}} \xrightarrow{G.\text{hb}} \xi_b.rl_{\mathbf{x}} \xrightarrow{G.\text{po}} a$. That is, we have $a \xrightarrow{G.\text{hb}_{\text{loc}}} a$, contradicting the assumption that G is RA-consistent. Similarly in case (ii) we have $a \xrightarrow{G.\text{rf}} \xi_a.rs_{\mathbf{x}} \xrightarrow{G.\text{po}} \xi_b.ru_{\mathbf{x}} \xrightarrow{G.\text{hb}} \xi_b.pl_{\mathbf{x}} \xrightarrow{G.\text{po}} a$. That is, we have $a \xrightarrow{G.\text{hb}_{\text{loc}}} a$, contradicting the assumption that G is RA-consistent. In case (iii) the desired result holds trivially.

RTS. (15)

Assume $(a, b) \in G'.\text{mo}$. From the definition of $G'.\text{mo}$ we then know $(a, b) \in G.\text{mo}$.

On the other hand, from [Lemma 2](#) we know that either i) $\xi_b.wu_x \xrightarrow{G.hb} \xi_a.rl_x$; or ii) $\xi_a.wu_x \xrightarrow{G.hb} \xi_b.rl_x$. In case (i) we then have $a \xrightarrow{G.mo} b \xrightarrow{G.po} \xi_b.wu_x \xrightarrow{G.hb} \xi_a.rl_x \xrightarrow{G.po} a$. That is, we have $a \xrightarrow{G.hb_{loc}} a$, contradicting the assumption that G is RA-consistent. In case (ii) the desired result holds trivially.

RTS. (16)

Assume $(a, b) \in G'.mo; rf$. We then know there exists w such that $(a, w) \in G'.mo$ and $(w, b) \in G'.rf$. From the definition of $G'.mo$ we then know $(a, w) \in G.mo$. There are now three cases to consider: 1) $w \in \xi_a$; or 2) $w \in \xi_b$; or 3) $w \in \xi_c \wedge \xi_c \neq \xi_a \wedge \xi_c \neq \xi_b$. In case (1) the desired result follows from [part 14](#). In case (2) since $(a, w) \in G'.mo$ the desired result follows from [part 15](#).

In case (3) from the proof of [part 15](#) we have $\xi_a.wu_x \xrightarrow{G.hb} \xi_c.rl_x$. Moreover, from the shape of G traces we have $\xi_c.rl_x \xrightarrow{G.po} \xi_c.wu_x$. On the other hand, from the proof of [part 14](#) we have $\xi_c.wu_x \xrightarrow{G.hb} \xi_b.rl_x$. We thus have $\xi_a.wu_x \xrightarrow{G.hb} \xi_c.rl_x \xrightarrow{G.po} \xi_c.wu_x \xrightarrow{G.hb} \xi_b.rl_x$. As $G.po \subseteq G.hb$ and $G.hb$ is transitively closed, we have $\xi_a.wu_x \xrightarrow{G.hb} \xi_b.rl_x$, as required.

RTS. (17)

Assume $(a, b) \in G'.rb$. From the definition of $G'.rb$ we then know $(\xi_a.rs_x, b), (\xi_a.vs_x, b) \in G.rb$. On the other hand, from [Lemma 2](#) we know that either i) $\xi_b.wu_x \xrightarrow{G.hb} \xi_a.rl_x$; or ii) $\mathbf{x} \notin WS_{\xi_a}$ and $\xi_a.ru_x \xrightarrow{G.hb} \xi_a.pl_x$; or iii) $\mathbf{x} \in WS_{\xi_a}$ and $\xi_a.wu_x \xrightarrow{G.hb} \xi_b.rl_x$. In case (i) we then have $b \xrightarrow{G.po} \xi_b.wu_x \xrightarrow{G.hb} \xi_a.rl_x \xrightarrow{G.po} \xi_a.rs_x \xrightarrow{G.rb} b$. That is, we have $b \xrightarrow{G.hb_{loc}} \xi_a.rs_x \xrightarrow{G.rb} b$, contradicting the assumption that G is RA-consistent. In cases (ii-iii) the desired result holds trivially.

Lemma 4. For all RA-consistent execution graphs G of the implementation and their counterpart SI execution graphs G' constructed as above,

1. $(G'.po_{\top} \subseteq G.po) \wedge (G'.po_{\top}; G'.rb_{\top} \subseteq G.hb)$
2. $(G'.mo_{\top} \subseteq G.hb) \wedge (G'.mo_{\top}; G'.rb_{\top} \subseteq G.hb)$
3. $(G'.rf_{\top} \subseteq G.hb) \wedge (G'.rf_{\top}; G'.rb_{\top} \subseteq G.hb)$

Proof. Pick an arbitrary RA-consistent execution graph G of the implementation and its counterpart SI execution graph G' constructed as above.

RTS. (Part 1)

The proof of the first conjunct is immediate from the definitions of $G'.po$ and $G.po$. For the second conjunct, pick arbitrary $(a, b) \in G'.(po_{\top}; rb_{\top})$. We then know there exist c such that $(a, c) \in G'.po_{\top}$ and $(c, b) \in G'.rb_{\top}$. Since $(a, c) \in G'.po_{\top}$, from the definition of $G'.po$ we also have $(a, c) \in G.po$ and thus $(a, c) \in G.hb$. Moreover, from the definition of $G'.rb_{\top}$ we know there exist ξ_1, ξ_2, r, w such that $\xi_1 \neq \xi_2$, $c, r \in \theta'_{\xi_1}$, $b, w \in \theta'_{\xi_2}$ and $(r, w) \in G'.rb$. Let $loc(r) = loc(w) = \mathbf{x}$. We then know that $\mathbf{x} \in RS_{\xi_1} \cup WS_{\xi_1}$, $\mathbf{x} \in WS_{\xi_2}$, and that there exists w_x such that

$(w_x, r) \in G'.\mathbf{rf}$ and $(w_x, w) \in G'.\mathbf{mo}$. From the construction of $G'.\mathbf{mo}$ we then have $(w_x, w) \in G.\mathbf{mo}$. From the construction of $G'.\mathbf{rf}$ there are now two cases to consider: 1) $(w_x, \xi_1.rs_x) \in G.\mathbf{rf}$; or 2) $\mathbf{x} \in \mathbf{WS}_{\xi_1}$, $w_x \in \xi_1$ and $(w_x, r) \in G.\mathbf{po}$.

In case (1) we then have $(\xi_1.rs_x, w) \in G.\mathbf{rb}$. Moreover, from [Lemma 3](#) we have that either: i) $\mathbf{x} \notin \mathbf{WS}_{\xi_1} \wedge \xi_1.ru_x \xrightarrow{G.\mathbf{hb}} \xi_2.pl_x$; or ii) $\mathbf{x} \in \mathbf{WS}_{\xi_1} \wedge \xi_1.wu_x \xrightarrow{G.\mathbf{hb}} \xi_2.rl_x$.

In case (1.i), since we have $(a, c) \in G'.\mathbf{po}_T$, from the construction of $G'.\mathbf{po}$ we have $(a, \xi_1.ru_x) \in G.\mathbf{po}$. We thus have $a \xrightarrow{G.\mathbf{po}} \xi_1.ru_x \xrightarrow{G.\mathbf{hb}} \xi_2.pl_x \xrightarrow{G.\mathbf{po}} b$. That is, as $G.\mathbf{po} \subseteq G.\mathbf{hb}$ and G is transitively closed, we have $(a, b) \in G.\mathbf{hb}$, as required. Similarly, in case (1.ii), since we have $(a, c) \in G'.\mathbf{po}_T$, from the construction of $G'.\mathbf{po}$ we have $(a, \xi_1.wu_x) \in G.\mathbf{po}$. We thus have $a \xrightarrow{G.\mathbf{po}} \xi_1.wu_x \xrightarrow{G.\mathbf{hb}} \xi_2.rl_x \xrightarrow{G.\mathbf{po}} b$. That is, as $G.\mathbf{po} \subseteq G.\mathbf{hb}$ and G is transitively closed, we have $(a, b) \in G.\mathbf{hb}$, as required.

In case (2), from [Lemma 3](#) we have $\xi_1.wu_x \xrightarrow{G.\mathbf{hb}} \xi_2.rl_x$. Moreover, we have $(c, \xi_1.wu_x) \in G.\mathbf{po}$ and $(\xi_2.rl_x, b) \in G.\mathbf{po}$. We thus have, $a \xrightarrow{G.\mathbf{hb}} c \xrightarrow{G.\mathbf{po}} \xi_1.wu_x \xrightarrow{G.\mathbf{hb}} \xi_2.rl_x \xrightarrow{G.\mathbf{po}} b$. That is, as $G.\mathbf{po} \subseteq G.\mathbf{hb}$ and G is transitively closed, we have $(a, b) \in G.\mathbf{hb}$, as required.

RTS. (Part 2)

For the first conjunct, pick an arbitrary $(a, b) \in G'.\mathbf{mo}_T$; we are then required to show that $(a, b) \in G.\mathbf{hb}$.

From the definition of $G'.\mathbf{mo}_T$ and the construction of G' we know there exist ξ_1, ξ_2, d, e such that $\xi_1 \neq \xi_2$, $(d, e) \in G'.\mathbf{mo}$, $a, d \in \mathcal{T}_{\xi_1}$ and $b, e \in \mathcal{T}_{\xi_2}$. Let $\mathbf{loc}(d) = \mathbf{loc}(e) = \mathbf{y}$. We then know $a \xrightarrow{G.\mathbf{po}} \xi_1.wu_y$, $\xi_1.rl_y \xrightarrow{G.\mathbf{po}} d \xrightarrow{G.\mathbf{po}} \xi_1.wu_y$, $\xi_2.rl_y \xrightarrow{G.\mathbf{po}} b$ and $\xi_2.rl_y \xrightarrow{G.\mathbf{po}} e \xrightarrow{G.\mathbf{po}} \xi_2.wu_y$.

From [Lemma 3](#) we then know $\xi_1.wu_y \xrightarrow{G.\mathbf{hb}} \xi_2.rl_y$. As such, we have $a \xrightarrow{G.\mathbf{po}} \xi_1.wu_y \xrightarrow{G.\mathbf{hb}} \xi_2.rl_y \xrightarrow{G.\mathbf{po}} b$. As $G.\mathbf{po} \in G.\mathbf{hb}$ and $G.\mathbf{hb}$ is transitively closed, we have $a \xrightarrow{G.\mathbf{hb}} b$, as required.

For the second conjunct, pick an arbitrary $(b, c) \in G'.\mathbf{rb}_T$. We are then required to show that $(a, c) \in G.\mathbf{hb}$. From the definition of $G'.\mathbf{rb}_T$ we then know there exist ξ_3, r, w such that $\xi_3 \neq \xi_2$, $r \in \mathcal{T}_{\xi_2}$, $c, w \in \mathcal{T}_{\xi_3}$ and $(r, w) \in G'.\mathbf{rb}$. Let $\mathbf{loc}(r) = \mathbf{loc}(w) = \mathbf{x}$. We then know that $\mathbf{x} \in \mathbf{RS}_{\xi_2} \cup \mathbf{WS}_{\xi_2}$, $\mathbf{x} \in \mathbf{WS}_{\xi_3}$, and that there exists w_x such that $(w_x, r) \in G'.\mathbf{rf}$ and $(w_x, w) \in G'.\mathbf{mo}$. From the construction of $G'.\mathbf{mo}$ we then have $(w_x, w) \in G.\mathbf{mo}$. From the construction of $G'.\mathbf{rf}$ there are now two cases to consider: 1) $(w_x, \xi_2.rs_x) \in G.\mathbf{rf}$; or 2) $\mathbf{x} \in \mathbf{WS}_{\xi_2}$, $w_x \in \xi_2$ and $(w_x, r) \in G.\mathbf{po}$.

In case (1) we then have $(\xi_2.rs_x, w) \in G.\mathbf{rb}$. Moreover, from [Lemma 3](#) we have either i) $\mathbf{x} \notin \mathbf{WS}_{\xi_2} \wedge \xi_2.ru_x \xrightarrow{G.\mathbf{hb}} \xi_3.pl_x$; or ii) $\mathbf{x} \in \mathbf{WS}_{\xi_2} \wedge \xi_2.wu_x \xrightarrow{G.\mathbf{hb}} \xi_3.rl_x$. In case (1.i), from the proof of the first conjunct recall that we have $a \xrightarrow{G.\mathbf{po}} \xi_1.wu_y \xrightarrow{G.\mathbf{hb}} \xi_2.rl_y$. Also, from the shape of G traces we know that $\xi_2.rl_y \xrightarrow{G.\mathbf{po}} \xi_2.ru_x$. As such, we have $a \xrightarrow{G.\mathbf{po}} \xi_1.wu_y \xrightarrow{G.\mathbf{hb}} \xi_2.rl_y \xrightarrow{G.\mathbf{po}} \xi_2.ru_x \xrightarrow{G.\mathbf{hb}} \xi_3.pl_x \xrightarrow{G.\mathbf{po}} c$. That is, as $G.\mathbf{po} \subseteq G.\mathbf{hb}$ and G is transitively closed, we have $(a, c) \in G.\mathbf{hb}$, as required.

Similarly, in case (1.ii), from the proof of the first conjunct recall that we have $a \xrightarrow{G.\text{po}} \xi_1.wu_y \xrightarrow{G.\text{hb}} \xi_2.rl_y$. Also, from the shape of G traces we know that $\xi_2.rl_y \xrightarrow{G.\text{po}} \xi_2.wu_x$. As such, we have $a \xrightarrow{G.\text{po}} \xi_1.wu_y \xrightarrow{G.\text{hb}} \xi_2.rl_y \xrightarrow{G.\text{po}} \xi_2.wu_x \xrightarrow{G.\text{hb}} \xi_3.rl_x \xrightarrow{G.\text{po}} c$. That is, as $G.\text{po} \subseteq G.\text{hb}$ and G is transitively closed, we have $(a, c) \in G.\text{hb}$, as required.

In case (2) from [Lemma 3](#) we have $\xi_2.wu_x \xrightarrow{G.\text{hb}} \xi_3.rl_x$. Recall that with the first conjunct we demonstrated that $(a, b) \in G.\text{hb}$. Moreover, we have $(b, \xi_2.wu_x) \in G.\text{po}$ and $(\xi_3.rl_x, c) \in G.\text{po}$. We thus have, $a \xrightarrow{G.\text{hb}} b \xrightarrow{G.\text{po}} \xi_2.wu_x \xrightarrow{G.\text{hb}} \xi_3.rl_x \xrightarrow{G.\text{po}} c$. That is, as $G.\text{po} \subseteq G.\text{hb}$ and G is transitively closed, we have $(a, b) \in G.\text{hb}$, as required.

RTS. (Part 3)

For the first conjunct, pick an arbitrary $(a, b) \in G'.\text{rf}_\top$; we are then required to show that $(a, b) \in G.\text{hb}$.

From the definition of $G'.\text{rf}_\top$ and the construction of G' we know there exist ξ_1, ξ_2, w_y, r_y such that $\xi_1 \neq \xi_2$, $(w_y, r_y) \in G'.\text{rf}$, $a, w_y \in \xi_1$ and $b, r_y \in \xi_2$. Let $\text{loc}(w) = \text{loc}(r) = y$. We then know $\xi_1.rl_y \xrightarrow{G.\text{po}} \xi_1.pl_y \xrightarrow{G.\text{po}} w_y \xrightarrow{G.\text{po}} \xi_1.wu_y$ and $a \xrightarrow{G.\text{po}} \xi_1.wu_y$.

Let $w_y = \mathbb{W}(y, v)$ and $r_y = \mathbb{R}(y, v)$. From the construction of G' we know that $(w_y, \xi_2.rs_y) \in G'.\text{rf}$. On the other hand, from [Lemma 3](#) we have $\xi_1.wu_y \xrightarrow{G.\text{hb}} \xi_2.rl_y$. We then have $a \xrightarrow{G.\text{po}} \xi_1.wu_y \xrightarrow{G.\text{hb}} \xi_2.rl_y \xrightarrow{G.\text{po}} b$. That is, as $G.\text{po} \subseteq G.\text{hb}$, we have $a \xrightarrow{G.\text{hb}} b$, as required.

For the second conjunct, pick an arbitrary $(b, c) \in G'.\text{rb}_\top$. We are then required to show that $(a, c) \in G.\text{hb}$. From the definition of $G'.\text{rb}_\top$ we then know there exist ξ_3, r, w such that $\xi_3 \neq \xi_2$, $r \in \mathcal{T}_{\xi_2}$, $c, w \in \mathcal{T}_{\xi_3}$ and $(r, w) \in G'.\text{rb}$. Let $\text{loc}(r) = \text{loc}(w) = x$. We then know that $x \in \text{RS}_{\xi_2} \cup \text{WS}_{\xi_2}$, $x \in \text{WS}_{\xi_3}$, and that there exists w_x such that $(w_x, r) \in G'.\text{rf}$ and $(w_x, w) \in G'.\text{mo}$. From the construction of $G'.\text{mo}$ we then have $(w_x, w) \in G.\text{mo}$. From the proof of the first conjunct recall that we have $a \xrightarrow{G.\text{po}} \xi_1.wu_y \xrightarrow{G.\text{hb}} \xi_2.rl_y$.

From the construction of $G'.\text{rf}$ there are now two cases to consider: 1) $(w_x, \xi_2.rs_x) \in G'.\text{rf}$; or 2) $x \in \text{WS}_{\xi_2}$, $w_x \in \xi_2$ and $(w_x, r) \in G.\text{po}$.

In case (1) we then have $(\xi_2.rs_x, w) \in G'.\text{rb}$. Moreover, from [Lemma 3](#) we have either: i) $\xi_2.ru_x \xrightarrow{G.\text{hb}} \xi_3.pl_x$; or ii) $\xi_2.wu_x \xrightarrow{G.\text{hb}} \xi_3.rl_x$.

In case (1.i), from the shape of G traces we know that $\xi_2.rl_y \xrightarrow{G.\text{po}} \xi_2.ru_x$. As such, we have $a \xrightarrow{G.\text{po}} \xi_1.wu_y \xrightarrow{G.\text{hb}} \xi_2.rl_y \xrightarrow{G.\text{po}} \xi_2.ru_x \xrightarrow{G.\text{hb}} \xi_3.pl_x \xrightarrow{G.\text{po}} c$. That is, as $G.\text{po} \subseteq G.\text{hb}$ and G is transitively closed, we have $(a, c) \in G.\text{hb}$, as required.

Similarly, in case (1.ii), from the shape of G traces we know that $\xi_2.rl_y \xrightarrow{G.\text{po}} \xi_2.wu_x$. As such, we have $a \xrightarrow{G.\text{po}} \xi_1.wu_y \xrightarrow{G.\text{hb}} \xi_2.rl_y \xrightarrow{G.\text{po}} \xi_2.wu_x \xrightarrow{G.\text{hb}} \xi_3.rl_x \xrightarrow{G.\text{po}} c$. That is, as $G.\text{po} \subseteq G.\text{hb}$ and G is transitively closed, we have $(a, c) \in G.\text{hb}$, as required.

In case (2), from [Lemma 3](#) we have $\xi_2.wu_x \xrightarrow{G.\text{hb}} \xi_3.rl_x$. Recall that with the first conjunct we demonstrated that $(a, b) \in G.\text{hb}$. Moreover, we have $(b, \xi_2.wu_x) \in G.\text{po}$ and $(\xi_3.rl_x, c) \in G.\text{po}$. We thus have, $a \xrightarrow{G.\text{hb}} b \xrightarrow{G.\text{po}} \xi_2.wu_x \xrightarrow{G.\text{hb}} \xi_3.rl_x \xrightarrow{G.\text{po}} c$. That is, as $G.\text{po} \subseteq G.\text{hb}$ and G is transitively closed, we have $(a, b) \in G.\text{hb}$, as required.

Theorem 7 (Soundness). *For all execution graphs G of the implementation and their counterpart SI execution graphs G' constructed as above,*

$$\text{RA-consistent}(G) \Rightarrow \text{si-consistent}(G')$$

Proof. Pick an arbitrary execution graph G of the implementation such that $\text{RA-consistent}(G)$ holds, and its associated SI execution graph G' constructed as described above.

RTS. $\text{acyclic}(G'.((\text{po}_T \cup \text{rf}_T \cup \text{mo}_T); \text{rb}_T^?))$

We proceed by contradiction. Let us assume

$$\text{RA-consistent}(G) \wedge \neg \text{acyclic}(G'.((\text{po}_T \cup \text{rf}_T \cup \text{mo}_T); \text{rb}_T^?)) \quad (18)$$

From the definition of $\text{RA-consistent}(\cdot)$ we then know that there exists a such that $(a, a) \in ((G'.\text{po}_T \cup G'.\text{rf}_T \cup G'.\text{mo}_T); G'.\text{rb}_T^?)^+$. Consequently, from [Lemma 4](#) we have $(a, a) \in G.\text{hb}^+$. That is, since $G.\text{hb}$ is transitively closed, we have $(a, a) \in G.\text{hb}$, contradicting our assumption that G is RA-consistent.

RTS. $\text{rf}_I \cup \text{mo}_I \cup \text{rb}_I \subseteq \text{po}$

Follows immediately from the construction of G' .

C.2 Implementation Completeness

In order to establish the completeness of our implementation, it suffices to show that given an SI-consistent execution graph $G' = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$, we can construct a corresponding RA-consistent execution graph G of the implementation. Before proceeding with the construction of a corresponding implementation graph, we describe several auxiliary definitions.

Given an abstract transaction class $\mathcal{T}_\xi \in G'.\mathcal{T}/\text{st}$, we write WS_ξ for the set of locations written to by \mathcal{T}_ξ : $\text{WS}_\xi = \bigcup_{e \in \mathcal{T}_\xi \cap \mathcal{W}} \text{loc}(e)$. Similarly, we write RS_ξ for the set of locations read from by \mathcal{T}_ξ , *prior to* being written by \mathcal{T}_ξ . For each location x read from by \mathcal{T}_ξ , we additionally record the first read event in \mathcal{T}_ξ that retrieved the value of x . That is,

$$\text{RS}_{\mathcal{T}_\xi} \triangleq \left\{ (x, r) \mid r \in \mathcal{T}_i \cap \mathcal{R}_x \wedge \neg \exists e \in \mathcal{T}_\xi \cap E_x. e \xrightarrow{\text{po}} r \right\}$$

Note that the execution trace for each transaction $\mathcal{T}_\xi \in G'.\mathcal{T}/\text{st}$ is of the form $\theta'_\xi = t'_1 \xrightarrow{\text{po}|_{\text{imm}}} \dots \xrightarrow{\text{po}|_{\text{imm}}} t'_k$ for some k , where each t'_i is a read or write event. As such, we have $G'.E = G'.\mathcal{T} = \bigcup_{\mathcal{T}_\xi \in G'.\mathcal{T}/\text{st}} \mathcal{T}_\xi = \theta'_\xi.E$.

Let $\mathbf{RS}_{\mathcal{T}_\xi} \cup \mathbf{WS}_{\mathcal{T}_\xi} = \{\mathbf{x}_1 \cdots \mathbf{x}_n\}$. We then construct the implementation trace θ_ξ as:

$$\theta_\xi = Rs \xrightarrow{\text{po|imm}} RUs \xrightarrow{\text{po|imm}} PLs \xrightarrow{\text{po|imm}} Ts \xrightarrow{\text{po|imm}} Us$$

where

- $Rs = rl_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} rl_{\mathbf{x}_n} \xrightarrow{\text{po|imm}} S_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} S_{\mathbf{x}_n}$, where the identifiers of all constituent events of Rs are picked fresh, and

$$rl_{\mathbf{x}_j} = \mathbf{R}(\mathbf{x}_j) \quad S_{\mathbf{x}_j} = \begin{cases} rs_{\mathbf{x}_j} \xrightarrow{\text{po|imm}} ws_{\mathbf{x}_j} & \text{if } \exists r. (\mathbf{x}_j, r) \in \mathbf{RS}_\xi \wedge \mathbf{val}_r(r) = v_j \\ \emptyset & \text{otherwise} \end{cases}$$

with $rs_{\mathbf{x}_j} \triangleq \mathbf{R}(\mathbf{x}_j, v_j)$ and $ws_{\mathbf{x}_j} \triangleq \mathbf{W}(\mathbf{s}[\mathbf{x}_j], v_j)$.

- $RUs = ru_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} ru_{\mathbf{x}_n}$, where the identifiers of all constituent events of RUs are picked fresh, and for all $j \in \{1 \cdots n\}$:

$$ru_{\mathbf{x}_j} = \begin{cases} \mathbf{R}(\mathbf{x}_j) & \text{if } \mathbf{x}_j \notin \mathbf{WS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- $PLs = pl_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} pl_{\mathbf{x}_n}$, where the identifiers of all constituent events of PLs are picked fresh, and for all $j \in \{1 \cdots n\}$:

$$pl_{\mathbf{x}_j} = \begin{cases} \mathbf{P}(\mathbf{x}_j) & \text{if } \mathbf{x}_j \in \mathbf{WS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- $Ts = t_1 \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} t_k$, where for all $m \in \{1 \cdots k\}$:

$$t_m = \begin{cases} \mathbf{R}(\mathbf{s}[\mathbf{x}_n], v_n) & \text{if } t'_m = \mathbf{R}(\mathbf{x}_n, v_n) \\ \mathbf{W}(\mathbf{x}_n, v_n) \xrightarrow{\text{po|imm}} \mathbf{W}(\mathbf{s}[\mathbf{x}_n], v_n) & \text{if } t'_m = \mathbf{W}(\mathbf{x}_n, v_n) \end{cases}$$

such that in the first case the identifier of t_m is that of t'_m ; and in the second case the identifier of the first event in t_m is that of t'_m and the identifier of the second event is picked fresh.

- $Us = wu_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} wu_{\mathbf{x}_n}$, where the identifiers of all constituent events of Us are picked fresh, and

$$wu_{\mathbf{x}_j} = \begin{cases} \mathbf{W}(\mathbf{U}(\mathbf{x}_j)) & \text{if } \mathbf{x}_j \in \mathbf{WS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

We use the ξ . prefix to project the various events of the implementation trace θ_ξ (e.g. $\xi.rl_{\mathbf{x}_j}$).

Given the transaction classes \mathcal{T}/st of G' , we construct a strict total order $\text{TCO} : \mathcal{T}/\text{st} \times \mathcal{T}/\text{st}$ as an extension of $G'.(\text{rb}_T^?; (\text{po}_T \cup \text{mo}_T \cup \text{rf}_T))^+$. That is, construct TCO as a total order such that:

$$\forall e, e'. (e, e') \in G'.(\text{rb}_T^?; (\text{po}_T \cup \text{mo}_T \cup \text{rf}_T))^+ \Rightarrow ([e]_{\text{st}}, [e']_{\text{st}}) \in \text{TCO}$$

Recall that since G' is SI-consistent, we know $\text{acyclic}((\text{po}_T \cup \text{mo}_T \cup \text{rf}_T); \text{rb}_T^?)$ holds, i.e. $\text{acyclic}(\text{rb}_T^?; (\text{po}_T \cup \text{mo}_T \cup \text{rf}_T))$. As such, it is always possible to extend $(\text{rb}_T^?; (\text{po}_T \cup \text{mo}_T \cup \text{rf}_T))^+$ to a total order as described above.

For each location \mathbf{x} , let $WT_{\mathbf{x}} \triangleq \{\mathcal{T}_{\xi} \in \mathcal{T}/\text{st} \mid \mathcal{W}_{\mathbf{x}} \cap \mathcal{T}_{\xi} \neq \emptyset\}$ denote those transactions that write to \mathbf{x} . We then define $\text{TMO}_{\mathbf{x}} \triangleq \text{TCO}|_{WT_{\mathbf{x}}}$ as a strict total order on $WT_{\mathbf{x}}$. Given a strict total order r , we write $r|_i$ for the i^{th} item of r , indexed from 0 (e.g. $\text{TMO}_{\mathbf{x}}|_i$). For $i \in \mathbb{N}$, we then define:

$$\begin{aligned} S_0^{\mathbf{x}} &\triangleq \{\mathcal{T}_{\xi} \mid \text{TMO}_{\mathbf{x}}|_0 = \mathcal{T}_{\xi} \vee \exists r \in \mathcal{T}_{\xi}. (\text{init}_{\mathbf{x}}, r) \in G'.\text{rf}\} \\ S_{i+1}^{\mathbf{x}} &\triangleq \{\mathcal{T}_{\xi} \mid \text{TMO}_{\mathbf{x}}|_{i+1} = \mathcal{T}_{\xi} \vee \exists r \in \mathcal{T}_{\xi}. \exists w \in \text{TMO}_{\mathbf{x}}|_i. (w, r) \in G'.\text{rf} \setminus \text{rf}_I\} \end{aligned}$$

where $\text{init}_{\mathbf{x}}$ denotes the write event initialising the value of \mathbf{x} (with zero).

For all locations \mathbf{x} and $i \in \mathbb{N}$, let $P_i^{\mathbf{x}} \triangleq \text{TCO}|_{S_i^{\mathbf{x}}}$. Note that for each $i \in \mathbb{N}$, the $S_i^{\mathbf{x}}$ contains at most one transaction that writes to \mathbf{x} . We denote this transaction by $P_i^{\mathbf{x}}.\text{writer}$. For each lock \mathbf{x} we then define:

$$\begin{aligned} \text{LO}_{\mathbf{x}} &\triangleq \left\{ \begin{array}{l} (P_i^{\mathbf{x}}.\text{writer}.pl_{\mathbf{x}}, P_i^{\mathbf{x}}.\text{writer}.wu_{\mathbf{x}}), \\ (P_i^{\mathbf{x}}.\text{writer}.pl_{\mathbf{x}}, P_k^{\mathbf{x}}|_j.rl_{\mathbf{x}}), \\ (P_i^{\mathbf{x}}.\text{writer}.pl_{\mathbf{x}}, P_k^{\mathbf{x}}|_j.ru_{\mathbf{x}}), \\ (P_i^{\mathbf{x}}.\text{writer}.pl_{\mathbf{x}}, P_k^{\mathbf{x}}.\text{writer}.pl_{\mathbf{x}}), \\ (P_i^{\mathbf{x}}.\text{writer}.pl_{\mathbf{x}}, P_k^{\mathbf{x}}.\text{writer}.wu_{\mathbf{x}}), \\ (P_i^{\mathbf{x}}.\text{writer}.wu_{\mathbf{x}}, P_k^{\mathbf{x}}|_j.rl_{\mathbf{x}}), \\ (P_i^{\mathbf{x}}.\text{writer}.wu_{\mathbf{x}}, P_k^{\mathbf{x}}|_j.ru_{\mathbf{x}}), \\ (P_i^{\mathbf{x}}.\text{writer}.wu_{\mathbf{x}}, P_k^{\mathbf{x}}.\text{writer}.pl_{\mathbf{x}}), \\ (P_i^{\mathbf{x}}.\text{writer}.wu_{\mathbf{x}}, P_k^{\mathbf{x}}.\text{writer}.wu_{\mathbf{x}}) \end{array} \mid i \in \mathbb{N} \wedge k > i \wedge 0 \leq j < |S_k^{\mathbf{x}}| \right\} \\ &\cup \left\{ \begin{array}{l} (P_i^{\mathbf{x}}|_j.rl_{\mathbf{x}}, P_k^{\mathbf{x}}.\text{writer}.pl_{\mathbf{x}}), \\ (P_i^{\mathbf{x}}|_j.ru_{\mathbf{x}}, P_k^{\mathbf{x}}.\text{writer}.pl_{\mathbf{x}}), \\ (P_i^{\mathbf{x}}|_j.rl_{\mathbf{x}}, P_k^{\mathbf{x}}.\text{writer}.wu_{\mathbf{x}}), \\ (P_i^{\mathbf{x}}|_j.ru_{\mathbf{x}}, P_k^{\mathbf{x}}.\text{writer}.wu_{\mathbf{x}}), \end{array} \mid i \in \mathbb{N} \wedge k \geq i \wedge 0 \leq j < |S_i^{\mathbf{x}}| \right\} \\ \text{RLO}_{\mathbf{x}} &\triangleq \text{LO}_{\mathbf{x}} \cup \left\{ \begin{array}{l} (P_i^{\mathbf{x}}|_j.rl_{\mathbf{x}}, P_i^{\mathbf{x}}|_j.ru_{\mathbf{x}}), \\ (P_i^{\mathbf{x}}|_j.rl_{\mathbf{x}}, P_{i'}^{\mathbf{x}}|_{j'}.rl_{\mathbf{x}}), \\ (P_i^{\mathbf{x}}|_j.rl_{\mathbf{x}}, P_{i'}^{\mathbf{x}}|_{j'}.ru_{\mathbf{x}}), \\ (P_i^{\mathbf{x}}|_j.ru_{\mathbf{x}}, P_{i'}^{\mathbf{x}}|_{j'}.rl_{\mathbf{x}}), \\ (P_i^{\mathbf{x}}|_j.ru_{\mathbf{x}}, P_{i'}^{\mathbf{x}}|_{j'}.ru_{\mathbf{x}}), \end{array} \mid \begin{array}{l} i, i' \in \mathbb{N} \wedge 0 \leq j < |S_i^{\mathbf{x}}| \wedge 0 \leq j' < |S_{i'}^{\mathbf{x}}| \\ \wedge ((i = i' \wedge j < j') \vee (i < i')) \end{array} \right\} \end{aligned}$$

Remark 4. Let $\text{lo}_1 \triangleq \bigcup_{\mathbf{x} \in \text{LOC}} \text{LO}_{\mathbf{x}}$ and $\text{lo}_2 \triangleq \bigcup_{\mathbf{x} \in \text{LOC}} \text{RLO}_{\mathbf{x}}$. Note that both lo_1 and lo_2 satisfy the conditions stated in [Def. 4](#). The lo_2 additionally satisfies the

‘read-read-synchronisation’ property in (RSYNC). In what follows we demonstrate that given an SI-consistent execution graph G' , it is always possible to construct an RA-consistent execution graph G of the implementation with its lock order defined as lo_2 . Note that as $\text{lo}_1 \subseteq \text{lo}_2$, it is straightforward to show that replacing lo_2 in such a G with lo_1 , preserves the RA-consistency of G , as defined in Def. 4. In other words, as $\text{lo}_1 \subseteq \text{lo}_2$, we have:

$$\text{acyclic}(\text{hb}_{loc} \cup \text{mo} \cup \text{rb}) \text{ with } \text{hb} \triangleq (\text{po} \cup \text{rf} \cup \text{lo}_2)^+ \Rightarrow \\ \text{acyclic}(\text{hb}_{loc} \cup \text{mo} \cup \text{rb}) \text{ with } \text{hb} \triangleq (\text{po} \cup \text{rf} \cup \text{lo}_1)^+$$

As such, by establishing the completeness of our implementation with respect to lo_2 , we also establish its completeness with respect to lo_1 . In other words, we demonstrate the completeness of our implementation with respect to both lock implementations presented earlier in Appendix A.

We now demonstrate the completeness of our implementation. Given an SI-consistent graph, we construct an implementation graph G as follows and demonstrate that it is RA-consistent.

- $G.E = \bigcup_{\mathcal{T}_\xi \in G'.\mathcal{T}/\text{st}} \theta_\xi.E$, with the $\text{tx}(e) = 0$, for all $e \in G.E$.
Observe that $G'.E \subseteq G.E$.
- $G.\text{po}$ is defined as $G'.\text{po}$ extended by the po for the additional events of G , given by each θ_ξ trace defined above. Note that $G.\text{po}$ does not introduce additional orderings between events of $G'.E$. That is, $\forall a, b \in G'.E. (a, b) \in G'.\text{po} \Leftrightarrow (a, b) \in G.\text{po}$.
- $G.\text{rf} = \bigcup_{\mathbf{x} \in \text{Locs}} \{(w, \xi.r\text{s}_\mathbf{x}) \mid \exists r. (\mathbf{x}, r) \in \text{RS}_\xi \wedge (w, r) \in G'.\text{rf}\}$.
- $G.\text{mo} = G'.\text{mo}$
- $G.\text{lo} = \bigcup_{\mathbf{x} \in \text{Loc}} \text{RLO}_\mathbf{x}$, with $\text{RLO}_\mathbf{x}$ as defined above.

Notation Given an implementation graph G as constructed above (with $G.\mathcal{T} = \emptyset$), and a relation $r \subseteq G.E \times G.E$, we override the r_E notation and write r_E for:

$$\{(a, b) \in r \mid \exists \xi_a, \xi_b. a \in \theta_{\xi_a}.E \wedge b \in \theta_{\xi_b}.E \wedge \xi_a \neq \xi_b\}$$

Analogously, we write r_I for $\{(a, b) \in r \mid \exists \xi. a, b \in \theta_\xi.E\}$.

Lemma 5. *For all SI-consistent execution graphs $G' = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$, and for all $\mathbf{x} \in \text{Loc}$, $\xi_a, \xi_b \in \text{TXID}$, and $i \in \mathbb{N}$:*

$$\xi_a \in S_i^\mathbf{x} \wedge \xi_b \in S_{i+1}^\mathbf{x} \Rightarrow (\xi_a, \xi_b) \in \text{TCO}$$

Proof. Pick arbitrary SI-consistent execution graph $G' = (E, \text{po}, \text{rf}, \text{mo}, \mathcal{T})$, $\mathbf{x} \in \text{Loc}$, $\xi_a, \xi_b \in \text{TXID}$, and $i \in \mathbb{N}$. There are now four cases to consider: 1) $\xi_a = P_i^\mathbf{x}.\text{writer}$ and $\xi_b = P_{i+1}^\mathbf{x}.\text{writer}$; 2) $\xi_a = P_i^\mathbf{x}.\text{writer}$ and $\xi_b \neq P_{i+1}^\mathbf{x}.\text{writer}$; 3) $\xi_a \neq P_i^\mathbf{x}.\text{writer}$ and $\xi_b = P_{i+1}^\mathbf{x}.\text{writer}$; and 4) $\xi_a \neq P_i^\mathbf{x}.\text{writer}$ and $\xi_b \neq P_{i+1}^\mathbf{x}.\text{writer}$.

In case (1) from the definitions of $S_i^\mathbf{x}$ and $P_i^\mathbf{x}.\text{writer}$ we have $\mathcal{T}_{\xi_a} \times \mathcal{T}_{\xi_b} \subseteq \text{mo}_\mathbf{T}$ and thus from the definition of TCO we have $(\xi_a, \xi_b) \in \text{TCO}$, as required. Similarly,

in case (2) from the definitions of S_i^x and $P_i^x.\text{writer}$ we have $\mathcal{T}_{\xi_a} \times \mathcal{T}_{\xi_b} \subseteq \text{rf}_T$ and thus from the definition of TCO we have $(\xi_a, \xi_b) \in \text{TCO}$, as required.

In case (3), from the definition of S_i^x we know there exist $r \in \mathcal{T}_{\xi_a} \cap \mathcal{R}_x$ and $w \in P_i^x.\text{writer} \cap \mathcal{W}_x$ such that $(r, w) \in \text{rb}$ and thus $(a, w) \in \text{rb}_T$. On the other hand from the definitions of S_i^x and S_{i+1}^x we know there exists $w' \in \mathcal{T}_{\xi_b} \cap \mathcal{W}_x$ such that $(w, w') \in \text{mo}$ and thus $(w, w') \in \text{mo}_T$. As such we have $(a, w') \in \text{rb}_T; \text{mo}_T$ and thus from the definition of TCO we have $(\xi_a, \xi_b) \in \text{TCO}$, as required.

In case (4), from the definition of S_i^x we know there exist $r \in \mathcal{T}_{\xi_a} \cap \mathcal{R}_x$ and $w \in P_i^x.\text{writer} \cap \mathcal{W}_x$ such that $(r, w) \in \text{rb}$ and thus $(a, w) \in \text{rb}_T$. On the other hand from the definitions of S_i^x and S_{i+1}^x we know there exists $r' \in \mathcal{T}_{\xi_b} \cap \mathcal{R}_x$ such that $(w, r') \in \text{rf}$ and thus $(w, r') \in \text{rf}_T$. As such we have $(a, r') \in \text{rb}_T; \text{rf}_T$ and thus from the definition of TCO we have $(\xi_a, \xi_b) \in \text{TCO}$, as required.

Corollary 1. For all SI-consistent execution graphs $G' = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$, and for all $x \in \text{Loc}$, $\xi_a, \xi_b \in \text{TXID}$, and $i, j \in \mathbb{N}$:

$$\xi_a \in S_i^x \wedge \xi_b \in S_j^x \wedge i < j \Rightarrow (\xi_a, \xi_b) \in \text{TCO}$$

Proof. Follows by induction from [Lemma 5](#).

Given an implementation graph $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$, let

$$\text{hb}' \triangleq \left(\text{po} \cup \text{rf} \cup \text{lo} \setminus \left\{ \begin{array}{l} (\xi_1.rl_x, \xi_2.pl_x) \in \text{lo}_x, \\ (\xi_1.ru_x, \xi_2.pl_x) \in \text{lo}_x, \\ (\xi_1.rl_x, \xi_2.wu_x) \in \text{lo}_x, \\ (\xi_1.ru_x, \xi_2.wu_x) \in \text{lo}_x \end{array} \middle| \exists i, j, k. P_i^x|_j = \xi_1 \wedge P_i^x|_k = \xi_2 \wedge j > k \right\} \right)^+$$

Lemma 6. For all implementation graphs $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ constructed as above,

$$\begin{aligned} \forall a, b, \xi_a, \xi_b. a \in \theta_{\xi_a}.E \wedge b \in \theta_{\xi_b}.E \wedge a \xrightarrow{\text{hb}'} b \\ \Rightarrow (a, b) \in G.\text{po}_I \vee (\xi_a, \xi_b) \in \text{TCO} \end{aligned}$$

Proof. Pick an arbitrary implementation graph $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ constructed as above. Since hb' is a transitive closure, it is straightforward to demonstrate that $\text{hb}' = \bigcup_{i \in \mathbb{N}} \text{hb}'_i$, where:

$$\text{hb}'_0 = \text{po} \cup \text{rf} \cup \text{lo} \setminus \left\{ \begin{array}{l} (\xi_1.rl_x, \xi_2.pl_x) \in \text{lo}_x, \\ (\xi_1.ru_x, \xi_2.pl_x) \in \text{lo}_x, \\ (\xi_1.rl_x, \xi_2.wu_x) \in \text{lo}_x, \\ (\xi_1.ru_x, \xi_2.wu_x) \in \text{lo}_x \end{array} \middle| \exists i, j, k. P_i^x|_j = \xi_1 \wedge P_i^x|_k = \xi_2 \wedge j > k \right\}$$

and $\text{hb}'_{i+1} = \text{hb}'_0; \text{hb}'_i$. It thus suffices to show:

$$\begin{aligned} \forall i \in \mathbb{N}. \forall a, b, \xi_a, \xi_b. a \in \theta_{\xi_a}.E \wedge b \in \theta_{\xi_b}.E \wedge a \xrightarrow{\text{hb}'_i} b \\ \Rightarrow (a, b) \in G.\text{po}_I \vee (\xi_a, \xi_b) \in \text{TCO} \end{aligned}$$

We proceed by induction over i .

Base case $i = 0$

Pick arbitrary a, b, ξ_a, ξ_b such that $a \in \theta_{\xi_a}.E$ and $b \in \theta_{\xi_b}.E$, $a \xrightarrow{\text{hb}'_0} b$. There are then four cases to consider: 1) $a \xrightarrow{G.\text{po}_1 \cup \text{rf}_1 \cup \text{lo}_1} b$; or 2) $a \xrightarrow{G.\text{poE}} b$; 3) $a \xrightarrow{G.\text{rfE}} b$; or 4) $a \xrightarrow{G.\text{loE}} b$.

In case (1), from the construction of $G.\text{rf}_1$ we have $G.\text{rf}_1 \subseteq G.\text{po}_1$; moreover, from the construction of $G.\text{lo}_1$ we have $G.\text{lo}_1 \subseteq G.\text{po}_1$. Consequently we have $a \xrightarrow{G.\text{po}_1} b$, as required.

In case (2) from the construction of $G.\text{po}$ we then know that $\mathcal{T}_{\xi_a} \times \mathcal{T}_{\xi_b} \subseteq G'.\text{po}$ and thus from the definition of TCO we have $(\xi_a, \xi_b) \in \text{TCO}$, as required.

In case (3) from the construction of $G.\text{rf}$ we then know that $\mathcal{T}_{\xi_a} \times \mathcal{T}_{\xi_b} \subseteq G'.\text{rf}_\top$ and thus from the definition of TCO we have $(\xi_a, \xi_b) \in \text{TCO}$, as required.

In case (4) from the definition of hb'_0 we know that there exists i, j, i', j' such that $\xi_a = P_i^x|_j$, $\xi_b = P_{i'}^x|_{j'}$, and either i) $i < i'$; or ii) $i = i' \wedge j < j'$.

In case (4.i) from [Corollary 1](#) we have $(\xi_a, \xi_b) \in \text{TCO}$, as required. In case (4.ii) from the definitions of P_i^x we have $(\xi_a, \xi_b) \in \text{TCO}$, as required.

Inductive case $i = n+1$

Pick arbitrary a, b, ξ_a, ξ_b such that $a \in \theta_{\xi_a}.E$, $b \in \theta_{\xi_b}.E$ and $a \xrightarrow{\text{hb}'_i} b$.

$$\begin{aligned} \forall j \leq n. \forall c, d, \xi_c, \xi_d. c \in \theta_{\xi_c}.E \wedge d \in \theta_{\xi_d}.E \wedge c \xrightarrow{\text{hb}'_j} d \\ \Rightarrow (c, d) \in G.\text{po}_1 \vee (\xi_c, \xi_d) \in \text{TCO} \end{aligned} \quad (\text{I.H.})$$

From the definition of hb'_{n+1} we then know there exists e such that $(a, e) \in \text{hb}'_0$ and $(e, b) \in \text{hb}'_n$. Let $e \in \xi_e.E$. Consequently, from the proof of the base case we then know that $(a, e) \in G.\text{po}_1 \vee (\xi_a, \xi_e) \in \text{TCO}$. Similarly, from (I.H.) we have $(e, b) \in G.\text{po}_1 \vee (\xi_e, \xi_b) \in \text{TCO}$. There are now four cases to consider: 1) $(a, e) \in G.\text{po}_1$ and $(e, b) \in G.\text{po}_1$; 2) $(a, e) \in G.\text{po}_1$ and $(\xi_e, \xi_b) \in \text{TCO}$; 3) $(\xi_a, \xi_e) \in \text{TCO}$ and $(e, b) \in G.\text{po}_1$; 4) $(\xi_a, \xi_e) \in \text{TCO}$ and $(\xi_e, \xi_b) \in \text{TCO}$.

In case (1) from the definition of $G.\text{po}_1$ we have $(a, b) \in G.\text{po}_1$. In cases (2-4) from the definitions of $G.\text{po}_1$, TCO we have $(\xi_a, \xi_b) \in \text{TCO}$, as required.

Given an implementation graph $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ constructed as above, let $\text{stg}(\cdot) : G.E \rightarrow \{1, 2, 3\}$ denote the *event stage* reflecting whether it is in the snapshot phase (1), reader lock release phase (2), or the update phase (3). That is,

$$\text{stg}(e) \triangleq \begin{cases} 1 & \exists \xi, x. e \in \{\xi.\text{rl}_x, \xi.\text{rs}_x, \xi.\text{ws}_x\} \\ 2 & \exists \xi, x. e = \xi.\text{ru}_x \\ 3 & \text{otherwise} \end{cases}$$

Lemma 7. For all implementation graphs $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ constructed as above, and for all $\mathbf{x} \in \text{LOC}$, $\xi_a, \xi_1, \xi_2 \in \text{TXID}$ and $a \in \theta_{\xi_a}.E$:

$$\left(\begin{array}{l} (\xi_1.rl_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}} a) \\ \vee (\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}} a) \\ \vee (\xi_1.rl_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.wu_{\mathbf{x}} \wedge \xi_2.wu_{\mathbf{x}} \xrightarrow{\text{hb}} a) \\ \vee (\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.wu_{\mathbf{x}} \wedge \xi_2.wu_{\mathbf{x}} \xrightarrow{\text{hb}} a) \end{array} \right) \wedge \exists j, k, l. P_j^{\mathbf{x}}|_k = \xi_1 \wedge P_j^{\mathbf{x}}|_l = \xi_2 \wedge k > l \\ \Rightarrow \text{stg}(\xi_1.ru_{\mathbf{x}}) < \text{stg}(a) \vee (\xi_1, \xi_a) \in \text{TCO}$$

Proof. Pick an arbitrary implementation graph $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ constructed as above, and pick arbitrary $\mathbf{x} \in \text{LOC}$, $\xi_a, \xi_1, \xi_2 \in \text{TXID}$ and $a \in \theta_{\xi_a}.E$ such that $((\xi_1.rl_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}} a) \vee (\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}} a) \vee (\xi_1.rl_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.wu_{\mathbf{x}} \wedge \xi_2.wu_{\mathbf{x}} \xrightarrow{\text{hb}} a) \vee (\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.wu_{\mathbf{x}} \wedge \xi_2.wu_{\mathbf{x}} \xrightarrow{\text{hb}} a))$ and $\exists j, k, l. P_j^{\mathbf{x}}|_k = \xi_1 \wedge P_j^{\mathbf{x}}|_l = \xi_2 \wedge k > l$. In case of the first disjunct from the construction of **lo** we then also have $\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}} a$. Similarly, in case of the third and fourth disjuncts from the construction of **lo** we also have $\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.wu_{\mathbf{x}} \xrightarrow{\text{hb}} a$. Moreover, since we have $\xi_2.pl_{\mathbf{x}} \xrightarrow{\text{po}} \xi_2.wu_{\mathbf{x}}$ and $\text{po} \subseteq \text{hb}$, we also have $\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}} a$. It thus suffices to show:

$$\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}} a \wedge \exists j, k, l. P_j^{\mathbf{x}}|_k = \xi_1 \wedge P_j^{\mathbf{x}}|_l = \xi_2 \wedge k > l \\ \Rightarrow \text{stg}(\xi_1.ru_{\mathbf{x}}) < \text{stg}(a) \vee (\xi_1, \xi_a) \in \text{TCO}$$

Since $G.\text{hb}$ is a transitive closure, it is straightforward to demonstrate that $G.\text{hb} = \bigcup_{i \in \mathbb{N}} \text{hb}_i$, where $\text{hb}_0 = G.\text{po} \cup G.\text{rf} \cup G.\text{lo}$ and $\text{hb}_{i+1} = \text{hb}_0; \text{hb}_i$. It thus suffices to show:

$$\forall i \in \mathbb{N}. \forall \mathbf{x} \in \text{LOC}, \xi_a, \xi_1, \xi_2 \in \text{TXID}, a \in \theta_{\xi_a}.E. \\ \xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}_i} a \wedge \exists j, k, l. P_j^{\mathbf{x}}|_k = \xi_1 \wedge P_j^{\mathbf{x}}|_l = \xi_2 \wedge k > l \\ \Rightarrow \text{stg}(\xi_1.ru_{\mathbf{x}}) < \text{stg}(a) \vee (\xi_1, \xi_a) \in \text{TCO}$$

We thus proceed by induction over i .

Base case $i = 0$

Pick arbitrary $\mathbf{x} \in \text{LOC}$, $\xi_a, \xi_1, \xi_2 \in \text{TXID}$ and $a \in \theta_{\xi_a}.E$ such that $\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}}$, $\xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}_0} a$ and $\exists j, k, l. P_j^{\mathbf{x}}|_k = \xi_1 \wedge P_j^{\mathbf{x}}|_l = \xi_2 \wedge k > l$. From the construction of G we know that there exists $r \in G'.\mathcal{T}_{\xi_1}$ and $w \in G'.\mathcal{T}_{\xi_2}$ such that $(r, w) \in G'.\text{rb}$. Moreover, since $k \neq l$ we know that $(r, w) \in G'.\text{rb}_{\top}$. Since we have $\xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}_0} a$, there are four cases to consider: 1) $\xi_2.pl_{\mathbf{x}} \xrightarrow{G.\text{po}_1 \cup \text{rf}_1 \cup \text{lo}_1} a$; or 2) $\xi_2.pl_{\mathbf{x}} \xrightarrow{G.\text{po}_E} a$; or 3) $\xi_2.pl_{\mathbf{x}} \xrightarrow{G.\text{rf}_E} a$; or 4) $\xi_2.pl_{\mathbf{x}} \xrightarrow{G.\text{lo}_E} a$.

In case (1), from the construction of $G.\text{rf}_1$ we have $G.\text{rf}_1 \subseteq G.\text{po}_1$; moreover, from the construction of $G.\text{lo}_1$ we have $G.\text{lo}_1 \subseteq G.\text{po}_1$. We thus have $\xi_2.pl_{\mathbf{x}} \xrightarrow{G.\text{po}_1} a$.

On the other hand, from the definition of $\text{stg}(\cdot)$ we know $\text{stg}(a) \geq \text{stg}(\xi_2.pl_x) > \text{stg}(\xi_1.ru_x)$, and thus $\text{stg}(a) > \text{stg}(\xi_1.ru_x)$, as required.

In case (2), since $\xi_2.pl_x \xrightarrow{\text{poE}} a$, we have $(w, a) \in G'.\text{po}_T$. As such, we have $(r, a) \in G'.(\text{rb}_T; \text{po}_T)$. From the definition of TCO we thus have $(\xi_1, \xi_a) \in \text{TCO}$, as required.

Case (3) cannot happen as there are no **rf** edges from lock events. In case (4), from the construction of **lo** we know there exists m, n such that $m > j$, and $\xi_a = P_m^x|_n$. As such, from *Corollary 1* we have $(\xi_1, \xi_a) \in \text{TCO}$, as required.

Inductive case $i = n+1$

Pick arbitrary $x \in \text{LOC}$ and $\xi_a, \xi_1, \xi_2 \in \text{TXID}$ and $a \in \theta_{\xi_a}.E$ such that $\xi_1.ru_x \xrightarrow{\text{lo}} \xi_2.pl_x$, $\xi_2.pl_x \xrightarrow{\text{hb}_j} a$ and $\exists j, k, l. P_j^x|_k = \xi_1 \wedge P_j^x|_l = \xi_2 \wedge k > l$.

$\forall j \leq n. \forall y \in \text{LOC}, \xi_{a'}, \xi_3, \xi_4 \in \text{TXID}, a' \in \theta_{\xi_{a'}}.E$.

$$\begin{aligned} (\xi_3.ru_x \xrightarrow{\text{lo}} \xi_4.pl_x \wedge \xi_4.pl_x \xrightarrow{\text{hb}_j} a') \wedge \exists j, k, l. P_j^x|_k = \xi_1 \wedge P_j^x|_l = \xi_2 \wedge k > l \\ \Rightarrow \text{stg}(\xi_3.ru_y) < \text{stg}(a') \vee (\xi_3, \xi_{a'}) \in \text{TCO} \end{aligned} \quad (\text{I.H.})$$

There are two cases to consider:

- 1) $\xi_2.pl_x \xrightarrow{(\text{po}_1 \cup \text{rf}_1 \cup \text{lo}_1)^+} a$; or
- 2) $\xi_2.pl_x \xrightarrow{(\text{po}_1 \cup \text{rf}_1 \cup \text{lo}_1)^*; (\text{po}_E \cup \text{rf}_E \cup \text{lo}_E); \text{hb}_m} a$, where $m \leq n$.

In case (1), from the construction of $G.\text{rf}_I$ we have $G.\text{rf}_I \subseteq G.\text{po}_I$; moreover, from the construction of $G.\text{lo}_I$ we have $G.\text{lo}_I \subseteq G.\text{po}_I$. We thus have $\xi_2.pl_x \xrightarrow{\text{po}_I^+} a$, i.e. $\xi_2.pl_x \xrightarrow{\text{po}_I} a$. As such, from the proof of the base case we have $\text{stg}(a) > \text{stg}(\xi_1.ru_x)$, as required.

In case (2) from the construction of G we know that there exists $r \in \theta_{\xi_1}.E$ and $w \in \theta_{\xi_2}.E$ such that $(r, w) \in G'.\text{rb}$. Moreover, since $k \neq l$, from the construction of G we know that $\xi_1 \neq \xi_2$ and thus we have $(r, w) \in G'.\text{rb}_T$. On the other hand, we know that there exist b, c such that $\xi_2.pl_x \xrightarrow{(\text{po}_1 \cup \text{rf}_1 \cup \text{lo}_1)^*} b \xrightarrow{\text{po}_E \cup \text{rf}_E \cup \text{lo}_E} c \xrightarrow{\text{hb}_m} a$. From the construction of $G.\text{rf}_I$ we have $G.\text{rf}_I \subseteq G.\text{po}_I$; moreover, from the construction of $G.\text{lo}_I$ we have $G.\text{lo}_I \subseteq G.\text{po}_I$. We thus have $\xi_2.pl_x \xrightarrow{\text{po}_I^*} b$. As such we have $\xi_2.pl_x \xrightarrow{\text{po}_I^*} b \xrightarrow{\text{po}_E \cup \text{rf}_E \cup \text{lo}_E} c$. Let $c \in \theta_{\xi_c}$. There are now three cases to consider: a) $b \xrightarrow{\text{po}_E} c$; or b) $b \xrightarrow{\text{rf}_E} c$; or c) $b \xrightarrow{\text{lo}_E} c$. We first demonstrate that in all three cases we have $(\xi_1, \xi_c) \in \text{TCO}$.

In case (2.a), since $\xi_2.pl_x \xrightarrow{\text{po}_E} c$, we know there exists $c' \in \mathcal{T}_{\xi_c}$ such that $(w, c') \in G'.\text{po}_T$. As such, we have $(r, c') \in G'.(\text{rb}_T; \text{po}_T)$. From the definition of TCO we thus have $(\xi_1, \xi_c) \in \text{TCO}$. In case (2.b), from the definition of $G.\text{rf}$ we know there exists $c' \in \mathcal{T}_{\xi_c}$ such that $(w, c') \in G'.\text{rf}_T$. As such, we have $(r, c') \in G'.(\text{rb}_T; \text{rf}_T)$. From the definition of TCO we thus have $(\xi_1, \xi_c) \in \text{TCO}$. In case (2.c) from the construction of **lo** we know there exists z such that either $b = \xi_2.pl_z$ or $b = \xi_2.wu_z$, and that there exist p, q such that $\xi_2 \in S_p^z$, $\xi_c \in S_q^z$ and $p < q$. Given the definition of S^z we then know that there exists $w_z \in G'.\mathcal{T}_{\xi_2}$ where either i) there exists $w'_z \in \mathcal{T}_{\xi_c}$ such that $(w_z, w'_z) \in G'.\text{mo}_T$; or ii) there exists $r_z \in \mathcal{T}_{\xi_c}$

such that $(w_z, r_z) \in G'.\mathbf{rf}_T$; or ii) there exists $r_z \in \mathcal{T}_{\xi_c}$ such that $(w_z, r_z) \in G'.(\mathbf{mo}_T; \mathbf{rf}_T)$. That is, we have either $(w, w'_z) \in G'.\mathbf{mo}_T$, or $(w, r_z) \in G'.\mathbf{rf}_T$, or $(w, r_z) \in G'.(\mathbf{mo}_T; \mathbf{rf}_T)$. Moreover, since we have $(r, w) \in G'.\mathbf{rb}_T$, we then have $(r, w'_z) \in G'.(\mathbf{rb}_T; \mathbf{mo}_T)$, or $(r, r_z) \in G'.(\mathbf{rb}_T; \mathbf{rf}_T)$, or $(r, r_z) \in G'.(\mathbf{rb}_T; \mathbf{mo}_T; \mathbf{rf}_T)$. From the definition of TCO we thus have $(\xi_1, \xi_c) \in \text{TCO}$.

Since $c \xrightarrow{\mathbf{hb}^n} a$, there are now two cases to consider: i) $c \xrightarrow{\mathbf{hb}'} a$; or ii) $c \not\xrightarrow{\mathbf{hb}'} a$.

In case (2.i), from Lemma 6 we have $(c, a) \in \text{po}_I \vee (\xi_c, \xi_a) \in \text{TCO}$. As we have $(\xi_1, \xi_c) \in \text{TCO}$, we thus have $(\xi_1, \xi_a) \in \text{TCO}$, as required.

In case (2.ii) let us split the path from c at the first occurrence of a non- \mathbf{hb}' edge. That is, pick $\xi_3, \xi_4, g, h, y, p, q, s, k$ such that $c \xrightarrow{\mathbf{hb}'} \xi_3.g \xrightarrow{\mathbf{lo}} \xi_4.h \xrightarrow{\mathbf{hb}^k} a$, $k < m$, $P_s^y|_p = \mathcal{T}_{\xi_3}$, $P_s^y|_q = \mathcal{T}_{\xi_4}$, $p > q$ and either a) $g = rly \wedge h = pl_y$; or b) $g = ru_y \wedge h = pl_y$; or c) $g = rly \wedge h = wu_y$; or d) $g = ru_y \wedge h = wu_y$. From Lemma 6 we then have $(c, \xi_3.g) \in \text{po}_I \vee (\xi_c, \xi_3) \in \text{TCO}$. As we have $(\xi_1, \xi_c) \in \text{TCO}$, we also have $(\xi_1, \xi_3) \in \text{TCO}$. We next demonstrate that in all cases (2.ii.a-2.ii.d) there exists t such that $t \leq m$ and $\xi_3.ru_y \xrightarrow{\mathbf{lo}} \xi_4.pl_y \xrightarrow{\mathbf{hb}^t} a$.

In case (2.ii.a) from the definition of \mathbf{lo} we also have $\xi_3.ru_y \xrightarrow{\mathbf{lo}} \xi_4.pl_x$. As such, we have $\xi_3.ru_y \xrightarrow{\mathbf{lo}} \xi_4.pl_y \xrightarrow{\mathbf{hb}^k} a$. In case (2.ii.b) the desired result holds immediately.

In cases (2.ii.c-2.ii.d) from the construction of \mathbf{lo} we have $\xi_3.ru_y \xrightarrow{\mathbf{lo}} \xi_4.pl_y$. Moreover, since we have $\xi_4.pl_y \xrightarrow{\mathbf{po}} \xi_4.wu_y$, we also have $\xi_4.pl_y \xrightarrow{\mathbf{po}} \xi_4.wu_y \xrightarrow{\mathbf{hb}^k} a$. As $\text{po} \subseteq \mathbf{hb}$ and \mathbf{hb} is transitively closed, we have $\xi_4.pl_y \xrightarrow{\mathbf{hb}^{k+1}} a$. As such, we have $\xi_3.ru_y \xrightarrow{\mathbf{lo}} \xi_4.pl_y \xrightarrow{\mathbf{hb}^{k+1}} a$. As $k < m$, the desired result holds immediately.

Consequently, from (I.H.) we have $\text{stg}(\xi_3.ru_y) < \text{stg}(a) \vee (\xi_3, \xi_a) \in \text{TCO}$. In the case of the first disjunct we have $\text{stg}(\xi_3.ru_y) = \text{stg}(\xi_1.ru_x) < \text{stg}(a)$, as required. In the case of the second disjunct, since we also have $(\xi_1, \xi_3) \in \text{TCO}$ and TCO is transitively closed, we have $(\xi_1, \xi_a) \in \text{TCO}$, as required.

Theorem 8 (Completeness). For all SI execution graphs G' and their counterpart implementation graphs G constructed as above,

$$\text{si-consistent}(G') \Rightarrow \text{RA-consistent}(G)$$

Proof. Pick an arbitrary SI execution graph G' and its counterpart implementation graph G constructed as above and assume $\text{si-consistent}(G')$ holds. From the definition of $\text{RA-consistent}(G)$ it suffices to show:

1. $\text{irreflexive}(G.\mathbf{hb})$
2. $\text{irreflexive}(G.\mathbf{mo}; G.\mathbf{hb})$
3. $\text{irreflexive}(G.\mathbf{rb}; G.\mathbf{hb})$

RTS. part 1

We proceed by contradiction. Let us assume that there exists a, θ_ξ such that $a \in \theta_\xi.E$ and $(a, a) \in G.\mathbf{hb}$. There are now two cases to consider: 1) $(a, a) \in \mathbf{hb}'$; or 2) $(a, a) \notin \mathbf{hb}'$.

In case (1), from [Lemma 6](#) we have $(a, a) \in G.\text{po}_I \vee (\mathcal{T}_\xi, \mathcal{T}_\xi) \in \text{TCO}$. The first disjunct leads to a contradiction as the construction of $G.\text{po}$ yields an acyclic relation. The second disjunct leads to a contradiction as TCO is a strict total order.

In case (2), let us split the $a \xrightarrow{\text{hb}} a$ at the first occurrence of a non- hb' edge. That is, pick $\xi_1, \xi_2, \mathbf{x}, i, j, k, g, h$ such that $a \xrightarrow{\text{hb}} \xi_1.g \xrightarrow{\text{lo}\backslash\text{hb}' } \xi_2.h \xrightarrow{\text{hb}} a$, $P_i^x|_j = \mathcal{T}_{\xi_1}$, $P_i^x|_k = \mathcal{T}_{\xi_2}$ and $j > k$. As we have $\xi_1.g \xrightarrow{\text{lo}\backslash\text{hb}' } \xi_2.h \xrightarrow{\text{hb}} a \xrightarrow{\text{hb}} \xi_1.g$, from [Lemma 7](#) and the definition of hb we then have $\text{stg}(\xi_1.g) < \xi_1.g \vee (\xi_1, \xi_1) \in \text{TCO}$, leading to a contradiction in both disjuncts (the second disjunct yields a contradiction as TCO is a strict total order).

RTS. part 2

We proceed by contradiction. Let us assume that there exists $a, \theta_{\xi_a}, b, \theta_{\xi_b}$ such that $a \in \theta_{\xi_a}.E$, $b \in \theta_{\xi_b}.E$, $(a, b) \in G.\text{hb}$ and $(b, a) \in G.\text{mo}$. Let $\text{loc}(a) = \text{loc}(b) = \mathbf{x}$ for some shared location \mathbf{x} . There are now two cases to consider: 1) $(b, a) \in G.\text{mo}_I$; or 2) $(b, a) \in G.\text{mo}_E$.

In case (1) we then have $(b, a) \in G'.\text{mo}_I \subseteq G'.\text{po}_I$. That is, we have $(b, a) \in G.\text{po} \subseteq G.\text{hb}$. We thus have $a \xrightarrow{G.\text{hb}} b \xrightarrow{G.\text{hb}} a$, contradicting our proof in part 1. In case (2), from the construction of $G.\text{mo}$ we have $(b, a) \in G'.\text{mo}_T$ and thus from the construction of $G.\text{lo}$ we then have $(\xi_b.wu_{\mathbf{x}}, \xi_a.rl_{\mathbf{x}}) \in G.\text{lo}$. As such we have $a \xrightarrow{G.\text{hb}} b \xrightarrow{G.\text{po}} \xi_b.wu_{\mathbf{x}} \xrightarrow{G.\text{lo}} \xi_a.rl_{\mathbf{x}} \xrightarrow{G.\text{po}} a$. That is, we have $a \xrightarrow{G.\text{hb}} a$, contradicting our proof in part 1.

RTS. part 3

We proceed by contradiction. Let us assume that there exists $a, \theta_{\xi_a}, b, \theta_{\xi_b}$ such that $a \in \theta_{\xi_a}.E$, $b \in \theta_{\xi_b}.E$, $(a, b) \in G.\text{hb}$ and $(b, a) \in G.\text{rb}$.

Let $\text{loc}(a) = \text{loc}(b) = \mathbf{x}$ for some shared location \mathbf{x} . There are now two cases to consider: 1) $(b, a) \in G.\text{rb}_I$; or 2) $(b, a) \in G.\text{rb}_E$.

In case (1) we then have $(b, a) \in G'.\text{rb}_I \subseteq G'.\text{po}_I$. That is, we have $(b, a) \in G.\text{po} \subseteq G.\text{hb}$. We thus have $a \xrightarrow{G.\text{hb}} b \xrightarrow{G.\text{hb}} a$, contradicting our proof in part 1. In case (2), from the construction of $G.\text{rb}$ we have $(b, a) \in G'.\text{rb}_T$ and thus from the construction of $G.\text{lo}$ we then have either $(b, \xi_b.wu_{\mathbf{x}}) \in G.\text{po}$ and $(\xi_b.wu_{\mathbf{x}}, \xi_a.rl_{\mathbf{x}}) \in G.\text{lo}$; or $(b, \xi_b.ru_{\mathbf{x}}) \in \text{po}$ and $(\xi_b.ru_{\mathbf{x}}, \xi_a.pl_{\mathbf{x}}) \in G.\text{lo}$. As such in both cases we have $(b, a) \in G.\text{hb}$. Consequently, we have $a \xrightarrow{G.\text{hb}} b \xrightarrow{G.\text{hb}} a$, contradicting our proof in part 1.

D Soundness and Completeness of the Lazy SI Implementation

Given an execution graph G of the lazy SI implementation, let us assign a transaction identifier to each transaction executed by the program; and given a transaction ξ . Let $RS_\xi^0 = WS_\xi^0 = \emptyset$. Observe that given a transaction ξ of the lazy SI implementation, the trace of ξ , written θ_ξ , is of the form:

$$Fs \xrightarrow{\text{pol}_{\text{imm}}} Is \xrightarrow{\text{pol}_{\text{imm}}} Ts \xrightarrow{\text{pol}_{\text{imm}}} RUs \xrightarrow{\text{pol}_{\text{imm}}} PLs \xrightarrow{\text{pol}_{\text{imm}}} Ws \xrightarrow{\text{pol}_{\text{imm}}} Us$$

where:

- Fs denotes the sequence of events failing to obtain the necessary locks, i.e. those iterations that do not succeed in promoting the writer locks;
- Is denotes the sequence of events initialising the values of **LS,RS** and **WS** with \emptyset , and initialising $\mathbf{s}[\mathbf{x}]$ with \perp for each location \mathbf{x} ;
- Ts denotes the sequence of events corresponding to the execution of $\langle \mathbf{T} \rangle$ and is of the form $t_1 \xrightarrow{\text{pol}_{\text{imm}}} \dots \xrightarrow{\text{pol}_{\text{imm}}} t_k$, where for all $m \in \{1 \dots k\}$:

$$t_m = \begin{cases} rd(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \xrightarrow{\text{pol}_{\text{imm}}} lr_{\mathbf{x}_m} & \text{if } O_m = \mathbf{R}(\mathbf{x}_m, v_m) \\ wr(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \xrightarrow{\text{pol}_{\text{imm}}} wws_{\mathbf{x}_m} \xrightarrow{\text{pol}_{\text{imm}}} lw_{\mathbf{x}_m} & \text{if } O_m = \mathbf{W}(\mathbf{x}_m, v_m) \end{cases}$$

where O_m denotes the m th event in the trace of the original \mathbf{T} ; $lr_{\mathbf{x}_m} \triangleq \mathbf{R}(\mathbf{s}[\mathbf{x}_m], v_m)$; $lw_{\mathbf{x}_m} \triangleq \mathbf{W}(\mathbf{s}[\mathbf{x}_m], v_m)$;

$$rd(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \triangleq \begin{cases} \mathbf{R}(\mathbf{s}[\mathbf{x}_m], \perp) & \text{if } \mathbf{x}_m \notin RS_{m-1} \cup WS_{m-1} \\ \xrightarrow{\text{pol}_{\text{imm}}} fs_m \\ \xrightarrow{\text{pol}_{\text{imm}}} rl_{\mathbf{x}_m} \\ \xrightarrow{\text{pol}_{\text{imm}}} wrs_{\mathbf{x}_m} \\ \xrightarrow{\text{pol}_{\text{imm}}} rs_{\mathbf{x}_m} \\ \xrightarrow{\text{pol}_{\text{imm}}} ws_{\mathbf{x}_m} \\ \emptyset & \text{otherwise} \end{cases}$$

fs_m denotes the sequence of events attempting (but failing) to acquire the read lock on \mathbf{x}_m , $rl_{\mathbf{x}_m} \triangleq \mathbf{RL}(\mathbf{x}_m)$, $wrs_{\mathbf{x}_m} \triangleq \mathbf{W}(\mathbf{RS}, RS_m)$, $rs_{\mathbf{x}_m} \triangleq \mathbf{R}(\mathbf{x}_m, v_m)$, and $ws_{\mathbf{x}_m} \triangleq \mathbf{W}(\mathbf{s}[\mathbf{x}_m], v_m)$; and for all $m > 0$:

$$RS_{m+1} \triangleq \begin{cases} RS_m \cup \{\mathbf{x}_m\} & \text{if } O_m = \mathbf{R}(\mathbf{x}_m, -) \\ RS_m & \text{otherwise} \end{cases}$$

and

$$wr(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \triangleq \begin{cases} \mathbf{R}(\mathbf{s}[\mathbf{x}_m], \perp) & \text{if } \mathbf{x}_m \notin RS_{m-1} \cup WS_{m-1} \\ \xrightarrow{\text{pol}_{\text{imm}}} fs_m \\ \xrightarrow{\text{pol}_{\text{imm}}} rl_{\mathbf{x}_m} \emptyset & \text{otherwise} \end{cases}$$

$lw_{x_m} = W(\mathbf{s}[\mathbf{x}_m], v_m)$, $wws_{x_m} \triangleq W(\mathbf{WS}, WS_m)$; fs_m and rl_{x_m} are as defined above; and

$$WS_{m+1} \triangleq \begin{cases} WS_m \cup \{\mathbf{x}_m\} & \text{if } O_m = W(\mathbf{x}_m, -) \\ WS_m & \text{otherwise} \end{cases}$$

Let $RS_\xi = RS_m$ and $WS_\xi = WS_m$; let $RS_\xi \cup WS_\xi$ be enumerated as $\{\mathbf{x}_1 \cdots \mathbf{x}_i\}$ for some i .

- RU_s denotes the sequence of events releasing the reader locks (when the given location is in the read set only) and is of the form $ru_{x_1} \xrightarrow{\text{po|imm}} \cdots \xrightarrow{\text{po|imm}} ru_{x_i}$, where for all $n \in \{1 \cdots i\}$:

$$ru_{x_n} = \begin{cases} RU(\mathbf{x}1_n) & \text{if } \mathbf{x}_n \notin WS_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- PL_s denotes the sequence of events promoting the reader locks to writer ones (when the given location is in the write set), and is of the form $pl_{x_1} \xrightarrow{\text{po|imm}} \cdots \xrightarrow{\text{po|imm}} pl_{x_i}$, where for all $n \in \{1 \cdots i\}$:

$$pl_{x_n} = \begin{cases} PL(\mathbf{x}1_n) & \text{if } \mathbf{x}_n \in WS_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- Ws denotes the sequence of events committing the writes of (\mathbf{T}) and is of the form $c_{x_1} \xrightarrow{\text{po|imm}} \cdots \xrightarrow{\text{po|imm}} c_{x_i}$, where for all $n \in \{1 \cdots i\}$:

$$c_{x_n} = \begin{cases} R(\mathbf{s}[\mathbf{x}_n], v_n) \xrightarrow{\text{po|imm}} w_{x_n} = W(\mathbf{x}_n, v_n) & \text{if } \mathbf{x}_n \in WS_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- Us denotes the sequence of events releasing the locks on the write set, and is of the form $wu_{x_1} \xrightarrow{\text{po|imm}} \cdots \xrightarrow{\text{po|imm}} wu_{x_i}$, where for all $n \in \{1 \cdots i\}$:

$$wu_{x_n} = \begin{cases} WU(\mathbf{x}1_n) & \text{if } \mathbf{x}_n \in WS_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

Given a transaction trace θ_ξ , we write e.g. $\xi.U_s$ to refer to its constituent Us sub-trace and write $Us.E$ for the set of events related by po in Us . Similarly, we write $\xi.E$ for the set of events related by po in θ_ξ . Note that $G.E = \bigcup_{\xi \in \text{TX}} \xi.E$.

Note that for each transaction ξ and each location \mathbf{x} , the $\xi.rl_{\mathbf{x}}$, $\xi.rs_{\mathbf{x}}$, $\xi.ru_{\mathbf{x}}$, $\xi.pl_{\mathbf{x}}$, $\xi.wu_{\mathbf{x}}$ and $\xi.w_{\mathbf{x}}$ are uniquely identified when they exist.

For each location $\mathbf{x} \in WS_\xi$, let $fw_{\mathbf{x}}$ denote the maximal write (in po order within ξ) logging a write for \mathbf{x} in $\mathbf{s}[\mathbf{x}]$. That is, when $\theta_\xi = t_1 \xrightarrow{\text{po|imm}} \cdots \xrightarrow{\text{po|imm}} t_m$, let $fw_{\mathbf{x}} = \mathbf{wmax}(\mathbf{x}, [t_1 \cdots t_m])$, where

$$\mathbf{wmax}(\mathbf{x}, L++[t]) \triangleq \begin{cases} \mathbf{wmax}(\mathbf{x}, [t]) & \text{undefined} \\ lw_{\mathbf{x}} & \text{if } t = wr(\mathbf{x}, -, -, -) \xrightarrow{\text{po}} lw_{\mathbf{x}} \xrightarrow{\text{po}} wws_{\mathbf{x}} \\ \mathbf{wmax}(\mathbf{x}, L) & \text{otherwise} \end{cases}$$

Similarly, for each location $\mathbf{x} \in \text{WS}_\xi$, let $iw_{\mathbf{x}}$ denote the minimal write (in po order within ξ) logging a write for \mathbf{x} in $\mathbf{s}[\mathbf{x}]$. That is, when $\theta_\xi = t_1 \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} t_m$, let $iw_{\mathbf{x}} = \text{wmin}(\mathbf{x}, [t_1 \dots t_m])$, where

$$\text{wmin}(\mathbf{x}, [t]++L) \triangleq \begin{cases} \text{wmin}(\mathbf{x}, []) \text{ undefined} & \\ lw_{\mathbf{x}} & \text{if } t = wr(\mathbf{x}, -, -, -) \xrightarrow{\text{po}} lw_{\mathbf{x}} \xrightarrow{\text{po}} wws_{\mathbf{x}} \\ \text{wmin}(\mathbf{x}, L) & \text{otherwise} \end{cases}$$

D.1 Implementation Soundness

In order to establish the soundness of our implementation, it suffices to show that given an RA-consistent execution graph of the implementation G , we can construct a corresponding SI-consistent execution graph G' with the same outcome.

Given a transaction $\xi \in \text{Tx}$ with $\text{RS}_\xi \cup \text{WS}_\xi = \{\mathbf{x}_1 \dots \mathbf{x}_i\}$ and trace θ_ξ as above with $Ts = t_1 \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} t_k$, we construct the corresponding SI execution trace θ'_ξ as follows:

$$\theta'_\xi \triangleq t'_1 \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} t'_k$$

where for all $m \in \{1 \dots k\}$:

$$\begin{aligned} t'_m &= R(\mathbf{x}_m, v_m) \text{ when } t_m = rd(\mathbf{x}_m, v_m, S_m) \xrightarrow{\text{po|imm}} lr_{\mathbf{x}_m} \\ t'_m &= W(\mathbf{x}_m, v_m) \text{ when } t_m = wr(\mathbf{x}_m, v_m, S_m) \xrightarrow{\text{po|imm}} lw_{\mathbf{x}_m} \xrightarrow{\text{po|imm}} wws_{\mathbf{x}_m} \end{aligned}$$

such that in the first case the identifier of t'_m is that of $lr_{\mathbf{x}_m}$; and in the second case the identifier of t'_m is that of $lw_{\mathbf{x}_m}$. We then define:

$$\text{RF}_\xi \triangleq \left\{ (w, t'_j) \left| \begin{array}{l} t'_j \in \theta'_\xi.E \wedge \exists \mathbf{x}, v. t'_j = R(\mathbf{x}, v) \wedge w = W(\mathbf{x}, v) \\ \wedge (w \in \xi.E \Rightarrow w \xrightarrow{\text{po}} t'_j \wedge \\ (\forall e \in \xi.E. w \xrightarrow{\text{po}} e \xrightarrow{\text{po}} t'_j \Rightarrow (\text{loc}(e) \neq \mathbf{x} \vee e \notin \mathcal{W}))) \\ \wedge (w \notin \xi.E \Rightarrow (\forall e \in \xi.E. (e \xrightarrow{\text{po}} t'_j \Rightarrow (\text{loc}(e) \neq \mathbf{x} \vee e \notin \mathcal{W}))) \\ \wedge \exists \xi'. (\xi'.w_{\mathbf{x}}, \xi'.rs_{\mathbf{x}}) \in G.\text{rf}) \wedge w = \xi'.fw_{\mathbf{x}} \end{array} \right. \right\}$$

Similarly, we define:

$$\text{MO} \triangleq \left\{ (w_1, w_2), (w_3, w_4) \left| \begin{array}{l} \exists \xi. w_1, w_2 \in \xi.E \cap \mathcal{W} \wedge \text{loc}(w_1) = \text{loc}(w_2) \wedge (w_1, w_2) \in G.\text{po} \\ \wedge \exists \xi_1, \xi_2, \mathbf{x}. w_3 = \xi_1.fw_{\mathbf{x}} \wedge w_4 = \xi_1.iw_{\mathbf{x}} \wedge (\xi_1.w_{\mathbf{x}}, \xi_2.w_{\mathbf{x}}) \in G.\text{mo} \end{array} \right. \right\}^+$$

We are now in a position to demonstrate the soundness of our implementation. Given an RA-consistent execution graph G of the implementation, we construct an SI execution graph G' as follows and demonstrate that $\text{si-consistent}(G')$ holds.

– $G'.E = \bigcup_{\xi \in \text{Tx}} \theta'_\xi.E$, with the $\text{tx}(\cdot)$ function defined as:

$$\text{tx}(a) \triangleq \xi \quad \text{where} \quad a \in \theta'_\xi$$

- $G'.\text{po} = G.\text{po}|_{G'.E}$
- $G'.\text{rf} = \bigcup_{\xi \in \text{Tx}} \text{RF}_\xi$
- $G'.\text{mo} = \text{MO}$
- $G'.\text{lo} = \emptyset$

Observe that the events of each θ'_ξ trace coincides with those of the equivalence class \mathcal{T}_ξ of G' . That is, $\theta'_\xi.E = \mathcal{T}_\xi$.

Lemma 8. *Given an RA-consistent execution graph G of the implementation and its corresponding SI execution graph G' constructed as above, for all $a, b, \xi_a, \xi_b, \mathbf{x}$:*

$$\xi_a \neq \xi_b \wedge a \in \xi_a.E \wedge b \in \xi_b.E \wedge \text{loc}(a) = \text{loc}(b) = \mathbf{x} \Rightarrow$$

$$((a, b) \in G'.\text{rf} \Rightarrow \xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x) \quad (19)$$

$$\wedge ((a, b) \in G'.\text{mo} \Rightarrow \xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x) \quad (20)$$

$$\wedge ((a, b) \in G'.\text{rb} \Rightarrow (\mathbf{x} \in \text{WS}_{\xi_a} \wedge \xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x) \vee (\mathbf{x} \notin \text{WS}_{\xi_a} \wedge \xi_a.ru_x \xrightarrow{G.\text{hb}} \xi_b.pl_x)) \quad (21)$$

Proof. Pick an arbitrary RA-consistent execution graph G of the implementation and its corresponding SI execution graph G' constructed as above. Pick an arbitrary $a, b, \xi_a, \xi_b, \mathbf{x}$ such that $\xi_a \neq \xi_b$, $a \in \xi_a.E$, $b \in \xi_b.E$, and $\text{loc}(a) = \text{loc}(b) = \mathbf{x}$.

RTS. (19)

Assume $(a, b) \in G'.\text{rf}$. From the definition of $G'.\text{rf}$ we then know $(\xi_a.w_x, \xi_b.rs_x) \in G.\text{rf}$. On the other hand, from Lemma 2 we know that either i) $\mathbf{x} \in \text{WS}_{\xi_b}$ and $\xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x$; or ii) $\mathbf{x} \notin \text{WS}_{\xi_b}$ and $\xi_b.ru_x \xrightarrow{G.\text{hb}} \xi_a.pl_x$; or iii) $\xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x$. In case (i) we then have $\xi_a.w_x \xrightarrow{G.\text{rf}} \xi_b.rs_x \xrightarrow{G.\text{po}} \xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x \xrightarrow{G.\text{po}} \xi_a.w_x$. That is, we have $\xi_a.w_x \xrightarrow{G.\text{hb}^{\text{loc}}} \xi_a.w_x$, contradicting the assumption that G is RA-consistent. Similarly in case (ii) we have $\xi_a.w_x \xrightarrow{G.\text{rf}} \xi_b.rs_x \xrightarrow{G.\text{po}} \xi_b.ru_x \xrightarrow{G.\text{hb}} \xi_a.pl_x \xrightarrow{G.\text{po}} \xi_a.w_x$. That is, we have $\xi_a.w_x \xrightarrow{G.\text{hb}^{\text{loc}}} \xi_a.w_x$, contradicting the assumption that G is RA-consistent. In case (iii) the desired result holds trivially.

RTS. (20)

Assume $(a, b) \in G'.\text{mo}$. From the definition of $G'.\text{mo}$ we then know $(\xi_a.w_x, \xi_b.w_x) \in G.\text{mo}$. On the other hand, from Lemma 2 we know that either i) $\xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x$; or ii) $\xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x$. In case (i) we then have $\xi_a.w_x \xrightarrow{G.\text{mo}} \xi_b.w_x \xrightarrow{G.\text{po}} \xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x \xrightarrow{G.\text{po}} \xi_a.w_x$. That is, we have $\xi_a.w_x \xrightarrow{G.\text{mo}} \xi_b.w_x \xrightarrow{G.\text{hb}^{\text{loc}}} \xi_a.w_x$, contradicting the assumption that G is RA-consistent. In case (ii) the desired result holds trivially.

RTS. (21)

Assume $(a, b) \in G'.\text{rb}$. From the definition of $G'.\text{rb}$ we then know that either $(\xi_a.w_x, \xi_b.w_x) \in G.\text{mo}$ or $(\xi_a.rs_x, \xi_b.w_x) \in G.\text{rb}$. In the former case the desired result follows immediately from the proof of part (20). In the latter

case, from [Lemma 2](#) we know that either i) $\xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x$; or ii) $\mathbf{x} \notin \text{WS}_{\xi_a}$ and $\xi_a.ru_x \xrightarrow{G.\text{hb}} \xi_a.pl_x$; or iii) $\mathbf{x} \in \text{WS}_{\xi_a}$ and $\xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x$. In case (i) we then have $\xi_b.w_x \xrightarrow{G.\text{po}} \xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x \xrightarrow{G.\text{po}} \xi_a.rs_x \xrightarrow{G.\text{rb}} \xi_b.w_x$. That is, we have $\xi_b.w_x \xrightarrow{G.\text{hb}^{\text{loc}}} \xi_a.rs_x \xrightarrow{G.\text{rb}} \xi_b.w_x$, contradicting the assumption that G is RA-consistent. In cases (ii-iii) the desired result holds trivially.

Lemma 9. For all RA-consistent execution graphs G of the implementation and their counterpart SI execution graphs G' constructed as above:

$$\begin{aligned} & \forall \xi_a, \xi_b. \forall a \in G'.\mathcal{T}_{\xi_a}, b \in G'.\mathcal{T}_{\xi_b}. \\ & (a, b) \in G'.(\text{po}_{\top} \cup \text{rf}_{\top} \cup \text{mo}_{\top})^+ \Rightarrow \\ & \quad \exists d \in G'.\mathcal{T}_{\xi_b}. \forall c \in G'.\mathcal{T}_{\xi_a}. (c, d) \in G.\text{hb} \\ & \quad \wedge (\exists \mathbf{y}. (\xi_a.wu_{\mathbf{y}}, d) \in G.\text{hb} \vee \forall c \in \xi_a.E. (c, d) \in G.\text{hb}) \end{aligned}$$

Proof. Let $S_0 = G'.(\text{po}_{\top} \cup \text{rf}_{\top} \cup \text{mo}_{\top})$, and $S_{n+1} = S_0; S_n$, for all $n \geq 0$. It is straightforward to demonstrate that $G'.(\text{po}_{\top} \cup \text{rf}_{\top} \cup \text{mo}_{\top})^+ = \bigcup_{i \in \mathbb{N}} S_i$. We thus demonstrate instead that:

$$\begin{aligned} & \forall i \in \mathbb{N}. \forall \xi_a, \xi_b. \forall a \in G'.\mathcal{T}_{\xi_a}, b \in G'.\mathcal{T}_{\xi_b}. \\ & (a, b) \in S_i \Rightarrow \exists d \in G'.\mathcal{T}_{\xi_b}. \forall c \in G'.\mathcal{T}_{\xi_a}. (c, d) \in G.\text{hb} \\ & \quad \wedge (\exists \mathbf{y}. (\xi_a.wu_{\mathbf{y}}, d) \in G.\text{hb} \vee \forall c \in \xi_a.E. (c, d) \in G.\text{hb}) \end{aligned}$$

We proceed by induction on i .

Base case $i = 0$

Pick arbitrary ξ_a, ξ_b , $a \in G'.\mathcal{T}_{\xi_a}, b \in G'.\mathcal{T}_{\xi_b}$ such that $(a, b) \in S_0$. There are now three cases to consider: 1) $(a, b) \in G'.\text{po}_{\top}$; or 2) $(a, b) \in G'.\text{rf}_{\top}$; or 3) $(a, b) \in G'.\text{mo}_{\top}$.

In case (1), pick an arbitrary $c \in G'.\mathcal{T}_{\xi_a}$. From the definition of $G'.\text{po}_{\top}$ we have $(c, b) \in G.\text{po} \subseteq G.\text{hb}$, as required. Pick an arbitrary $c \in \xi_a.E$. From the definition of $G'.\text{po}_{\top}$ we have $(c, b) \in G.\text{po} \subseteq G.\text{hb}$, as required.

In case (2), we then know there exists $w \in G'.\mathcal{T}_{\xi_a}$ and $r \in G'.\mathcal{T}_{\xi_b}$ such that $(w, r) \in G'.\text{rf}$. Let $\text{loc}(w) = \text{loc}(r) = \mathbf{x}$. From [Lemma 8](#) we then have $\xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x$. Pick an arbitrary $c \in G'.\mathcal{T}_{\xi_a}$. As such we have $c \xrightarrow{G.\text{po}} \xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x \xrightarrow{G.\text{po}} \xi_b.rs_x$. That is, we have $(c, \xi_b.rs_x), (\xi_a.wu_x, \xi_b.rs_x) \in G.\text{hb}$, as required.

In case (3), we then know there exists $w \in G'.\mathcal{T}_{\xi_a}$ and $w' \in G'.\mathcal{T}_{\xi_b}$ such that $(w, w') \in G'.\text{mo}$. Let $\text{loc}(w) = \text{loc}(w') = \mathbf{x}$. From [Lemma 8](#) we then have $\xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x$. Pick an arbitrary $c \in G'.\mathcal{T}_{\xi_a}$. As we also have $\xi_b.rl_x \xrightarrow{G.\text{po}} w'$, and $c \xrightarrow{G.\text{po}} \xi_a.wu_x$, we then have $(c, w'), (\xi_a.wu_x, w') \in G.\text{hb}$, as required.

Inductive case $i = n + 1$

$$\begin{aligned} & \forall j \in \mathbb{N}. \forall \xi_a, \xi_b. \forall a \in G'.\mathcal{T}_{\xi_a}, b \in G'.\mathcal{T}_{\xi_b}. \\ & (a, b) \in S_j \wedge j \leq n \Rightarrow \\ & \quad \exists d \in G'.\mathcal{T}_{\xi_b}. \forall c \in G'.\mathcal{T}_{\xi_a}. (c, d) \in G.\text{hb} \\ & \quad \wedge (\exists \mathbf{y}. (\xi_a.wu_{\mathbf{y}}, d) \in G.\text{hb} \vee \forall c \in \xi_a.E. (c, d) \in G.\text{hb}) \end{aligned} \tag{I.H.}$$

Pick arbitrary ξ_a, ξ_b , $a \in G'.\mathcal{T}_{\xi_a}, b \in G'.\mathcal{T}_{\xi_b}$ such that $(a, b) \in S_i$. From the definition of S_i we then know there exist e, ξ_e such that $e \in G'.\mathcal{T}_{\xi_e}$, $(a, e) \in S_0$ and $(e, b) \in S_n$. Since $(e, b) \in S_n$, from (I.H.) we know there exists $d \in G'.\mathcal{T}_{\xi_b}$ such that $\forall f' \in G'.\mathcal{T}_{\xi_e}. (f', d) \in G.\mathbf{hb}$. On the other hand, from the proof of the base case we know there exists $f \in G'.\mathcal{T}_{\xi_e}$ such that $\forall c' \in G'.\mathcal{T}_{\xi_a}. (c', f) \in G.\mathbf{hb}$; and that $\exists y. (\xi_e.wu_y, f) \in G.\mathbf{hb} \vee \forall c' \in \xi_a.E. (c', f) \in G.\mathbf{hb}$. Pick an arbitrary $c \in G'.\mathcal{T}_{\xi_a}$. We thus know that $(c, f) \in G.\mathbf{hb}$. As $f \in G'.\mathcal{T}_{\xi_e}$, we thus have $(f, d) \in G.\mathbf{hb}$. As $G.\mathbf{hb}$ is transitively closed and we have $(c, f), (f, d) \in G.\mathbf{hb}$ and $\exists y. (\xi_e.wu_y, f) \in G.\mathbf{hb} \vee \forall c' \in \xi_a.E. (c', f) \in G.\mathbf{hb}$, we then have $(c, d) \in G.\mathbf{hb}$, and $\exists y. (\xi_e.wu_y, d) \in G.\mathbf{hb} \vee \forall c' \in \xi_a.E. (c', d) \in G.\mathbf{hb}$, as required.

Theorem 9 (Soundness). *For all execution graphs G of the implementation and their counterpart SI execution graphs G' constructed as above,*

$$\text{RA-consistent}(G) \Rightarrow \text{si-consistent}(G')$$

Proof. Pick an arbitrary execution graph G of the implementation such that $\text{RA-consistent}(G)$ holds, and its associated SI execution graph G' constructed as described above.

RTS. irreflexive(si-hb)

We proceed by contradiction. Let us assume $\neg \text{irreflexive}(G'.(\text{po}_T \cup \text{rf}_T \cup \text{mo}_T \cup \text{si-rb}))$. Let $S = G'.(\text{po}_T \cup \text{rf}_T \cup \text{mo}_T)^+$. There are two cases to consider: either there is a cycle without a **si-rb** edge; or there is a cycle with one or more **si-rb** edges. That is, either 1) there exists a such that $(a, a) \in S$; or 2) there exist $a_1, b_1, \dots, a_n, b_n$ such that $a_1 \xrightarrow{G'.\text{si-rb}} b_1 \xrightarrow{S} a_2 \xrightarrow{G'.\text{si-rb}} b_2 \xrightarrow{S} \dots \xrightarrow{S} a_n \xrightarrow{G'.\text{si-rb}} b_n \xrightarrow{S} a_1$.

In case (1) we then know there exists ξ such that $a \in G'.\mathcal{T}_{\xi}$. As such, from **Lemma 9** we know that there exists $d \in G'.\mathcal{T}_{\xi}$ such that for all $c \in G'.\mathcal{T}_{\xi}$, $(c, d) \in G.\mathbf{hb}$. As such, we have $(d, d) \in G.\mathbf{hb}$, contradicting the assumption that G is RA-consistent.

In case (2), for an arbitrary $i \in \{1 \dots n\}$, let $j = i+1$ when $i \neq n$; and $j = 1$ when $i = n$. As $a_i \xrightarrow{G'.\text{si-rb}} b_i$, we know there exists ξ_{a_i}, ξ_{b_i} such that $a_i \in \mathcal{T}_{\xi_{a_i}}$, $b_i \in \mathcal{T}_{\xi_{b_i}}$, and that there exist $r_i \in \mathcal{T}_{\xi_{a_i}} \cap \mathcal{R}_E$ and $w_i \in \mathcal{T}_{\xi_{b_i}} \cap \mathcal{W}$ such that $(r_i, w_i) \in G'.\mathbf{rb}$. Let $\text{loc}(r_i) = \text{loc}(w_i) = \mathbf{x}_i$. From **Lemma 8** we then know that either i) $\xi_{a_i}.ru_{\mathbf{x}_i} \xrightarrow{G.\mathbf{hb}} \xi_{b_i}.pl_{\mathbf{x}_i}$; or ii) $\xi_{a_i}.wu_{\mathbf{x}_i} \xrightarrow{G.\mathbf{hb}} \xi_{b_i}.rl_{\mathbf{x}_i}$. Note that for all y such that $\xi_{b_i}.wu_y$ exists, we know $\xi_{b_i}.rl_{\mathbf{x}_i} \xrightarrow{G.\text{po}} \xi_{b_i}.wu_y$ and $\xi_{b_i}.pl_{\mathbf{x}_i} \xrightarrow{G.\text{po}} \xi_{b_i}.wu_y$. As such, as $(b_i, a_j) \in S$, from **Lemma 9** and since $G.\mathbf{hb}$ is transitively closed, we know there exists $d_j \in \mathcal{T}_{\xi_{a_j}}$ such that either $\xi_{a_i}.ru_{\mathbf{x}_i} \xrightarrow{G.\mathbf{hb}} d_j$, or $\xi_{a_i}.wu_{\mathbf{x}_i} \xrightarrow{G.\mathbf{hb}} d_j$. That is, $\xi_{a_i}.u_{\mathbf{x}_i} \xrightarrow{G.\mathbf{hb}} d_j$, where either $u_{\mathbf{x}_i} = ru_{\mathbf{x}_i}$ or $u_{\mathbf{x}_i} = wu_{\mathbf{x}_i}$. On the other hand, observe that for all $d_i \in \mathcal{T}_{\xi_{a_i}}$ we have $d_i \xrightarrow{G.\text{po}} \xi_{a_i}.u_{\mathbf{x}_i}$ i.e. $d_i \xrightarrow{G.\mathbf{hb}} \xi_{a_i}.u_{\mathbf{x}_i}$. As such, we have $d_j \xrightarrow{G.\mathbf{hb}} \xi_{a_j}.u_{\mathbf{x}_j}$. As we also have $\xi_{a_i}.u_{\mathbf{x}_i} \xrightarrow{G.\mathbf{hb}} d_j$ and $G.\mathbf{hb}$ is transitively closed, we have $\xi_{a_i}.u_{\mathbf{x}_i} \xrightarrow{G.\mathbf{hb}} \xi_{a_j}.u_{\mathbf{x}_j}$. We then have $\xi_{a_1}.u_{\mathbf{x}_1} \xrightarrow{G.\mathbf{hb}} \xi_{a_2}.u_{\mathbf{x}_2} \xrightarrow{G.\mathbf{hb}} \dots \xrightarrow{G.\mathbf{hb}} \xi_{a_n}.u_{\mathbf{x}_n} \xrightarrow{G.\mathbf{hb}} \xi_{a_1}.u_{\mathbf{x}_1}$. That is, $\xi_{a_1}.u_{\mathbf{x}_1} \xrightarrow{G.\mathbf{hb}} \xi_{a_1}.u_{\mathbf{x}_1}$, contradicting the assumption

that G is RA-consistent.

RTS. $\text{rf}_I \cup \text{mo}_I \cup \text{rb}_I \subseteq \text{po}$

Follows immediately from the construction of G' .

D.2 Implementation Completeness

In order to establish the completeness of our implementation, it suffices to show that given an SI-consistent execution graph $G' = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$, we can construct a corresponding RA-consistent execution graph G of the implementation.

Note that the execution trace for each transaction $\mathcal{T}_\xi \in G'.\mathcal{T}/\text{st}$ is of the form $\theta'_\xi = t'_1 \xrightarrow{\text{po}|_{\text{imm}}} \dots \xrightarrow{\text{po}|_{\text{imm}}} t'_k$ for some k , where each t'_i is a read or write event. As such, we have $G'.E = G'.\mathcal{T} = \bigcup_{\mathcal{T}_\xi \in G'.\mathcal{T}/\text{st}} \mathcal{T}_\xi = \theta'_\xi.E$. We thus construct the implementation trace θ_ξ as follows:

$$I_s \xrightarrow{\text{po}|_{\text{imm}}} T_s \xrightarrow{\text{po}|_{\text{imm}}} RUs \xrightarrow{\text{po}|_{\text{imm}}} PLs \xrightarrow{\text{po}|_{\text{imm}}} Ws \xrightarrow{\text{po}|_{\text{imm}}} Us$$

where:

- I_s denotes the sequence of events initialising the values of RS and WS with \emptyset , and initialising $\mathbf{s}[\mathbf{x}]$ with \perp for each location \mathbf{x} ;
- T_s is of the form $t_1 \xrightarrow{\text{po}|_{\text{imm}}} \dots \xrightarrow{\text{po}|_{\text{imm}}} t_k$, where for all $m \in \{1 \dots k\}$:

$$t_m = \begin{cases} rd(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \xrightarrow{\text{po}|_{\text{imm}}} lr_{\mathbf{x}_m} & \text{if } t'_m = \text{R}(\mathbf{x}_m, v_m) \\ wr(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \xrightarrow{\text{po}|_{\text{imm}}} wws_{\mathbf{x}_m} \xrightarrow{\text{po}|_{\text{imm}}} lw_{\mathbf{x}_m} & \text{if } t'_m = \text{W}(\mathbf{x}_m, v_m) \end{cases}$$

where $lr_{\mathbf{x}_m} \triangleq \text{R}(\mathbf{s}[\mathbf{x}_m], v_m)$; $lw_{\mathbf{x}_m} \triangleq \text{W}(\mathbf{s}[\mathbf{x}_m], v_m)$; the identifiers of $lr_{\mathbf{x}_m}$ and $lw_{\mathbf{x}_m}$ are those of t'_m , whilst the identifiers of other events in t_m are picked fresh;

$$rd(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \triangleq \begin{cases} \text{R}(\mathbf{s}[\mathbf{x}_m], \perp) & \text{if } \mathbf{x}_m \notin RS_{m-1} \cup WS_{m-1} \\ \xrightarrow{\text{po}|_{\text{imm}}} wrs_{\mathbf{x}_m} \\ \xrightarrow{\text{po}|_{\text{imm}}} fs_m \\ \xrightarrow{\text{po}|_{\text{imm}}} rl_{\mathbf{x}_m} \\ \xrightarrow{\text{po}|_{\text{imm}}} rs_{\mathbf{x}_m} \\ \xrightarrow{\text{po}|_{\text{imm}}} ws_{\mathbf{x}_m} \\ \emptyset & \text{otherwise} \end{cases}$$

$wrs_{\mathbf{x}_m} \triangleq \text{W}(\text{RS}, RS_m)$, fs_m denotes the sequence of events attempting (but failing) to acquire the read lock on \mathbf{x}_m , $rl_{\mathbf{x}_m} \triangleq \text{RL}(\mathbf{x}_m)$, $rs_{\mathbf{x}_m} \triangleq \text{R}(\mathbf{x}_m, v_m)$, and $ws_{\mathbf{x}_m} \triangleq \text{W}(\mathbf{s}[\mathbf{x}_m], v_m)$; $RS_0 = \emptyset$ and for all $m > 0$:

$$RS_{m+1} \triangleq \begin{cases} RS_m \cup \{\mathbf{x}_m\} & \text{if } t'_m = \text{R}(\mathbf{x}_m, -) \\ RS_m & \text{otherwise} \end{cases}$$

and

$$wr(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \triangleq \begin{cases} \mathbf{R}(\mathbf{s}[\mathbf{x}_m], \perp) & \text{if } \mathbf{x}_m \notin RS_{m-1} \cup WS_{m-1} \\ \xrightarrow{\text{po|imm}} fs_m & \\ \xrightarrow{\text{po|imm}} rl_{\mathbf{x}_m} \emptyset & \text{otherwise} \end{cases}$$

$lw_{\mathbf{x}_m} = \mathbf{W}(\mathbf{s}[\mathbf{x}_m], v_m)$, $wws_{\mathbf{x}_m} \triangleq \mathbf{W}(WS, WS_m)$; fs_m and $rl_{\mathbf{x}_m}$ are as defined above; and

$$WS_{m+1} \triangleq \begin{cases} WS_m \cup \{\mathbf{x}_m\} & \text{if } t'_m = \mathbf{W}(\mathbf{x}_m, -) \\ WS_m & \text{otherwise} \end{cases}$$

Let $RS_\xi = RS_m$ and $WS_\xi = WS_m$; let $RS_\xi \cup WS_\xi$ be enumerated as $\{\mathbf{x}_1 \cdots \mathbf{x}_i\}$ for some i .

- RUs denotes the sequence of events releasing the reader locks (when the given location is in the read set only) and is of the form $ru_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \cdots \xrightarrow{\text{po|imm}} ru_{\mathbf{x}_i}$, where for all $n \in \{1 \cdots i\}$:

$$ru_{\mathbf{x}_n} = \begin{cases} \mathbf{RU}(\mathbf{x}1_n) & \text{if } \mathbf{x}_n \notin WS_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

with the identifier of each $ru_{\mathbf{x}_n}$ picked fresh;

- PLs denotes the sequence of events promoting the reader locks to writer ones (when the given location is in the write set), and is of the form $pl_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \cdots \xrightarrow{\text{po|imm}} pl_{\mathbf{x}_i}$, where for all $n \in \{1 \cdots i\}$:

$$pl_{\mathbf{x}_n} = \begin{cases} \mathbf{PL}(\mathbf{x}1_n) & \text{if } \mathbf{x}_n \in WS_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

with the identifier of each $ru_{\mathbf{x}_n}$ picked fresh;

- Ws denotes the sequence of events committing the writes of (T) and is of the form $c_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \cdots \xrightarrow{\text{po|imm}} c_{\mathbf{x}_i}$, where for all $n \in \{1 \cdots i\}$:

$$c_{\mathbf{x}_n} = \begin{cases} \mathbf{R}(\mathbf{s}[\mathbf{x}_n], v_n) \xrightarrow{\text{po|imm}} w_{\mathbf{x}_n} = \mathbf{W}(\mathbf{x}_n, v_n) & \text{if } \mathbf{x}_n \in WS_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

with the identifiers of events in each $c_{\mathbf{x}_n}$ picked fresh;

- Us denotes the sequence of events releasing the locks on the write set, and is of the form $wu_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \cdots \xrightarrow{\text{po|imm}} wu_{\mathbf{x}_i}$, where for all $n \in \{1 \cdots i\}$:

$$wu_{\mathbf{x}_n} = \begin{cases} \mathbf{WU}(\mathbf{x}1_n) & \text{if } \mathbf{x}_n \in WS_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

with the identifier of each $wu_{\mathbf{x}_n}$ picked fresh.

We use the ξ . prefix to project the various events of the implementation trace θ_ξ (e.g. $\xi.rl_{x_j}$). Note that for each transaction ξ and each location \mathbf{x} , the $\xi.rl_{\mathbf{x}}$, $\xi.rs_{\mathbf{x}}$, $\xi.ru_{\mathbf{x}}$, $\xi.pl_{\mathbf{x}}$, $\xi.wu_{\mathbf{x}}$ and $\xi.w_{\mathbf{x}}$ are uniquely identified when they exist.

For each location $\mathbf{x} \in \text{WS}_\xi$, let $fw_{\mathbf{x}}$ denote the maximal write (in po order within ξ) logging a write for \mathbf{x} in $\mathbf{s}[\mathbf{x}]$. That is, when $\theta_\xi = t_1 \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} t_m$, let $fw_{\mathbf{x}} = \text{wmax}(\mathbf{x}, [t_1 \dots t_m])$, where

$$\text{wmax}(\mathbf{x}, L++[t]) \triangleq \begin{cases} \text{wmax}(\mathbf{x}, []) \text{ undefined} & \\ lw_{\mathbf{x}} & \text{if } t=wr(\mathbf{x}, -, -, -) \xrightarrow{\text{po}} lw_{\mathbf{x}} \xrightarrow{\text{po}} wws_{\mathbf{x}} \\ \text{wmax}(\mathbf{x}, L) & \text{otherwise} \end{cases}$$

Similarly, for each location $\mathbf{x} \in \text{WS}_\xi$, let $iw_{\mathbf{x}}$ denote the minimal write (in po order within ξ) logging a write for \mathbf{x} in $\mathbf{s}[\mathbf{x}]$. That is, when $\theta_\xi = t_1 \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} t_m$, let $iw_{\mathbf{x}} = \text{wmin}(\mathbf{x}, [t_1 \dots t_m])$, where

$$\text{wmin}(\mathbf{x}, [t]++L) \triangleq \begin{cases} \text{wmin}(\mathbf{x}, []) \text{ undefined} & \\ lw_{\mathbf{x}} & \text{if } t=wr(\mathbf{x}, -, -, -) \xrightarrow{\text{po}} lw_{\mathbf{x}} \xrightarrow{\text{po}} wws_{\mathbf{x}} \\ \text{wmin}(\mathbf{x}, L) & \text{otherwise} \end{cases}$$

Analogously, for each location $\mathbf{x} \in \text{RS}_\xi$, let $ir_{\mathbf{x}}$ denote the minimal read (in po order within ξ) reading the value of \mathbf{x} . That is, when $\theta_\xi = t_1 \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} t_m$, let $ir_{\mathbf{x}} = \text{rmin}(\mathbf{x}, [t_1 \dots t_m])$, where

$$\text{rmin}(\mathbf{x}, [t]++L) \triangleq \begin{cases} \text{rmin}(\mathbf{x}, []) \text{ undefined} & \\ lr_{\mathbf{x}} & \text{if } t=rd(\mathbf{x}, -, -, -) \xrightarrow{\text{po}} lr_{\mathbf{x}} \\ \text{rmin}(\mathbf{x}, L) & \text{otherwise} \end{cases}$$

Given the G' classes \mathcal{T}/st , let us construct the strict total order $\text{TCO} : \mathcal{T}/\text{st} \times \mathcal{T}/\text{st}$ as described in [Appendix C.2](#). For each location \mathbf{x} and $i \in \mathbb{N}$, let us similarly define $S_i^{\mathbf{x}}$, $P_i^{\mathbf{x}}$, $P_i^{\mathbf{x}}.\text{writer}$, $\text{LO}_{\mathbf{x}}$, $\text{RLO}_{\mathbf{x}}$, lo_1 and lo_2 .

Remark 5. Recall that both lo_1 and lo_2 satisfy the conditions stated in [Def. 4](#); lo_2 additionally satisfies the ‘read-read-synchronisation’ property in [\(RSYNC\)](#). As before, we demonstrate that given an SI-consistent execution graph G' , it is always possible to construct an RA-consistent execution graph G of the implementation with its lock order defined as lo_2 . Recall that as $\text{lo}_1 \subseteq \text{lo}_2$, it is straightforward to show that replacing lo_2 in such a G with lo_1 , preserves the RA-consistency of G , as defined in [Def. 4](#). In other words, as $\text{lo}_1 \subseteq \text{lo}_2$, we have:

$$\text{acyclic}(\text{hb}_{loc} \cup \text{mo} \cup \text{rb}) \text{ with } \text{hb} \triangleq (\text{po} \cup \text{rf} \cup \text{lo}_2)^+ \Rightarrow \text{acyclic}(\text{hb}_{loc} \cup \text{mo} \cup \text{rb}) \text{ with } \text{hb} \triangleq (\text{po} \cup \text{rf} \cup \text{lo}_1)^+$$

As such, by establishing the completeness of our implementation with respect to lo_2 , we also establish its completeness with respect to lo_1 . In other words, we demonstrate the completeness of our implementation with respect to both lock implementations presented earlier in [Appendix A](#).

We now demonstrate the completeness of our implementation. Given an SI-consistent graph, we construct an implementation graph G as follows and demonstrate that it is RA-consistent.

- $G.E = \bigcup_{\mathcal{T}_\xi \in G'.\mathcal{T}/\text{st}} \theta_\xi.E$, with the $\text{tx}(e) = 0$, for all $e \in G.E$.
- $G.\text{po}$ is defined as $G'.\text{po}$ extended by the po for the additional events of G , given by each θ_ξ trace defined above. Note that $G.\text{po}$ does not introduce additional orderings between events of $G'.E$. That is, $\forall a, b \in G'.E. (a, b) \in G'.\text{po} \Leftrightarrow (a, b) \in G.\text{po}$.
- $G.\text{rf} = \bigcup_{\xi \in \text{TX}} \text{RF}_\xi$ with $\text{RF}_\xi \triangleq \{(w, \xi.rs_x) \mid x \in \text{Loc} \wedge (w, \xi.ir_x) \in G'.\text{rf}\}$.
- $G.\text{mo} = \{(\xi_1.w_x, \xi_2.w_x) \mid \xi_1, \xi_2, \in \text{TX} \wedge x \in \text{Loc} \wedge (\xi_1.fw_x, \xi_2.iw_x) \in G'.\text{mo}\}^+$.
- $G.\text{lo} = \bigcup_{x \in \text{Loc}} \text{RLO}_x$, with RLO_x as defined above.

Notation Given an implementation graph G as constructed above (with $G.\mathcal{T} = \emptyset$), and a relation $r \subseteq G.E \times G.E$, we override the r_E notation and write r_E for:

$$\{(a, b) \in r \mid \exists \xi_a, \xi_b. a \in \theta_{\xi_a}.E \wedge b \in \theta_{\xi_b}.E \wedge \xi_a \neq \xi_b\}$$

Analogously, we write r_I for $\{(a, b) \in r \mid \exists \xi. a, b \in \theta_\xi.E\}$.

Lemma 10. *For all SI-consistent execution graphs $G' = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$, and for all $x \in \text{Loc}$, $\xi_a, \xi_b \in \text{TXID}$, and $i, j \in \mathbb{N}$:*

$$\xi_a \in S_i^x \wedge \xi_b \in S_j^x \wedge i < j \Rightarrow (\xi_a, \xi_b) \in \text{TCO}$$

Proof. Follows immediately from [Corollary 1](#).

Given an implementation graph $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$, let us define hb' and hb'_i as described in [Appendix C.2](#) for all $i \in \mathbb{N}$.

Lemma 11. *For all implementation graphs $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ constructed as above,*

$$\begin{aligned} \forall a, b, \xi_a, \xi_b. a \in \theta_{\xi_a}.E \wedge b \in \theta_{\xi_b}.E \wedge a \xrightarrow{\text{hb}'} b \\ \Rightarrow (a, b) \in G.\text{po}_I \vee (\xi_a, \xi_b) \in \text{TCO} \end{aligned}$$

Proof. Pick an arbitrary implementation graph $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ constructed as above. Since hb' is a transitive closure, it is straightforward to demonstrate that $\text{hb}' = \bigcup_{i \in \mathbb{N}} \text{hb}'_i$. It thus suffices to show:

$$\begin{aligned} \forall i \in \mathbb{N}. \forall a, b, \xi_a, \xi_b. a \in \theta_{\xi_a}.E \wedge b \in \theta_{\xi_b}.E \wedge a \xrightarrow{\text{hb}'_i} b \\ \Rightarrow (a, b) \in G.\text{po}_I \vee (\xi_a, \xi_b) \in \text{TCO} \end{aligned}$$

We proceed by induction over i .

Base case $i = 0$

Pick arbitrary a, b, ξ_a, ξ_b such that $a \in \theta_{\xi_a}.E$ and $b \in \theta_{\xi_b}.E$, $a \xrightarrow{\text{hb}'_0} b$. There are then four cases to consider: 1) $a \xrightarrow{G.\text{po}_1 \cup \text{rf}_1 \cup \text{lo}_1} b$; or 2) $a \xrightarrow{G.\text{poE}} b$; 3) $a \xrightarrow{G.\text{rfE}} b$; or 4) $a \xrightarrow{G.\text{loE}} b$.

In case (1), from the construction of $G.\text{rf}_1$ we have $G.\text{rf}_1 \subseteq G.\text{po}_1$; moreover, from the construction of $G.\text{lo}_1$ we have $G.\text{lo}_1 \subseteq G.\text{po}_1$. Consequently we have $a \xrightarrow{G.\text{po}_1} b$, as required.

In case (2) from the construction of $G.\text{po}$ we then know that $\mathcal{T}_{\xi_a} \times \mathcal{T}_{\xi_b} \subseteq G'.\text{po}$ and thus from the definition of TCO we have $(\xi_a, \xi_b) \in \text{TCO}$, as required.

In case (3) from the construction of $G.\text{rf}$ we then know that $\mathcal{T}_{\xi_a} \times \mathcal{T}_{\xi_b} \subseteq G'.\text{rf}_\top$ and thus from the definition of TCO we have $(\xi_a, \xi_b) \in \text{TCO}$, as required.

In case (4) from the definition of hb'_0 we know that there exists i, j, i', j' such that $\xi_a = P_i^x|_j$, $\xi_b = P_{i'}^x|_{j'}$, and either i) $i < i'$; or ii) $i = i' \wedge j < j'$.

In case (4.i) from [Lemma 10](#) we have $(\xi_a, \xi_b) \in \text{TCO}$, as required. In case (4.ii) from the definitions of P_i^x we have $(\xi_a, \xi_b) \in \text{TCO}$, as required.

Inductive case $i = n+1$

Pick arbitrary a, b, ξ_a, ξ_b such that $a \in \theta_{\xi_a}.E$, $b \in \theta_{\xi_b}.E$ and $a \xrightarrow{\text{hb}'_i} b$.

$$\begin{aligned} \forall j \leq n. \forall c, d, \xi_c, \xi_d. c \in \theta_{\xi_c}.E \wedge d \in \theta_{\xi_d}.E \wedge c \xrightarrow{\text{hb}'_j} d \\ \Rightarrow (c, d) \in G.\text{po}_1 \vee (\xi_c, \xi_d) \in \text{TCO} \end{aligned} \quad (\text{I.H.})$$

From the definition of hb'_{n+1} we then know there exists e such that $(a, e) \in \text{hb}'_0$ and $(e, b) \in \text{hb}'_n$. Let $e \in \xi_e.E$. Consequently, from the proof of the base case we then know that $(a, e) \in G.\text{po}_1 \vee (\xi_a, \xi_e) \in \text{TCO}$. Similarly, from (I.H.) we have $(e, b) \in G.\text{po}_1 \vee (\xi_e, \xi_b) \in \text{TCO}$. There are now four cases to consider: 1) $(a, e) \in G.\text{po}_1$ and $(e, b) \in G.\text{po}_1$; 2) $(a, e) \in G.\text{po}_1$ and $(\xi_e, \xi_b) \in \text{TCO}$; 3) $(\xi_a, \xi_e) \in \text{TCO}$ and $(e, b) \in G.\text{po}_1$; 4) $(\xi_a, \xi_e) \in \text{TCO}$ and $(\xi_e, \xi_b) \in \text{TCO}$.

In case (1) from the definition of $G.\text{po}_1$ we have $(a, b) \in G.\text{po}_1$. In cases (2-4) from the definitions of $G.\text{po}_1$, TCO we have $(\xi_a, \xi_b) \in \text{TCO}$, as required.

Given an implementation graph $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ constructed as above, let $\text{stg}(\cdot) : G.E \rightarrow \{1, 2, 3\}$ denote the *event stage* as follows:

$$\text{stg}(e) \triangleq \begin{cases} 1 & \exists \xi. e \in \xi.Ts \\ 2 & \exists \xi, \mathbf{x}. e = \xi.ru_{\mathbf{x}} \\ 3 & \text{otherwise} \end{cases}$$

Lemma 12. For all implementation graphs $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ constructed as above, and for all $\mathbf{x} \in \text{LOC}$, $\xi_a, \xi_1, \xi_2 \in \text{TXID}$ and $a \in \theta_{\xi_a}.E$:

$$\left(\begin{array}{l} (\xi_1.rl_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}} a) \\ \vee (\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}} a) \\ \vee (\xi_1.rl_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.wu_{\mathbf{x}} \wedge \xi_2.wu_{\mathbf{x}} \xrightarrow{\text{hb}} a) \\ \vee (\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.wu_{\mathbf{x}} \wedge \xi_2.wu_{\mathbf{x}} \xrightarrow{\text{hb}} a) \end{array} \right) \wedge \exists j, k, l. P_j^{\mathbf{x}}|_k = \xi_1 \wedge P_j^{\mathbf{x}}|_l = \xi_2 \wedge k > l \\ \Rightarrow \text{stg}(\xi_1.ru_{\mathbf{x}}) < \text{stg}(a) \vee (\xi_1, \xi_a) \in \text{TCO}$$

Proof. Pick an arbitrary implementation graph $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ constructed as above, and pick arbitrary $\mathbf{x} \in \text{LOC}$, $\xi_a, \xi_1, \xi_2 \in \text{TXID}$ and $a \in \theta_{\xi_a}.E$ such that $((\xi_1.rl_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}} a) \vee (\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}} a) \vee (\xi_1.rl_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.wu_{\mathbf{x}} \wedge \xi_2.wu_{\mathbf{x}} \xrightarrow{\text{hb}} a) \vee (\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.wu_{\mathbf{x}} \wedge \xi_2.wu_{\mathbf{x}} \xrightarrow{\text{hb}} a))$ and $\exists j, k, l. P_j^{\mathbf{x}}|_k = \xi_1 \wedge P_j^{\mathbf{x}}|_l = \xi_2 \wedge k > l$. In case of the first disjunct from the construction of **lo** we then also have $\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}} a$. Similarly, in case of the third and fourth disjuncts from the construction of **lo** we also have $\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.wu_{\mathbf{x}} \xrightarrow{\text{hb}} a$. Moreover, since we have $\xi_2.pl_{\mathbf{x}} \xrightarrow{\text{po}} \xi_2.wu_{\mathbf{x}}$ and $\text{po} \subseteq \text{hb}$, we also have $\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}} a$. It thus suffices to show:

$$\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}} a \wedge \exists j, k, l. P_j^{\mathbf{x}}|_k = \xi_1 \wedge P_j^{\mathbf{x}}|_l = \xi_2 \wedge k > l \\ \Rightarrow \text{stg}(\xi_1.ru_{\mathbf{x}}) < \text{stg}(a) \vee (\xi_1, \xi_a) \in \text{TCO}$$

Since $G.\text{hb}$ is a transitive closure, it is straightforward to demonstrate that $G.\text{hb} = \bigcup_{i \in \mathbb{N}} \text{hb}_i$, where $\text{hb}_0 = G.\text{po} \cup G.\text{rf} \cup G.\text{lo}$ and $\text{hb}_{i+1} = \text{hb}_0; \text{hb}_i$. It thus suffices to show:

$$\forall i \in \mathbb{N}. \forall \mathbf{x} \in \text{LOC}, \xi_a, \xi_1, \xi_2 \in \text{TXID}, a \in \theta_{\xi_a}.E. \\ \xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}} \wedge \xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}_i} a \wedge \exists j, k, l. P_j^{\mathbf{x}}|_k = \xi_1 \wedge P_j^{\mathbf{x}}|_l = \xi_2 \wedge k > l \\ \Rightarrow \text{stg}(\xi_1.ru_{\mathbf{x}}) < \text{stg}(a) \vee (\xi_1, \xi_a) \in \text{TCO}$$

We thus proceed by induction over i .

Base case $i = 0$

Pick arbitrary $\mathbf{x} \in \text{LOC}$, $\xi_a, \xi_1, \xi_2 \in \text{TXID}$ and $a \in \theta_{\xi_a}.E$ such that $\xi_1.ru_{\mathbf{x}} \xrightarrow{\text{lo}} \xi_2.pl_{\mathbf{x}}$, $\xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}_0} a$ and $\exists j, k, l. P_j^{\mathbf{x}}|_k = \xi_1 \wedge P_j^{\mathbf{x}}|_l = \xi_2 \wedge k > l$. From the construction of G we know that there exists $r \in G'.\mathcal{T}_{\xi_1}$ and $w \in G'.\mathcal{T}_{\xi_2}$ such that $(r, w) \in G'.\text{rb}$. Moreover, since $k \neq l$ we know that $(r, w) \in G'.\text{rb}_{\top}$. Since we have $\xi_2.pl_{\mathbf{x}} \xrightarrow{\text{hb}_0} a$, there are four cases to consider: 1) $\xi_2.pl_{\mathbf{x}} \xrightarrow{G.\text{po}_1 \cup G.\text{rf}_1 \cup G.\text{lo}_1} a$; or 2) $\xi_2.pl_{\mathbf{x}} \xrightarrow{G.\text{po}_E} a$; or 3) $\xi_2.pl_{\mathbf{x}} \xrightarrow{G.\text{rf}_E} a$; or 4) $\xi_2.pl_{\mathbf{x}} \xrightarrow{G.\text{lo}_E} a$.

In case (1), from the construction of $G.\text{rf}_1$ we have $G.\text{rf}_1 \subseteq G.\text{po}_1$; moreover, from the construction of $G.\text{lo}_1$ we have $G.\text{lo}_1 \subseteq G.\text{po}_1$. We thus have $\xi_2.pl_{\mathbf{x}} \xrightarrow{G.\text{po}_1} a$.

On the other hand, from the definition of $\text{stg}(\cdot)$ we know $\text{stg}(a) \geq \text{stg}(\xi_2.pl_x) > \text{stg}(\xi_1.ru_x)$, and thus $\text{stg}(a) > \text{stg}(\xi_1.ru_x)$, as required.

In case (2), since $\xi_2.pl_x \xrightarrow{\text{poE}} a$, we know there exists $a' \in \mathcal{T}_{\xi_a}$ such that $(w, a') \in G'.\text{po}_T$. As such, we have $(r, a') \in G'.(\text{rb}_T; \text{po}_T)$. From the definition of TCO we thus have $(\xi_1, \xi_a) \in \text{TCO}$, as required.

Case (3) cannot happen as there are no **rf** edges from lock events. In case (4), from the construction of **lo** we know there exists m, n such that $m > j$, and $\xi_a = P_m^x|_n$. As such, from [Lemma 10](#) we have $(\xi_1, \xi_a) \in \text{TCO}$, as required.

Inductive case $i = n+1$

Pick arbitrary $x \in \text{LOC}$ and $\xi_a, \xi_1, \xi_2 \in \text{TXID}$ and $a \in \theta_{\xi_a}.E$ such that $\xi_1.ru_x \xrightarrow{\text{lo}} \xi_2.pl_x$, $\xi_2.pl_x \xrightarrow{\text{hb}_j} a$ and $\exists j, k, l. P_j^x|_k = \xi_1 \wedge P_j^x|_l = \xi_2 \wedge k > l$.

$\forall j \leq n. \forall y \in \text{LOC}, \xi_{a'}, \xi_3, \xi_4 \in \text{TXID}, a' \in \theta_{\xi_{a'}}.E$.

$$\begin{aligned} (\xi_3.ru_x \xrightarrow{\text{lo}} \xi_4.pl_x \wedge \xi_4.pl_x \xrightarrow{\text{hb}_j} a') \wedge \exists j, k, l. P_j^x|_k = \xi_1 \wedge P_j^x|_l = \xi_2 \wedge k > l \\ \Rightarrow \text{stg}(\xi_3.ru_y) < \text{stg}(a') \vee (\xi_3, \xi_{a'}) \in \text{TCO} \end{aligned} \quad (\text{I.H.})$$

There are two cases to consider:

- 1) $\xi_2.pl_x \xrightarrow{(\text{po}_1 \cup \text{rf}_1 \cup \text{lo}_1)^+} a$; or
- 2) $\xi_2.pl_x \xrightarrow{(\text{po}_1 \cup \text{rf}_1 \cup \text{lo}_1)^*; (\text{po}_E \cup \text{rf}_E \cup \text{lo}_E); \text{hb}_m} a$, where $m \leq n$.

In case (1), from the construction of $G.\text{rf}_1$ we have $G.\text{rf}_1 \subseteq G.\text{po}_1$; moreover, from the construction of $G.\text{lo}_1$ we have $G.\text{lo}_1 \subseteq G.\text{po}_1$. We thus have $\xi_2.pl_x \xrightarrow{\text{po}_1^+} a$, i.e. $\xi_2.pl_x \xrightarrow{\text{po}_1} a$. As such, from the proof of the base case we have $\text{stg}(a) > \text{stg}(\xi_1.ru_x)$, as required.

In case (2) from the construction of G we know that there exists $r \in \theta_{\xi_1}.E$ and $w \in \theta_{\xi_2}.E$ such that $(r, w) \in G'.\text{rb}$. Moreover, since $k \neq l$, from the construction of G we know that $\xi_1 \neq \xi_2$ and thus we have $(r, w) \in G'.\text{rb}_T$. On the other hand, we know that there exist b, c such that $\xi_2.pl_x \xrightarrow{(\text{po}_1 \cup \text{rf}_1 \cup \text{lo}_1)^*} b \xrightarrow{\text{po}_E \cup \text{rf}_E \cup \text{lo}_E} c \xrightarrow{\text{hb}_m} a$. From the construction of $G.\text{rf}_1$ we have $G.\text{rf}_1 \subseteq G.\text{po}_1$; moreover, from the construction of $G.\text{lo}_1$ we have $G.\text{lo}_1 \subseteq G.\text{po}_1$. We thus have $\xi_2.pl_x \xrightarrow{\text{po}_1^*} b$. As such we have $\xi_2.pl_x \xrightarrow{\text{po}_1^*} b \xrightarrow{\text{po}_E \cup \text{rf}_E \cup \text{lo}_E} c$. Let $c \in \theta_{\xi_c}$. There are now three cases to consider: a) $b \xrightarrow{\text{po}_E} c$; or b) $b \xrightarrow{\text{rf}_E} c$; or c) $b \xrightarrow{\text{lo}_E} c$. We first demonstrate that in all three cases we have $(\xi_1, \xi_c) \in \text{TCO}$.

In case (2.a), since $\xi_2.pl_x \xrightarrow{\text{po}_E} c$, we know there exists $c' \in \mathcal{T}_{\xi_c}$ such that $(w, c') \in G'.\text{po}_T$. As such, we have $(r, c') \in G'.(\text{rb}_T; \text{po}_T)$. From the definition of TCO we thus have $(\xi_1, \xi_c) \in \text{TCO}$. In case (2.b), from the definition of $G.\text{rf}$ we know there exists $c' \in \mathcal{T}_{\xi_c}$ such that $(w, c') \in G'.\text{rf}_T$. As such, we have $(r, c') \in G'.(\text{rb}_T; \text{rf}_T)$. From the definition of TCO we thus have $(\xi_1, \xi_c) \in \text{TCO}$. In case (2.c) from the construction of **lo** we know there exists z such that either $b = \xi_2.pl_z$ or $b = \xi_2.wu_z$, and that there exist p, q such that $\xi_2 \in S_p^z$, $\xi_c \in S_q^z$ and $p < q$. Given the definition of S^z we then know that there exists $w_z \in G'.\mathcal{T}_{\xi_2}$ where either i) there exists $w'_z \in \mathcal{T}_{\xi_c}$ such that $(w_z, w'_z) \in G'.\text{mo}_T$; or ii) there exists $r_z \in \mathcal{T}_{\xi_c}$

such that $(w_z, r_z) \in G'.\mathbf{rf}_T$; or ii) there exists $r_z \in \mathcal{T}_{\xi_c}$ such that $(w_z, r_z) \in G'.(\mathbf{mo}_T; \mathbf{rf}_T)$. That is, we have either $(w, w'_z) \in G'.\mathbf{mo}_T$, or $(w, r_z) \in G'.\mathbf{rf}_T$, or $(w, r_z) \in G'.(\mathbf{mo}_T; \mathbf{rf}_T)$. Moreover, since we have $(r, w) \in G'.\mathbf{rb}_T$, we then have $(r, w'_z) \in G'.(\mathbf{rb}_T; \mathbf{mo}_T)$, or $(r, r_z) \in G'.(\mathbf{rb}_T; \mathbf{rf}_T)$, or $(r, r_z) \in G'.(\mathbf{rb}_T; \mathbf{mo}_T; \mathbf{rf}_T)$. From the definition of TCO we thus have $(\xi_1, \xi_c) \in \text{TCO}$.

Since $c \xrightarrow{\mathbf{hb}^n} a$, there are now two cases to consider: i) $c \xrightarrow{\mathbf{hb}'} a$; or ii) $c \not\xrightarrow{\mathbf{hb}'} a$.

In case (2.i), from [Lemma 11](#) we have $(c, a) \in \text{po}_I \vee (\xi_c, \xi_a) \in \text{TCO}$. As we have $(\xi_1, \xi_c) \in \text{TCO}$, we thus have $(\xi_1, \xi_a) \in \text{TCO}$, as required.

In case (2.ii) let us split the path from c at the first occurrence of a non- \mathbf{hb}' edge. That is, pick $\xi_3, \xi_4, g, h, y, p, q, s, k$ such that $c \xrightarrow{\mathbf{hb}'} \xi_3.g \xrightarrow{\mathbf{lo}} \xi_4.h \xrightarrow{\mathbf{hb}^k} a$, $k < m$, $P_s^y|_p = \mathcal{T}_{\xi_3}$, $P_s^y|_q = \mathcal{T}_{\xi_4}$, $p > q$ and either a) $g = \mathbf{rl}_y \wedge h = \mathbf{pl}_y$; or b) $g = \mathbf{ru}_y \wedge h = \mathbf{pl}_y$; or c) $g = \mathbf{rl}_y \wedge h = \mathbf{wu}_y$; or d) $g = \mathbf{ru}_y \wedge h = \mathbf{wu}_y$. From [Lemma 11](#) we then have $(c, \xi_3.g) \in \text{po}_I \vee (\xi_c, \xi_3) \in \text{TCO}$. As we have $(\xi_1, \xi_c) \in \text{TCO}$, we also have $(\xi_1, \xi_3) \in \text{TCO}$. We next demonstrate that in all cases (2.ii.a-2.ii.d) there exists t such that $t \leq m$ and $\xi_3.\mathbf{ru}_y \xrightarrow{\mathbf{lo}} \xi_4.\mathbf{pl}_y \xrightarrow{\mathbf{hb}^t} a$.

In case (2.ii.a) from the definition of \mathbf{lo} we also have $\xi_3.\mathbf{ru}_y \xrightarrow{\mathbf{lo}} \xi_4.\mathbf{pl}_x$. As such, we have $\xi_3.\mathbf{ru}_y \xrightarrow{\mathbf{lo}} \xi_4.\mathbf{pl}_y \xrightarrow{\mathbf{hb}^k} a$. In case (2.ii.b) the desired result holds immediately.

In cases (2.ii.c-2.ii.d) from the construction of \mathbf{lo} we have $\xi_3.\mathbf{ru}_y \xrightarrow{\mathbf{lo}} \xi_4.\mathbf{pl}_y$. Moreover, since we have $\xi_4.\mathbf{pl}_y \xrightarrow{\mathbf{po}} \xi_4.\mathbf{wu}_y$, we also have $\xi_4.\mathbf{pl}_y \xrightarrow{\mathbf{po}} \xi_4.\mathbf{wu}_y \xrightarrow{\mathbf{hb}^k} a$. As $\text{po} \subseteq \mathbf{hb}$ and \mathbf{hb} is transitively closed, we have $\xi_4.\mathbf{pl}_y \xrightarrow{\mathbf{hb}^{k+1}} a$. As such, we have $\xi_3.\mathbf{ru}_y \xrightarrow{\mathbf{lo}} \xi_4.\mathbf{pl}_y \xrightarrow{\mathbf{hb}^{k+1}} a$. As $k < m$, the desired result holds immediately.

Consequently, from (I.H.) we have $\text{stg}(\xi_3.\mathbf{ru}_y) < \text{stg}(a) \vee (\xi_3, \xi_a) \in \text{TCO}$. In the case of the first disjunct we have $\text{stg}(\xi_3.\mathbf{ru}_y) = \text{stg}(\xi_1.\mathbf{ru}_x) < \text{stg}(a)$, as required. In the case of the second disjunct, since we also have $(\xi_1, \xi_3) \in \text{TCO}$ and TCO is transitively closed, we have $(\xi_1, \xi_a) \in \text{TCO}$, as required.

Theorem 10 (Completeness). For all SI execution graphs G' and their counterpart implementation graphs G constructed as above,

$$\text{si-consistent}(G') \Rightarrow \text{RA-consistent}(G)$$

Proof. Pick an arbitrary SI execution graph G' and its counterpart implementation graph G constructed as above and assume $\text{si-consistent}(G')$ holds. From the definition of $\text{RA-consistent}(G)$ it suffices to show:

1. $\text{irreflexive}(G.\mathbf{hb})$
2. $\text{irreflexive}(G.\mathbf{mo}; G.\mathbf{hb})$
3. $\text{irreflexive}(G.\mathbf{rb}; G.\mathbf{hb})$

RTS. part 1

We proceed by contradiction. Let us assume that there exists a, θ_ξ such that $a \in \theta_\xi.E$ and $(a, a) \in G.\mathbf{hb}$. There are now two cases to consider: 1) $(a, a) \in \mathbf{hb}'$; or 2) $(a, a) \notin \mathbf{hb}'$.

In case (1), from [Lemma 11](#) we have $(a, a) \in G.\text{po}_1 \vee (\mathcal{T}_\xi, \mathcal{T}_\xi) \in \text{TCO}$. The first disjunct leads to a contradiction as the construction of $G.\text{po}$ yields an acyclic relation. The second disjunct leads to a contradiction as TCO is a strict total order.

In case (2), let us split the $a \xrightarrow{\text{hb}} a$ at the first occurrence of a non- hb' edge. That is, pick $\xi_1, \xi_2, \mathbf{x}, i, j, k, g, h$ such that $a \xrightarrow{\text{hb}} \xi_1.g \xrightarrow{\text{lo}\backslash\text{hb}'}$ $\xi_2.h \xrightarrow{\text{hb}} a$, $P_i^x|_j = \mathcal{T}_{\xi_1}$, $P_i^x|_k = \mathcal{T}_{\xi_2}$ and $j > k$. As we have $\xi_1.g \xrightarrow{\text{lo}\backslash\text{hb}'}$ $\xi_2.h \xrightarrow{\text{hb}} a \xrightarrow{\text{hb}} \xi_1.g$, from [Lemma 12](#) and the definition of hb we then have $\text{stg}(\xi_1.g) < \xi_1.g \vee (\xi_1, \xi_1) \in \text{TCO}$, leading to a contradiction in both disjuncts (the second disjunct yields a contradiction as TCO is a strict total order).

RTS. part 2

We proceed by contradiction. Let us assume that there exists $a, \theta_{\xi_a}, b, \theta_{\xi_b}$ such that $a \in \theta_{\xi_a}.E$, $b \in \theta_{\xi_b}.E$, $(a, b) \in G.\text{hb}$ and $(b, a) \in G.\text{mo}$. Let $\text{loc}(a) = \text{loc}(b) = \mathbf{x}$ for some shared location \mathbf{x} . There are now two cases to consider: 1) $(b, a) \in G.\text{mo}_I$; or 2) $(b, a) \in G.\text{mo}_E$.

In case (1) we then have $(b, a) \in G'.\text{mo}_I \subseteq G'.\text{po}_I$. That is, we have $(b, a) \in G.\text{po} \subseteq G.\text{hb}$. We thus have $a \xrightarrow{G.\text{hb}} b \xrightarrow{G.\text{hb}}$ a , contradicting our proof in part 1. In case (2), from the construction of $G.\text{mo}$ we have $(b, a) \in G'.\text{mo}_T$ and thus from the construction of $G.\text{lo}$ we then have $(\xi_b.wu_{\mathbf{x}}, \xi_a.rl_{\mathbf{x}}) \in G.\text{lo}$. As such we have $a \xrightarrow{G.\text{hb}} b \xrightarrow{G.\text{po}} \xi_b.wu_{\mathbf{x}} \xrightarrow{G.\text{lo}} \xi_a.rl_{\mathbf{x}} \xrightarrow{G.\text{po}}$ a . That is, we have $a \xrightarrow{G.\text{hb}}$ a , contradicting our proof in part 1.

RTS. part 3

We proceed by contradiction. Let us assume that there exists $a, \theta_{\xi_a}, b, \theta_{\xi_b}$ such that $a \in \theta_{\xi_a}.E$, $b \in \theta_{\xi_b}.E$, $(a, b) \in G.\text{hb}$ and $(b, a) \in G.\text{rb}$.

Let $\text{loc}(a) = \text{loc}(b) = \mathbf{x}$ for some shared location \mathbf{x} . There are now two cases to consider: 1) $(b, a) \in G.\text{rb}_I$; or 2) $(b, a) \in G.\text{rb}_E$.

In case (1) we then have $(b, a) \in G'.\text{rb}_I \subseteq G'.\text{po}_I$. That is, we have $(b, a) \in G.\text{po} \subseteq G.\text{hb}$. We thus have $a \xrightarrow{G.\text{hb}} b \xrightarrow{G.\text{hb}}$ a , contradicting our proof in part 1. In case (2), from the construction of $G.\text{rb}$ we have $(b, a) \in G'.\text{rb}_T$ and thus from the construction of $G.\text{lo}$ we then have either $(b, \xi_b.wu_{\mathbf{x}}) \in G.\text{po}$ and $(\xi_b.wu_{\mathbf{x}}, \xi_a.rl_{\mathbf{x}}) \in G.\text{lo}$; or $(b, \xi_b.ru_{\mathbf{x}}) \in \text{po}$ and $(\xi_b.ru_{\mathbf{x}}, \xi_a.pl_{\mathbf{x}}) \in G.\text{lo}$. As such in both cases we have $(b, a) \in G.\text{hb}$. Consequently, we have $a \xrightarrow{G.\text{hb}} b \xrightarrow{G.\text{hb}}$ a , contradicting our proof in part 1.

E Soundness and Completeness of the Eager RSI Implementation

Notation Given an execution graph $(E, \text{po}, \text{rf}, \text{mo}, \text{lo})$ we write \mathcal{T}/st for the set of equivalence classes of \mathcal{T} induced by **st**; $[a]_{\text{st}}$ for the equivalence class that contains a ; and \mathcal{T}_ξ for the equivalence class of transaction $\xi \in \text{TXID}$: $\mathcal{T}_\xi \triangleq \{a \mid \text{tx}(a) = \xi\}$. We write $\text{rsi-consistent}(G')$ to denote that G' is RSI-consistent; and write $\text{RA-consistent}(G)$ to denote that G is RA-consistent.

Given an execution graph G of the RSI implementation in Fig. 3, let us assign a transaction identifier to each transaction executed by the program; and given a transaction ξ , let RS_ξ and WS_ξ denote its read and write sets, respectively. Observe that given a transaction ξ of the RSI implementation in Fig. 3 with $\text{RS}_\xi \cup \text{WS}_\xi = \{\mathbf{x}_1, \dots, \mathbf{x}_i\}$, the trace of ξ , written θ_ξ , is of the following form:

$$\theta_\xi = FS^* \xrightarrow{\text{po|imm}} Rs \xrightarrow{\text{po|imm}} RUs \xrightarrow{\text{po|imm}} PLs \xrightarrow{\text{po|imm}} Ts \xrightarrow{\text{po|imm}} Us$$

where

- FS_ξ^* denotes the sequence of events failing to obtain a valid snapshot.
- Rs denotes the sequence of events acquiring a valid snapshot, and is of the form $rl_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} rl_{\mathbf{x}_i} \xrightarrow{\text{po|imm}} S_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} S_{\mathbf{x}_i} \xrightarrow{\text{po|imm}} V_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} V_{\mathbf{x}_i}$, where for all $n \in \{1 \dots i\}$:

$$rl_{\mathbf{x}_n} = \text{RL}(\mathbf{x}_n) \quad S_{\mathbf{x}_n} = \begin{cases} rs_{\mathbf{x}_n} \xrightarrow{\text{po|imm}} ws_{\mathbf{x}_n} & \text{if } \mathbf{x} \in \text{RS}_\xi \\ \emptyset & \text{otherwise} \end{cases} \quad V_{\mathbf{x}_n} = \begin{cases} vs_{\mathbf{x}_n} & \text{if } \mathbf{x} \in \text{RS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

with $rs_{\mathbf{x}_n} \triangleq \text{R}(\mathbf{x}_n, v_n)$, $ws_{\mathbf{x}_n} \triangleq \text{W}(\mathbf{s}[\mathbf{x}_n], v_n)$ and $vs_{\mathbf{x}_n} \triangleq \text{R}(\mathbf{x}_n, v_n)$, for some v_n .

- RUs denotes the sequence of events releasing the reader locks (when the given location is in the read set only), and is of the form $ru_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} ru_{\mathbf{x}_i}$, where for all $n \in \{1 \dots i\}$:

$$ru_{\mathbf{x}_n} = \begin{cases} \text{RU}(\mathbf{x}_n) & \text{if } \mathbf{x}_n \notin \text{WS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- PLs denotes the sequence of events promoting the reader locks to writer ones (when the given location is in the write set), and is of the form $pl_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} pl_{\mathbf{x}_i}$, where for all $n \in \{1 \dots i\}$:

$$pl_{\mathbf{x}_n} = \begin{cases} \text{PL}(\mathbf{x}_n) & \text{if } \mathbf{x}_n \in \text{WS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- Ts denotes the sequence of events corresponding to the execution of (\mathbb{T}) in Fig. 3 and is of the form $t_1 \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} t_k$, where for all $m \in \{1 \dots k\}$:

$$t_m = \begin{cases} \mathbf{R}(\mathbf{s}[\mathbf{x}_n], v_n) & \text{if } O_m = \mathbf{R}(\mathbf{x}_n, v_n) \\ \mathbf{W}(\mathbf{x}_n, v_n) \xrightarrow{\text{po|imm}} \mathbf{W}(\mathbf{s}[\mathbf{x}_n], v_n) & \text{if } O_m = \mathbf{W}(\mathbf{x}_n, v_n) \end{cases}$$

where O_m denotes the m th event in the trace of the original \mathbb{T} ;

- Us denotes the sequence of events releasing the locks on the write set. That is, the events in Us correspond to the execution of the last line of the implementation in Fig. 3, and is of the form $wu_{x_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} wu_{x_i}$, where for all $n \in \{1 \dots i\}$:

$$wu_{x_n} = \begin{cases} \mathbf{WU}(y_{1_n}) & \text{if } \mathbf{x}_n \in \mathbf{WS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

Given a transaction trace θ_ξ , we write e.g. $\xi.Ls$ to refer to its constituent Ls sub-trace and write $Ls.E$ for the set of events related by po in Ls . Similarly, we write $\xi.E$ for the set of events related by po in θ_ξ . Note that $G.E = \bigcup_{\xi \in \text{Tx}} \xi.E$.

E.1 Implementation Soundness

In order to establish the soundness of our implementation, it suffices to show that given an RA-consistent execution graph of the implementation $G = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$, we can construct a corresponding RSI-consistent execution graph G' with the same outcome.

Given a transaction $\xi \in \text{Tx}$ with $\mathbf{RS}_\xi \cup \mathbf{WS}_\xi = \{\mathbf{x}_1 \dots \mathbf{x}_i\}$ and trace $\theta_\xi = F_s^* \xrightarrow{\text{po|imm}} R_s \xrightarrow{\text{po|imm}} RUs \xrightarrow{\text{po|imm}} PLs \xrightarrow{\text{po|imm}} Ts \xrightarrow{\text{po|imm}} Us$, with $Ts = t_1 \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} t_k$, we define $\theta'_\xi \triangleq t'_1 \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} t'_k$, such that for all $m \in \{1 \dots k\}$:

$$\begin{aligned} t'_m &= \mathbf{R}(\mathbf{x}_n, rb_n) \text{ when } t_m = \mathbf{R}(\mathbf{s}[\mathbf{x}_n], rb_n) \\ t'_m &= \mathbf{W}(\mathbf{x}_n, rb_n) \text{ when } t_m = \mathbf{W}(\mathbf{x}_n, rb_n) \xrightarrow{\text{po|imm}} \mathbf{W}(\mathbf{s}[\mathbf{x}_n], rb_n) \end{aligned}$$

such that in the first case the identifier of t'_m is that of t_m ; and in the second case the identifier of t'_m is that of the first event in t_m . We then define:

$$\mathbf{RF}_\xi \triangleq \left\{ (w, t'_j) \left| \begin{array}{l} t'_j \in Ts'_\xi \wedge \exists \mathbf{x}, v. t'_j = \mathbf{R}(\mathbf{x}, v) \wedge w = \mathbf{W}(\mathbf{x}, v) \\ \wedge (w \in \xi.E \Rightarrow w \xrightarrow{\text{po}} t'_j \wedge \\ \quad (\forall e \in \xi.E. w \xrightarrow{\text{po}} e \xrightarrow{\text{po}} t'_j \Rightarrow (\mathbf{loc}(e) \neq \mathbf{x} \vee e \notin \mathcal{W}))) \\ \wedge (w \notin \xi.E \Rightarrow (\forall e \in \xi.E. (e \xrightarrow{\text{po}} t'_j \Rightarrow (\mathbf{loc}(e) \neq \mathbf{x} \vee e \notin \mathcal{W}))) \\ \quad \wedge (w, \xi.rs_x), (w, \xi.vs_x) \in G.\text{rf}) \end{array} \right. \right\}$$

We are now in a position to demonstrate the soundness of our implementation. Let $G.\mathcal{NT} \triangleq G.E \setminus (\bigcup_{\xi \in \text{Tx}} \theta_\xi.E)$. Given an RA-consistent execution graph G of the implementation, we construct an RSI execution graph G' as follows and demonstrate that $\text{rsi-consistent}(G')$ holds.

– $G'.E = \bigcup_{\xi \in \text{Tx}} \theta'_\xi.E$, with the $\text{tx}(\cdot)$ function defined as:

$$\text{tx}(a) \triangleq \begin{cases} \xi & \text{if } a \in \theta'_\xi \\ 0 & \text{otherwise} \end{cases}$$

– $G'.\text{po} = G.\text{po}|_{G'.E}$
– $G'.\text{rf} = \left(\bigcup_{\xi \in \text{Tx}} \text{RF}_\xi\right) \cup (G.\text{rf} \cap G.E \times G.\mathcal{NT})$
– $G'.\text{mo} = G.\text{mo}$
– $G'.\text{lo} = \emptyset$

Observe that the events of each θ'_ξ trace coincides with those of the equivalence class \mathcal{T}_ξ of G' . That is, $\theta'_\xi.E = \mathcal{T}_\xi$.

Lemma 13. *Given an RA-consistent execution graph G of the implementation and its corresponding RSI execution graph G' constructed as above, for all $a, b, \xi_a, \xi_b, \mathbf{x}$:*

$$\xi_a \neq \xi_b \wedge a \in \xi_a.E \wedge b \in \xi_b.E \wedge \text{loc}(a) = \text{loc}(b) = \mathbf{x} \Rightarrow$$

$$((a, b) \in G'.\text{rf} \Rightarrow \xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x) \quad (22)$$

$$\wedge ((a, b) \in G'.\text{mo} \Rightarrow \xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x) \quad (23)$$

$$\wedge ((a, b) \in G'.\text{mo}; \text{rf} \Rightarrow \xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x) \quad (24)$$

$$\wedge ((a, b) \in G'.\text{rb} \Rightarrow (\mathbf{x} \in \text{WS}_{\xi_a} \wedge \xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x) \vee (\mathbf{x} \notin \text{WS}_{\xi_a} \wedge \xi_a.ru_x \xrightarrow{G.\text{hb}} \xi_b.pl_x)) \quad (25)$$

Proof. Pick an arbitrary RA-consistent execution graph G of the implementation and its corresponding RSI execution graph G' constructed as above. Pick an arbitrary $a, b, \xi_a, \xi_b, \mathbf{x}$ such that $\xi_a \neq \xi_b$, $a \in \xi_a.E$, $a \in \xi_a.E$, and $\text{loc}(a) = \text{loc}(b) = \mathbf{x}$.

RTS. (22)

Assume $(a, b) \in G'.\text{rf}$. From the definition of $G'.\text{rf}$ we then know $(a, \xi_a.rs_x), (a, \xi_a.vs_x) \in G.\text{rf}$. On the other hand, from Lemma 2 we know that either i) $\mathbf{x} \in \text{WS}_{\xi_b}$ and $\xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x$; or ii) $\mathbf{x} \notin \text{WS}_{\xi_b}$ and $\xi_b.ru_x \xrightarrow{G.\text{hb}} \xi_a.pl_x$; or iii) $\xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x$. In case (i) we then have $a \xrightarrow{G.\text{rf}} \xi_a.rs_x \xrightarrow{G.\text{po}} \xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x \xrightarrow{G.\text{po}} a$. That is, we have $a \xrightarrow{G.\text{hbloc}} a$, contradicting the assumption that G is RA-consistent. Similarly in case (ii) we have $a \xrightarrow{G.\text{rf}} \xi_a.rs_x \xrightarrow{G.\text{po}} \xi_b.ru_x \xrightarrow{G.\text{hb}} \xi_a.pl_x \xrightarrow{G.\text{po}} a$. That is, we have $a \xrightarrow{G.\text{hbloc}} a$, contradicting the assumption that G is RA-consistent. In case (iii) the desired result holds trivially.

RTS. (23)

Assume $(a, b) \in G'.\text{mo}$. From the definition of $G'.\text{mo}$ we then know $(a, b) \in G.\text{mo}$. On the other hand, from Lemma 2 we know that either i) $\xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x$; or ii) $\xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x$. In case (i) we then have $a \xrightarrow{G.\text{mo}} b \xrightarrow{G.\text{po}} \xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x \xrightarrow{G.\text{po}} a$.

That is, we have $a \xrightarrow{G.\text{hb}_{loc}} a$, contradicting the assumption that G is RA-consistent. In case (ii) the desired result holds trivially.

RTS. (24)

Assume $(a, b) \in G'.\text{mo}; \text{rf}$. We then know there exists w such that $(a, w) \in G'.\text{mo}$ and $(w, b) \in G'.\text{rf}$. From the definition of $G'.\text{mo}$ we then know $(a, w) \in G.\text{mo}$. There are now two cases to consider: 1) $w \in \xi_b$; or 2) $w \notin \xi_b$. In case (1) since $(a, w) \in G'.\text{mo}$ the desired result follows from part 23.

In case (2) from the definition of $G'.\text{rf}$ we know that $(w, \xi_b.\text{rs}_x) \in G.\text{rf}$. From Lemma 2 we know that either i) $x \in \text{WS}_{\xi_b}$ and $\xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x$; or ii) $x \notin \text{WS}_{\xi_b}$ and $\xi_b.ru_x \xrightarrow{G.\text{hb}} \xi_a.pl_x$; or iii) $\xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x$. In case (i) we then have $a \xrightarrow{G.\text{mo}} w \xrightarrow{G.\text{rf}} \xi_a.\text{rs}_x \xrightarrow{G.\text{po}} \xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x \xrightarrow{G.\text{po}} a$. That is, we have $a \xrightarrow{G.\text{mo}} w \xrightarrow{G.\text{hb}_{loc}} a$, contradicting the assumption that G is RA-consistent. Similarly in case (ii) we have $a \xrightarrow{G.\text{mo}} w \xrightarrow{G.\text{rf}} \xi_a.\text{rs}_x \xrightarrow{G.\text{po}} \xi_b.ru_x \xrightarrow{G.\text{hb}} \xi_a.pl_x \xrightarrow{G.\text{po}} a$. That is, we have $a \xrightarrow{G.\text{mo}} w \xrightarrow{G.\text{hb}_{loc}} a$, contradicting the assumption that G is RA-consistent. In case (iii) the desired result holds trivially.

RTS. (25)

Assume $(a, b) \in G'.\text{rb}$. From the definition of $G'.\text{rb}$ we know $(\xi_a.\text{rs}_x, b), (\xi_a.\text{vs}_x, b) \in G.\text{rb}$. On the other hand, from Lemma 2 we know that either i) $\xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x$; or ii) $x \notin \text{WS}_{\xi_a}$ and $\xi_a.ru_x \xrightarrow{G.\text{hb}} \xi_a.pl_x$; or iii) $x \in \text{WS}_{\xi_a}$ and $\xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x$. In case (i) we then have $b \xrightarrow{G.\text{po}} \xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x \xrightarrow{G.\text{po}} \xi_a.\text{rs}_x \xrightarrow{G.\text{rb}} b$. That is, we have $b \xrightarrow{G.\text{hb}_{loc}} \xi_a.\text{rs}_x \xrightarrow{G.\text{rb}} b$, contradicting the assumption that G is RA-consistent. In cases (ii-iii) the desired result holds trivially.

Let

$$\text{rsi-hb}' \triangleq \text{rsi-hb} \setminus (\text{si-rb}; \text{rsi-hb}')$$

Lemma 14. For all RA-consistent execution graphs G of the implementation and their counterpart RSI execution graphs G' constructed as above,

1. $(G'.\text{rsi-po} \subseteq G.\text{hb}) \wedge (G'.(\text{rsi-po}; \text{si-rb}) \subseteq G.\text{hb})$
2. $(G'.\text{rsi-mo} \subseteq G.\text{hb}) \wedge (G'.(\text{rsi-mo}; \text{si-rb}) \subseteq G.\text{hb})$
3. $(G'.\text{rsi-rf} \subseteq G.\text{hb}) \wedge (G'.(\text{rsi-rf}; \text{si-rb}) \subseteq G.\text{hb})$
4. $(G'.\text{rsi-hb}' \subseteq G.\text{hb})$

Proof (Proof (Part 1)). Pick an arbitrary RA-consistent execution graph G of the implementation and its counterpart RSI execution graph G' constructed as above. Pick arbitrary $(a, b) \in G'.\text{rsi-po}$ and $(b, c) \in G'.\text{si-rb}$. There are then two cases to consider: 1) $(a, b) \notin G'.\text{rsi-po}_1$; or 2) $(a, b) \in G'.\text{rsi-po}_1$.

In case (1) from the construction of rsi-po we then have $(a, b) \in G.\text{po} \subseteq G.\text{hb}$, as required by the first conjunct. For the second conjunct, we know there exists ξ_b, ξ_c, r, w such that $b, r \in \xi_b.\mathcal{R}_E$, $c, w \in \xi_c.\mathcal{W}$, and $(r, w) \in G'.\text{rb}$. Let $\text{loc}(r) = \text{loc}(w) = \mathbf{x}$. Given the definition of \mathcal{R}_E and the construction of G' we know $(\xi_b.\text{rs}_\mathbf{x}, w) \in G.\text{rb}$.

As we have $(a, b) \notin G'.\text{rsi-po}_I$ and b and $\xi_b.\text{rs}_\mathbf{x}$ are both transactional events in ξ_b , from the definition of $G'.\text{po}$ and $G.\text{po}$ we then know $(a, \xi_b.\text{rs}_\mathbf{x}) \in G.\text{po}$. On the other hand, from [Lemma 13](#) we have either *i*) $\mathbf{x} \in \text{WS}_{\xi_b}$ and $\xi_b.wu_\mathbf{x} \xrightarrow{G.\text{hb}} \xi_c.\text{rl}_\mathbf{x}$; or *ii*) $\mathbf{x} \notin \text{WS}_{\xi_b}$ and $\xi_b.ru_\mathbf{x} \xrightarrow{G.\text{hb}} \xi_c.\text{pl}_\mathbf{x}$. In case (i) we have $a \xrightarrow{G.\text{po}} \text{rs}_\mathbf{x} \xrightarrow{G.\text{po}} \xi_a.wu_\mathbf{x} \xrightarrow{G.\text{hb}} \xi_b.\text{rl}_\mathbf{x} \xrightarrow{G.\text{po}} c$, i.e. $a \xrightarrow{G.\text{hb}} c$, as required. Similarly, in case (ii) have $a \xrightarrow{G.\text{po}} \text{rs}_\mathbf{x} \xrightarrow{G.\text{po}} \xi_a.ru_\mathbf{x} \xrightarrow{G.\text{hb}} \xi_b.\text{pl}_\mathbf{x} \xrightarrow{G.\text{po}} c$, i.e. $a \xrightarrow{G.\text{hb}} c$, as required.

In case (2) from the definition of $G'.\text{rsi-po}_I$ we then know that $a, b \in G'.\mathcal{W}$ and thus from the construction of $G'.\text{rsi-po}_I$ we have $(a, b) \in G.\text{po} \subseteq G.\text{hb}$, as required by the first conjunct. For the second conjunct, since we have $b \in G'.\mathcal{W}$, we cannot have $(b, c) \in G'.\text{si-rb}$ and thus the desired result holds vacuously.

Proof (Proof (Part 2)). Pick an arbitrary RA-consistent execution graph G of the implementation and its counterpart RSI execution graph G' constructed as above. For the first conjunct pick arbitrary $(a, b) \in G'.\text{rsi-mo}$. From the definition of $G'.\text{rsi-mo}$ we then know there exists w, w', ξ_a, ξ_b such that $\xi_a \neq \xi_b$, $w, a \in \xi_a.E$, $w', b \in \xi_b.E$ and $(w, w') \in G'.\text{mo}$. From the definition of $G'.\text{mo}$ we then have $(w, w') \in G.\text{mo}$. Let $\text{loc}(w) = \text{loc}(w') = \mathbf{x}$. From [Lemma 13](#) we then know that $\xi_a.wu_\mathbf{x} \xrightarrow{G.\text{hb}} \xi_b.\text{rl}_\mathbf{x}$. We then have $a \xrightarrow{G.\text{po}} \xi_a.wu_\mathbf{x} \xrightarrow{G.\text{hb}} \xi_b.\text{rl}_\mathbf{x} \xrightarrow{G.\text{po}} b$, i.e. $a \xrightarrow{G.\text{hb}} b$, as required.

For the second conjunct pick arbitrary c such that $(b, c) \in \text{si-rb}$. From the definition of si-rb and the construction of G' we then know there exist $\mathbf{y}, r_y, w_y, \xi_c$ such that $\xi_c \neq \xi_b$, $\text{loc}(w_y) = \text{loc}(r_y) = \mathbf{y}$, $(r_y, w_y) \in G'.\text{rb}$, and $w_y, c \in \xi_c.\mathcal{W}$. As we demonstrated for the first conjunct we have $a \xrightarrow{G.\text{po}} \xi_a.wu_\mathbf{x} \xrightarrow{G.\text{hb}} \xi_b.\text{rl}_\mathbf{x}$. That is, $a \xrightarrow{G.\text{hb}} \xi_b.\text{rl}_\mathbf{x}$. On the other hand, from [Lemma 13](#) we then know that either *i*) $\mathbf{y} \in \text{WS}_{\xi_b}$ and $\xi_b.wu_\mathbf{y} \xrightarrow{G.\text{hb}} \xi_c.\text{rl}_\mathbf{y}$; or *ii*) $\mathbf{y} \notin \text{WS}_{\xi_b}$ and $\xi_b.ru_\mathbf{y} \xrightarrow{G.\text{hb}} \xi_c.\text{pl}_\mathbf{y}$. In case (i) we have $a \xrightarrow{G.\text{hb}} \xi_b.\text{rl}_\mathbf{x} \xrightarrow{G.\text{po}} \xi_a.wu_\mathbf{y} \xrightarrow{G.\text{hb}} \xi_b.\text{rl}_\mathbf{y} \xrightarrow{G.\text{po}} c$, i.e. $a \xrightarrow{G.\text{hb}} c$, as required. Similarly, in case (ii) we have $a \xrightarrow{G.\text{hb}} \xi_b.\text{rl}_\mathbf{x} \xrightarrow{G.\text{po}} \xi_a.ru_\mathbf{y} \xrightarrow{G.\text{hb}} \xi_b.\text{pl}_\mathbf{y} \xrightarrow{G.\text{po}} c$, i.e. $a \xrightarrow{G.\text{hb}} c$, as required.

Proof (Proof (Part 3)). Pick an arbitrary RA-consistent execution graph G of the implementation and its counterpart RSI execution graph G' constructed as above. It suffices to show that:

$$G'.([\mathcal{N}\mathcal{T}]; \text{rf}; \text{st}) \subseteq G.\text{hb} \quad (26)$$

$$G'.([\mathcal{N}\mathcal{T}]; \text{rf}; \text{st}); G'.\text{si-rb} \subseteq G.\text{hb} \quad (27)$$

$$G'.\mathbf{rf}_T \subseteq G.\mathbf{hb} \quad (28)$$

$$G'.\mathbf{rf}_T; G'.\mathbf{si}\text{-}\mathbf{rb} \subseteq G.\mathbf{hb} \quad (29)$$

$$G'.(\mathbf{mo}; \mathbf{rf})_T \subseteq G.\mathbf{hb} \quad (30)$$

$$G'.(\mathbf{mo}; \mathbf{rf})_T; G'.\mathbf{si}\text{-}\mathbf{rb} \subseteq G.\mathbf{hb} \quad (31)$$

$$G'.(\mathbf{rf}; [\mathcal{NT}]) \subseteq G.\mathbf{hb} \quad (32)$$

$$G'.(\mathbf{rf}; [\mathcal{NT}]); G'.\mathbf{si}\text{-}\mathbf{rb} \subseteq G.\mathbf{hb} \quad (33)$$

RTS. (26)

Pick arbitrary $(w, a) \in G'.([\mathcal{NT}]; \mathbf{rf}; \mathbf{st})$ where $a \in \xi_a.E$. Let $\text{loc}(w) = \mathbf{x}$. From the construction of $G'.\mathbf{rf}$ we then know $(w, \xi_a.rs_x), (w, \xi_a.vs_x) \in G.\mathbf{rf}$ and $(\xi_a.rs_x, a), (\xi_a.vs_x, a) \in G.\text{po}$. As such we have $w \xrightarrow{G.\mathbf{rf}} \xi_a.rs_x \xrightarrow{G.\text{po}} a$, i.e. $w \xrightarrow{G.\mathbf{hb}} a$, as required.

RTS. (27)

Pick arbitrary $(w, c) \in G'.([\mathcal{NT}]; \mathbf{rf}; \mathbf{st}); G'.\mathbf{si}\text{-}\mathbf{rb}$. We then know there exist a, ξ_a, ξ_c such that $(w, a) \in G'.([\mathcal{NT}]; \mathbf{rf}; \mathbf{st})$, $(a, c) \in G'.\mathbf{si}\text{-}\mathbf{rb}$. $a \in \xi_a.E$, $c \in \xi_c.E$ and $\xi_a \neq \xi_c$. Let $\text{loc}(w) = \mathbf{x}$. As we demonstrated in the previous part we then know $(w, \xi_a.rs_x), (w, \xi_a.vs_x) \in G.\mathbf{rf}$ and $(\xi_a.rs_x, a), (\xi_a.vs_x, a) \in G.\text{po}$. Moreover, from the definition of $\mathbf{si}\text{-}\mathbf{rb}$ and the construction of G' we know there exist y, r_y, w_y, ξ_c such that $\text{loc}(w_y) = \text{loc}(r_y) = \mathbf{y}$, $(r_y, w_y) \in G'.\mathbf{rb}$, and $w_y, c \in \xi_c.W$. From the construction of $G'.\mathbf{rb}$ we then know that $(\xi_a.rs_y, w_y), (\xi_a.vs_y, w_y) \in G.\mathbf{rb}$. On the other hand, from Lemma 13 we know that either i) $\mathbf{y} \in \text{WS}_{\xi_a}$ and $\xi_a.wu_y \xrightarrow{G.\mathbf{hb}} \xi_c.rl_y$; or ii) $\mathbf{y} \notin \text{WS}_{\xi_a}$ and $\xi_a.ru_y \xrightarrow{G.\mathbf{hb}} \xi_c.pl_y$. In case (i) we have $w \xrightarrow{G.\mathbf{rf}} \xi_a.rs_x \xrightarrow{G.\text{po}} \xi_a.vs_y \xrightarrow{G.\text{po}} \xi_a.wu_y \xrightarrow{G.\mathbf{hb}} \xi_b.rl_y \xrightarrow{G.\text{po}} c$, i.e. $a \xrightarrow{G.\mathbf{hb}} c$, as required. Similarly, in case (ii) we have $w \xrightarrow{G.\mathbf{rf}} \xi_a.rs_x \xrightarrow{G.\text{po}} \xi_a.vs_y \xrightarrow{G.\text{po}} \xi_a.ru_y \xrightarrow{G.\mathbf{hb}} \xi_b.pl_y \xrightarrow{G.\text{po}} c$, i.e. $w \xrightarrow{G.\mathbf{hb}} c$, as required.

RTS. (28)

Pick arbitrary $(a, b) \in G'.\mathbf{rf}_T$. We then know there exist w, r, ξ_a, ξ_b such that $w, a \in \xi_a.E$, $r, b \in \xi_b.E$, $\xi_a \neq \xi_b$ and $(w, r) \in G'.\mathbf{rf}$. Let $\text{loc}(w) = \text{loc}(r) = \mathbf{x}$. From Lemma 13 we then know that $\xi_a.wu_x \xrightarrow{G.\mathbf{hb}} \xi_b.rl_x$. Consequently from the structure of G and the construction of G' we have $a \xrightarrow{G.\text{po}} \xi_a.wu_x \xrightarrow{G.\mathbf{hb}} \xi_b.rl_x \xrightarrow{G.\text{po}} b$, i.e. $a \xrightarrow{G.\mathbf{hb}} b$, as required.

RTS. (29)

Pick arbitrary $(a, c) \in G'.\mathbf{rf}_T; G'.\mathbf{si}\text{-}\mathbf{rb}$. b such that $(a, b) \in G'.\mathbf{rf}_T$ and $(b, c) \in G'.\mathbf{si}\text{-}\mathbf{rb}$. From the definition of $G'.\mathbf{rf}_T$ we then know there exist w, r, ξ_a, ξ_b such that $w, a \in \xi_a.E$, $r, b \in \xi_b.E$, $\xi_a \neq \xi_b$ and $(w, r) \in G'.\mathbf{rf}$. Let $\text{loc}(w) = \text{loc}(r) = \mathbf{x}$. From Lemma 13 we then know that $\xi_a.wu_x \xrightarrow{G.\mathbf{hb}} \xi_b.rl_x$.

On the other hand, from the definition of $G'.\mathbf{si}\text{-}\mathbf{rb}$ we know there exist r', w', ξ_c such that $r' \in \xi_b.E$, $w', c \in \xi_c.W$, $\xi_c \neq \xi_b$ and $(r', w') \in G'.\mathbf{rb}$. Let $\text{loc}(w') =$

$\text{loc}(r') = y$. From [Lemma 13](#) we then know that either i) $y \in \text{WS}_{\xi_b}$ and $\xi_b.wu_y \xrightarrow{G.\text{hb}} \xi_c.rl_y$; or ii) $y \notin \text{WS}_{\xi_b}$ and $\xi_b.ru_y \xrightarrow{G.\text{hb}} \xi_c.pl_y$.

In case (i) we have $a \xrightarrow{G.\text{po}} \xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x \xrightarrow{G.\text{po}} \xi_b.wu_y \xrightarrow{G.\text{hb}} \xi_c.rl_y \xrightarrow{G.\text{po}} c$. That is, $(a, c) \in G.\text{hb}$, as required. Similarly, In case (ii) we have $a \xrightarrow{G.\text{po}} \xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x \xrightarrow{G.\text{po}} \xi_b.ru_y \xrightarrow{G.\text{hb}} \xi_c.pl_y \xrightarrow{G.\text{po}} c$. That is, $(a, c) \in G.\text{hb}$, as required.

RTS. (30)

Pick arbitrary $(a, b) \in G'.(\text{mo}; \text{rf})_{\top}$. We then know there exist ξ_a, ξ_b, w, w', r such that $w, a \in \xi_a.E$, $r, b \in \xi_b.E$, $\xi_a \neq \xi_b$, $(w, w') \in G'.\text{mo}$ and $(w', r) \in G'.\text{rf}$.

From [Lemma 13](#) we then have $\xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x$. Consequently from the structure of G and the construction of G' we have $a \xrightarrow{G.\text{po}} \xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x \xrightarrow{G.\text{po}} b$, i.e. $a \xrightarrow{G.\text{hb}} b$, as required.

RTS. (31)

Pick arbitrary $(a, c) \in G'.(\text{mo}; \text{rf})_{\top}$; $G'.\text{si-rb}$. We then know there exist b such that $(a, b) \in G'.(\text{mo}; \text{rf})_{\top}$ and $(b, c) \in G'.\text{si-rb}$. From the definition of $G'.(\text{mo}; \text{rf})_{\top}$ we then know there exist ξ_a, ξ_b, w, w', r such that $w, a \in \xi_a.E$, $r, b \in \xi_b.E$, $\xi_a \neq \xi_b$, $(w, w') \in G'.\text{mo}$ and $(w', r) \in G'.\text{rf}$. From [Lemma 13](#) we then have $\xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x$.

On the other hand, from the definition of $G'.\text{si-rb}$ we know there exist r'', w'', ξ_c such that $r'' \in \xi_b.E$, $w'', c \in \xi_c.W$, $\xi_c \neq \xi_b$ and $(r'', w'') \in G'.\text{rb}$. Let $\text{loc}(w'') = \text{loc}(r'') = y$. From [Lemma 13](#) we then know that either i) $y \in \text{WS}_{\xi_b}$ and $\xi_b.wu_y \xrightarrow{G.\text{hb}} \xi_c.rl_y$; or ii) $y \notin \text{WS}_{\xi_b}$ and $\xi_b.ru_y \xrightarrow{G.\text{hb}} \xi_c.pl_y$.

In case (i) we then have $a \xrightarrow{G.\text{po}} \xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x \xrightarrow{G.\text{po}} \xi_b.wu_y \xrightarrow{G.\text{hb}} \xi_c.rl_y \xrightarrow{G.\text{po}} c$. That is, $(a, c) \in G.\text{hb}$, as required. Similarly, In case (ii) we then have $a \xrightarrow{G.\text{po}} \xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x \xrightarrow{G.\text{po}} \xi_b.ru_y \xrightarrow{G.\text{hb}} \xi_c.pl_y \xrightarrow{G.\text{po}} c$. That is, $(a, c) \in G.\text{hb}$, as required.

RTS. (32)

Pick arbitrary $(w, r) \in G'.(\text{rf}; [\mathcal{N}\mathcal{T}])$. As $r \in G'.\mathcal{N}\mathcal{T}$, from the construction of $G'.\text{rf}$ we have $(w, r) \in G.\text{rf} \subseteq G.\text{hb}$, as required.

RTS. (33)

The desired result holds trivially as $G'.(\text{rf}; [\mathcal{N}\mathcal{T}]); G'.\text{si-rb} = \emptyset$.

Proof (Proof (part 4)). Let $\text{rsi-hb}_0 \triangleq \text{rsi-po} \cup \text{rsi-rf} \cup \text{rsi-mo} \cup \text{si-rb}$ and $\text{rsi-hb}_{n+1} \triangleq \text{rsi-hb}_0$; rsi-hb_n , for all $n \geq 0$. Similarly, let $\text{rsi-hb}'_0 \triangleq \text{rsi-hb}_0 \setminus \text{si-rb}$ and $\text{rsi-hb}'_{n+1} \triangleq \text{rsi-hb}'_0$; $\text{rsi-hb}'_n$, for all $n \geq 0$. It is then straightforward to demonstrate that $\text{rsi-hb}' \triangleq \bigcup_{n \in \mathbb{N}} \text{rsi-hb}'_n$. It thus suffices to show that:

$$\forall n \in \mathbb{N}. \text{rsi-hb}'_n \subseteq G.\text{hb}$$

We proceed by induction on n .

Base case $n = 0$

Follows immediately from the definition of $\text{rsi-hb}'_0$ and the results established in 1-3.

Inductive case $n = m+1$

$$\forall i \in \mathbb{N}. i < n \Rightarrow \text{rsi-hb}'_i \subseteq G.\text{hb} \quad (\text{I.H.})$$

Pick arbitrary $(a, b) \in \text{rsi-hb}'_n$. From the definition of $\text{rsi-hb}'_n$ we then know there exists c such that $(a, c) \in \text{rsi-hb}'_0$ and $(c, b) \in \text{rsi-hb}_m$. Let $\text{rsi-hb}_{-1} \triangleq \text{id}$. There are now two cases to consider: 1) $(c, b) \in \text{rsi-hb}'_0$; rsi-hb_{m-1} ; or $(c, b) \in \text{si-rb}$; rsi-hb_{m-1} . In case (1) from the proof of base case we have $(a, c) \in G.\text{hb}$. On the other hand from the definition of $\text{rsi-hb}'_m$ we have $(c, b) \in \text{rsi-hb}'_m$ and thus from (I.H.) we have $(c, b) \in G.\text{hb}$. Consequently, since $G.\text{hb}$ is transitively closed we have $(a, b) \in G.\text{hb}$ as required.

In case (2) we know there exists d such that $(c, d) \in \text{si-rb}$; and $(d, b) \in \text{rsi-hb}_{m-1}$. Since we have $(a, c) \in \text{rsi-hb}'_0$ and $(c, d) \in \text{si-rb}$, from the definition of $\text{rsi-hb}'_0$ and the proofs of parts 1-3 we have $(a, d) \in G.\text{hb}$. Moreover we either have i) $m = 0$; or ii) $m > 0$. In case (2.i) since $(d, b) \in \text{rsi-hb}_{m-1}$, from the definition of rsi-hb_{m-1} we have $b = d$ and thus $(a, b) \in G.\text{hb}$, as required. In case (2.ii) from (I.H.) we then have $(d, b) \in G.\text{hb}$. As such, since $G.\text{hb}$ is transitively closed, we have $(a, b) \in G.\text{hb}$, as required.

Lemma 15. For all RA-consistent execution graphs G of the implementation and their counterpart RSI execution graphs G' constructed as above:

$$\begin{aligned} \forall n \in \mathbb{N}. \forall a, b, w, \xi, \mathbf{x}. \\ (a, b) \in G'.\text{rsi-hb}_n \wedge a \in G'.\mathcal{W} \wedge b \in G'.(\mathcal{R} \cap \mathcal{T}_\xi) \Rightarrow \\ (w \in G'.(\mathcal{W} \cap \mathcal{T}_\xi) \Rightarrow (a, w) \in G.\text{hb}) \wedge (\xi.\text{vs}_\mathbf{x} \text{ defined} \Rightarrow (a, \xi.\text{vs}_\mathbf{x}) \in G.\text{hb}) \end{aligned}$$

where $\text{rsi-hb}_0 \triangleq \text{rsi-po} \cup \text{rsi-rf} \cup \text{rsi-mo} \cup \text{si-rb}$ and $\text{rsi-hb}_{n+1} \triangleq \text{rsi-hb}_n$; rsi-hb_0 , for all $n \geq 0$.

Proof. Pick an arbitrary RA-consistent execution graph G of the implementation and its counterpart RSI execution graph G' constructed as above. We then proceed by induction on n .

Base case $n = 0$

Pick arbitrary a, b, ξ, \mathbf{x} such that $(a, b) \in (\text{rsi-hb}_0)$, $a \in G'.\mathcal{W}$, $b \in G'.(\mathcal{R} \cap \mathcal{T}_\xi)$, $\xi.\text{vs}_\mathbf{x}$ defined and $w \in G'.(\mathcal{W} \cap \mathcal{T}_\xi)$. From the definition of rsi-hb_0 we then have either 1) $(a, b) \in G'.\text{rsi-po} \wedge a \notin G'.\mathcal{T}_\xi$; or 2) $(a, b) \in G'.(\text{mo}_\top \cup \text{rf}_\top \cup (\text{mo}; \text{rf})_\top)$; or 3) $(a, b) \in G'.([\mathcal{N}\mathcal{T}]; \text{rf}; \text{st})$.

In case (1) from the definition of $G'.\text{rsi-po}$ we then know that $\{a\} \times \theta_\xi.E \subseteq G.\text{po}$ and thus $(a, \xi.\text{vs}_\mathbf{x}), (a, w) \in G.\text{po} \subseteq G.\text{hb}$, as required.

In case (2) from the definitions of rf_\top , mo_\top , $(\text{mo}; \text{rf})_\top$ we know there exists ξ', c, d, \mathbf{y} such that $a, c \in G'.\mathcal{T}_{\xi'}$, $a, c \in \xi'.E$, $b, d \in G'.\mathcal{T}_\xi$, $b, d \in \xi.E$, $(c, d) \in$

$G'.(\text{moUrf} \cup (\text{mo}; \text{rf}))$ and $\text{loc}(c) = \text{loc}(d) = \mathbf{y}$. As such, from [Lemma 13](#) we know $\xi'.wu_{\mathbf{y}} \xrightarrow{G.\text{hb}} \xi.rl_{\mathbf{y}}$. On the other hand we have $a \xrightarrow{G.\text{po}} \xi'.wu_{\mathbf{y}}$, $\xi.rl_{\mathbf{y}} \xrightarrow{G.\text{po}} \xi.vs_{\mathbf{x}}$ and $\xi.rl_{\mathbf{y}} \xrightarrow{G.\text{po}} w$. As such we have $a \xrightarrow{G.\text{po}} \xi'.wu_{\mathbf{y}} \xrightarrow{G.\text{hb}} \xi.rl_{\mathbf{y}} \xrightarrow{G.\text{po}} \xi.vs_{\mathbf{x}}$, i.e. $a \xrightarrow{G.\text{hb}} \xi.vs_{\mathbf{x}}$, as required. Similarly we have $a \xrightarrow{G.\text{po}} \xi'.wu_{\mathbf{y}} \xrightarrow{G.\text{hb}} \xi.rl_{\mathbf{y}} \xrightarrow{G.\text{po}} w$, i.e. $a \xrightarrow{G.\text{hb}} w$, as required.

In case (3) from the construction of G' we know there exists \mathbf{y} such that $\text{loc}(a) = \mathbf{y}$, and $(a, \xi.rs_{\mathbf{y}}), (a, \xi.vs_{\mathbf{y}}) \in G.\text{rf}$. As such we have $a \xrightarrow{G.\text{rf}} \xi.rs_{\mathbf{y}} \xrightarrow{G.\text{po}} \xi.vs_{\mathbf{x}}$, i.e. $a \xrightarrow{G.\text{hb}} \xi.vs_{\mathbf{x}}$, as required. Moreover, since we have $\xi.rs_{\mathbf{y}} \xrightarrow{G.\text{po}} w$, we have $a \xrightarrow{G.\text{hb}} \xi.rs_{\mathbf{y}} \xrightarrow{G.\text{po}} w$, i.e. $a \xrightarrow{G.\text{hb}} w$, as required.

Inductive case $n = m+1$

Pick arbitrary a, b, ξ, \mathbf{x} such that $(a, b) \in (\text{rsi-hb}_n)$, $a \in G'.\mathcal{W}$, $b \in G'.(\mathcal{R} \cap \mathcal{T}_{\xi})$, $\xi.vs_{\mathbf{x}}$ defined and $w \in G'.(\mathcal{W} \cap \mathcal{T}_{\xi})$. From the definition of rsi-hb_n we then know there exists c such that $(a, c) \in \text{rsi-hb}_m$ and $(c, b) \in \text{rsi-hb}_0$. There are then two cases to consider: 1) $c \in \mathcal{T}_{\xi}$; or 2) $c \notin \mathcal{T}_{\xi}$.

Case (1) leads to contradiction as $(c, b) \in \text{rsi-hb}_0$, $c, b \in \mathcal{T}_{\xi}$, $b \in \mathcal{R}$ and rsi-hb_0 does not include any internal edges to read events.

In case (2) since a is a write event, from the definition of $\text{rsi-hb}'$ we then have $(a, c) \in \text{rsi-hb}'$ and thus from [Lemma 14](#) we have $(a, c) \in G.\text{hb}$. On the other hand, since $c \notin \mathcal{T}_{\xi}$, $(c, b) \in \text{rsi-hb}_0$ and b is a read event, from the definition of rsi-hb_0 we know that either 1) $(c, b) \in G'.\text{rsi-po} \wedge c \notin \mathcal{T}_{\xi}$; or 2) $(c, b) \in G'.(\text{mo}_{\top} \cup \text{rf}_{\top} \cup (\text{mo}; \text{rf})_{\top})$; or 3) $(c, b) \in G'.([\mathcal{N}\mathcal{T}]; \text{rf}; \text{st})$. Following an analogous argument as that in the base case, we then have $c \xrightarrow{G.\text{hb}} \xi.vs_{\mathbf{x}}$ and $c \xrightarrow{G.\text{hb}} w$. As such, we have $a \xrightarrow{G.\text{hb}} c \xrightarrow{G.\text{hb}} \xi.vs_{\mathbf{x}}$, i.e. $a \xrightarrow{G.\text{hb}} \xi.vs_{\mathbf{x}}$, as required. Similarly, we have $a \xrightarrow{G.\text{hb}} c \xrightarrow{G.\text{hb}} w$, i.e. $a \xrightarrow{G.\text{hb}} w$, as required.

Theorem 11 (Soundness). For all execution graphs G of the implementation and their counterpart RSI execution graphs G' constructed as above,

$$\text{RA-consistent}(G) \Rightarrow \text{rsi-consistent}(G')$$

Proof. Pick an arbitrary execution graph G of the implementation such that $\text{RA-consistent}(G)$, and its associated RSI execution graph G' constructed as described above. It then suffices to show 1) $\text{rf}_{\top} \cup \text{mo}_{\top} \cup \text{rb}_{\top} \subseteq \text{po}$; 2) $\text{acyclic}(G'.\text{rsi-hb}_{\text{loc}})$; 3) $\text{acyclic}(G'.(\text{rsi-hb}_{\text{loc}}; \text{mo}))$; and 4) $\text{acyclic}(G'.(\text{rsi-hb}_{\text{loc}}; \text{rb}))$.

RTS. $\text{rf}_{\top} \cup \text{mo}_{\top} \cup \text{rb}_{\top} \subseteq \text{po}$

Follows immediately from the construction of G' and the RA-consistency of G .

RTS. $\text{acyclic}(G'.\text{rsi-hb}_{\text{loc}})$

We proceed by contradiction. Let us assume there exists a such that $(a, a) \in G'.\text{rsi-hb}_{\text{loc}}$. There are then two cases to consider: i) $(a, a) \in \text{rsi-hb}'$; or ii) $(a, a) \in \text{rsi-hb} \setminus \text{rsi-hb}'$. In case (1) from [Lemma 14](#) part 4 we then have $(a, a) \in G.\text{hb}$, contradicting the assumption that G is RA-consistent. In case (2) we

then know there exists b such that $a \in G'.\mathcal{R}_E$, $b \in G'.\mathcal{W}$, $(a, b) \in \text{si-rb}$ and $(b, a) \in \text{rsi-hb}$. Moreover, since $b \in G'.\mathcal{W}$, from the definition of si-rb and $\text{rsi-hb}'$ we have $(b, a) \in \text{rsi-hb}'$. Since we also have $(a, b) \in \text{si-rb} \subseteq \text{rsi-hb}$, from the definition $\text{rsi-hb}'$ we have $(b, b) \in \text{rsi-hb}'$. Consequently, from [Lemma 14](#) part 4 we have $(b, b) \in G'.\text{hb}$, contradicting the assumption that G is RA-consistent.

RTS. $\text{acyclic}(G'.(\text{rsi-hb}_{loc}; \text{mo}))$

We proceed by contradiction. Let us assume there exist a, b such that $(a, b) \in G'.\text{rsi-hb}_{loc}$ and $(b, a) \in G'.\text{mo}$. From the definition of $G'.\text{mo}$ we then know that $a, b \in G'.\mathcal{W}$ and $(a, b) \in G'.\text{mo}$. On the other hand, since $a \in G'.\mathcal{W}$ and $(a, b) \in \text{rsi-hb}$, from the definition of $\text{rsi-hb}'$ we have $(a, b) \in \text{rsi-hb}'$. Consequently, from [Lemma 14](#) part 4 we have $(a, b) \in G'.\text{hb}$. We then have $a \xrightarrow{G'.\text{hb}} b \xrightarrow{G'.\text{mo}} a$, contradicting the assumption that G is RA-consistent.

RTS. $\text{acyclic}(G'.(\text{rsi-hb}_{loc}; \text{rb}))$

We proceed by contradiction. Let us assume there exist w, r such that $(w, r) \in G'.\text{rsi-hb}_{loc}$ and $(r, w) \in G'.\text{rb}$. From the definition of $G'.\text{rb}$ we then know that $w \in G'.\mathcal{W}$, $r \in G'.\mathcal{R}$. Since $w \in G'.\mathcal{W}$, from the definition of $\text{rsi-hb}'$ we have $(w, r) \in \text{rsi-hb}'$. Consequently, from [Lemma 14](#) part 4 we have $(w, r) \in G'.\text{hb}$. Let $\text{loc}(w) = \text{loc}(r) = x$. There are then two cases to consider: 1) $r \in G'.\mathcal{NT}$; or 2) $\exists \xi_r. r \in \xi_r.E$.

In case (1) from the definition of $G'.\text{rb}$ we know $(r, w) \in G'.\text{rb}$. As such, we have $w \xrightarrow{G'.\text{hb}} r \xrightarrow{G'.\text{rb}} w$, contradicting the assumption that G is RA-consistent.

In case (2) we then know there exists w' such that $(w', r) \in G'.\text{rf}$ ($w', w) \in G'.\text{mo}$). There are now three cases to consider: a) $w \in \xi_r.E$; or b) $w, w' \notin \xi_r.E \wedge w' \in \xi_r.E$; or c) $w, w' \notin \xi_r.E$. In case (2.a.) from the construction of G' (in particular, $G'.\text{rf}$) we know that $(r, w) \in G'.\text{po}_1 \subseteq G.\text{po}_1 \subseteq G'.\text{hb}$. As such, we have $w \xrightarrow{G'.\text{hb}} r \xrightarrow{G'.\text{hb}} w$, contradicting our assumption that G is RA-consistent.

In case (2.b), since w' is a write event, $w' \in \xi_r.E$, r is a read event, w is a write event and $(w, r) \in G'.\text{rsi-hb}$, from [Lemma 15](#) we have $(w, w') \in G'.\text{hb}$. Moreover, since $(w', w) \in G'.\text{mo}$, from the definition of $G'.\text{mo}$ we also have $(w', w) \in G'.\text{mo}$. As such we have $w \xrightarrow{G'.\text{hb}} w' \xrightarrow{G'.\text{mo}} w$, contradicting the assumption that G is RA-consistent.

In case (2.c), from the construction of G' we then know $(\xi_r.rs_x, w) \in G'.\text{rb}$ and $(\xi_r.vs_x, w) \in G'.\text{rb}$. On the other hand, since r is a read event, w is a write event and $(w, r) \in G'.\text{rsi-hb}$, from [Lemma 15](#) we have $(w, \xi_r.vs_x) \in G'.\text{hb}$. As such we have $w \xrightarrow{G'.\text{hb}} \xi_r.vs_x \xrightarrow{G'.\text{rb}} w$, contradicting the assumption that G is RA-consistent.

E.2 Implementation Completeness

In order to establish the completeness of our implementation, it suffices to show that given an RSI-consistent execution graph $G' = (E, \text{po}, \text{rf}, \text{mo}, \text{lo})$, we can construct a corresponding RA-consistent execution graph G of the implementation. Before proceeding with the construction of a corresponding implementation graph, we describe several auxiliary definitions.

Given a transaction class $\mathcal{T}_\xi \in G'.\mathcal{T}/\text{st}$, we write WS_ξ for the set of locations written to by \mathcal{T}_ξ : $\text{WS}_\xi = \bigcup_{e \in \mathcal{T}_\xi \cap \mathcal{W}} \text{loc}(e)$. Similarly, we write RS_ξ for the set of locations read from by \mathcal{T}_ξ , *prior to* being written by \mathcal{T}_ξ . For each location \mathbf{x} read from by \mathcal{T}_ξ , we additionally record the first read event in \mathcal{T}_ξ that retrieved the value of \mathbf{x} . That is,

$$\text{RS}_{\mathcal{T}_\xi} \triangleq \left\{ (\mathbf{x}, r) \mid r \in \mathcal{T}_i \cap \mathcal{R}_x \wedge \neg \exists e \in \mathcal{T}_\xi \cap E_x. e \xrightarrow{\text{po}} r \right\}$$

Note that the execution trace for each transaction $\mathcal{T}_\xi \in G'.\mathcal{T}/\text{st}$ is of the form $\theta'_\xi = t'_1 \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} t'_k$, comprising a series of read or write events. As such, we have $G'.E = G'.\mathcal{T} = \bigcup_{\mathcal{T}_\xi \in G'.\mathcal{T}/\text{st}} \mathcal{T}_\xi = \theta'_\xi.E$. Let $\text{RS}_{\mathcal{T}_\xi} \cup \text{WS}_{\mathcal{T}_\xi} = \{\mathbf{x}_1 \dots \mathbf{x}_n\}$. We then construct the implementation trace θ_ξ as:

$$\theta_\xi = Rs \xrightarrow{\text{po|imm}} RUs \xrightarrow{\text{po|imm}} PLs \xrightarrow{\text{po|imm}} Ts \xrightarrow{\text{po|imm}} Us$$

where

- $Rs = rl_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} rl_{\mathbf{x}_n} \xrightarrow{\text{po|imm}} S_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} S_{\mathbf{x}_n} \xrightarrow{\text{po|imm}} V_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} V_{\mathbf{x}_n}$, where the identifiers of all constituent events of Rs are picked fresh, and

$$rl_{\mathbf{x}_j} = \text{RL}(\mathbf{x}_j) \quad S_{\mathbf{x}_j} = \begin{cases} rs_{\mathbf{x}_j} \xrightarrow{\text{po|imm}} ws_{\mathbf{x}_j} & \text{if } \exists r. (\mathbf{x}_j, r) \in \text{RS}_\xi \wedge \text{val}_r(r) = v_j \\ \emptyset & \text{otherwise} \end{cases}$$

$$V_{\mathbf{x}_j} = \begin{cases} vs_{\mathbf{x}_j} & \text{if } \exists r. (\mathbf{x}_j, r) \in \text{RS}_\xi \wedge \text{val}_r(r) = v_j \\ \emptyset & \text{otherwise} \end{cases}$$

with $rs_{\mathbf{x}_j} \triangleq \text{R}(\mathbf{x}_j, v_j)$, $vs_{\mathbf{x}_j} \triangleq \text{R}(\mathbf{x}_j, v_j)$ and $ws_{\mathbf{x}_j} \triangleq \text{W}(\mathbf{s}[\mathbf{x}_j], v_j)$.

- $RUs = ru_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} ru_{\mathbf{x}_n}$, where the identifiers of all constituent events of RUs are picked fresh, and for all $j \in \{1 \dots n\}$:

$$ru_{\mathbf{x}_j} = \begin{cases} \text{RU}(\mathbf{x}_j) & \text{if } \mathbf{x}_j \notin \text{WS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- $PLs = pl_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} pl_{\mathbf{x}_n}$, where the identifiers of all constituent events of PLs are picked fresh, and for all $j \in \{1 \dots n\}$:

$$pl_{\mathbf{x}_j} = \begin{cases} \text{PL}(\mathbf{x}_j) & \text{if } \mathbf{x}_j \in \text{WS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- $Ts = t_1 \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} t_k$, where for all $m \in \{1 \dots k\}$:

$$t_m = \begin{cases} \text{R}(\mathbf{s}[\mathbf{x}_n], v_n) & \text{if } t'_m = \text{R}(\mathbf{x}_n, v_n) \\ \text{W}(\mathbf{x}_n, v_n) \xrightarrow{\text{po|imm}} \text{W}(\mathbf{s}[\mathbf{x}_n], v_n) & \text{if } t'_m = \text{W}(\mathbf{x}_n, v_n) \end{cases}$$

such that in the first case the identifier of t_m is that of t'_m ; and in the second case the identifier of the first event in t_m is that of t'_m and the identifier of the second event is picked fresh.

- $Us = wu_{x_1} \xrightarrow{\text{po}|_{\text{imm}}} \dots \xrightarrow{\text{po}|_{\text{imm}}} wu_{x_n}$, where the identifiers of all constituent events of Us are picked fresh, and

$$wu_{x_j} = \begin{cases} \text{WU}(x_j) & \text{if } x_j \in \text{WS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

In what follows we write E_ξ as a shorthand for the events in the implementation trace of θ_ξ , i.e. $E_\xi \triangleq \{a \mid a \in \theta_\xi.E\}$. We use the ξ . prefix to project the various events of the implementation trace θ_ξ (e.g. $\xi.rl_{x_j}$).

For each location \mathbf{x} we then define:

$$\text{LO}_x \triangleq \left\{ \begin{array}{l} \{(\xi.rl_x, \xi.pl_x), (\xi.rl_x, \xi.wu_x), (\xi.pl_x, \xi.wu_x) \mid G'.\mathcal{T}_\xi \cap \mathcal{W}_x \neq \emptyset\} \\ \cup \left\{ \begin{array}{l} (\xi.rl_x, \xi'.pl_x), (\xi.rl_x, \xi'.wu_x), \\ (\xi.ru_x, \xi'.pl_x), (\xi.ru_x, \xi'.wu_x) \end{array} \middle| \begin{array}{l} \xi \neq \xi' \wedge \exists a, b, x. \\ a \in G'.\mathcal{T}_\xi \wedge b \in G'.\mathcal{T}_{\xi'} \\ \wedge \text{loc}(a) = \text{loc}(b) = \mathbf{x} \\ \wedge a \in G'.\mathcal{R}_E \wedge (a, b) \in G'.\text{rb} \end{array} \right\} \\ \cup \left\{ \begin{array}{l} (\xi.rl_x, \xi'.pl_x), (\xi.rl_x, \xi'.wu_x), \\ (\xi.pl_x, \xi'.rl_x), (\xi.pl_x, \xi'.pl_x), \\ (\xi.pl_x, \xi'.wu_x), \\ (\xi.wu_x, \xi'.rl_x), (\xi.wu_x, \xi'.pl_x), \\ (\xi.wu_x, \xi'.wu_x) \end{array} \middle| \begin{array}{l} \xi \neq \xi' \wedge \exists a, b, x. \\ a \in G'.\mathcal{T}_\xi \wedge b \in G'.\mathcal{T}_{\xi'} \\ \wedge \text{loc}(a) = \text{loc}(b) = \mathbf{x} \\ \wedge (a, b) \in G'.\text{mo} \end{array} \right\} \\ \cup \left\{ \begin{array}{l} (\xi.pl_x, \xi'.rl_x), (\xi.pl_x, \xi'.ru_x), \\ (\xi.wu_x, \xi'.rl_x), (\xi.wu_x, \xi'.ru_x) \end{array} \middle| \begin{array}{l} \xi \neq \xi' \wedge \exists a, b, x. \\ a \in G'.\mathcal{T}_\xi \wedge b \in G'.\mathcal{T}_{\xi'} \\ \wedge G'.\mathcal{T}_{\xi'} \cap \mathcal{W}_x = \emptyset \\ \wedge \text{loc}(a) = \text{loc}(b) = \mathbf{x} \\ \wedge (a, b) \in G'.(\text{mo}^?; \text{rf}) \end{array} \right\} \end{array} \right\}$$

Note that each LO_x satisfies the conditions in Def. 4.

We are now in a position to demonstrate the completeness of our implementation. Given an RSI-consistent execution graph G' , we construct an execution graph G of the implementation as follows and demonstrate that $\text{RA-consistent}(G)$ holds.

- $G.E = \bigcup_{\mathcal{T}_\xi \in G'.\mathcal{T}/\text{st}} \theta_\xi.E \cup G'.\mathcal{NT}$. Observe that $G'.E \subseteq G.E$.
- $G.\text{po}$ is defined as $G'.\text{po}$ extended by the po for the additional events of G , given by each θ_ξ trace defined above. Note that $G.\text{po}$ does not introduce additional orderings between events of $G'.E$. That is, $\forall a, b \in G'.E. (a, b) \in G'.\text{po} \Leftrightarrow (a, b) \in G.\text{po}$.
- $G.\text{rf} = \bigcup_{\mathbf{x} \in \text{Locs}} \left\{ \begin{array}{l} (w, \xi.rs_x), \\ (w, \xi.vs_x) \end{array} \middle| \begin{array}{l} \exists r. (\mathbf{x}, r) \in \text{RS}_\xi \\ \wedge (w, r) \in G'.\text{rf} \end{array} \right\} \cup (G'.\text{rf} \cap G'.E \times G'.\mathcal{NT})$.
- $G.\text{mo} = G'.\text{mo}$.
- $G.\text{lo} = \bigcup_{\mathbf{x} \in \text{Locs}} \text{LO}_x$, with LO_x as defined above.

Notation In what follows, given an RSI implementation graph G as constructed above we write $G.\mathcal{NT}$ for the non-transactional events of G , i.e. $G.\mathcal{NT} \triangleq \{a \mid \neg \exists \xi. a \in G.E_\xi\}$. Moreover, as before, given a relation $r \subseteq G.E \times G.E$, we override the r_I notation and write r_I for $\{(a, b) \in r \mid \exists \xi. a, b \in \theta_\xi.E\}$.

Lemma 16. *Given an RSI-consistent execution graph G' and its corresponding implementation graph G constructed as above, for all a, b, ξ_a, ξ_b :*

$$\begin{aligned}
& (a, b) \in G.\mathbf{hb} \Rightarrow \\
& \quad (\exists \xi. a, b \in G.E_\xi \Rightarrow (a, b) \in G.\mathbf{po}) \\
& \quad \wedge \left(\neg \exists \xi. a, b \in G.E_\xi \Rightarrow \right. \\
& \quad \quad \exists A, B. \emptyset \subset A \times B \subseteq G'.\mathbf{rsi-hb} \\
& \quad \quad \wedge (a \in G.\mathcal{NT} \Rightarrow A = \{a\}) \wedge (b \in G.\mathcal{NT} \Rightarrow B = \{b\}) \\
& \quad \quad \wedge [a \in G.E_{\xi_a} \Rightarrow \\
& \quad \quad \quad A = G'.\mathcal{T}_{\xi_a} \vee (\mathbf{stg}(a, \xi_a) \leq 2 \wedge A = G'.\mathcal{T}_{\xi_a} \cap \mathcal{R}_E) \\
& \quad \quad \quad \vee (\mathbf{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\mathbf{po}^?} d \wedge A = \{d\})] \\
& \quad \quad \wedge [b \in G.E_{\xi_b} \Rightarrow \\
& \quad \quad \quad B = G'.\mathcal{T}_{\xi_b} \vee (\mathbf{stg}(b, \xi_b) \geq 3 \wedge B = G'.\mathcal{T}_{\xi_b} \cap \mathcal{W})] \quad \left. \right)
\end{aligned}$$

where

$$\mathbf{stg}(a, \xi_a) \triangleq \begin{cases} 1 & \text{if } a \in \xi_a.Rs_{\xi_a} \\ 2 & \text{if } a \in \xi_a.RUs_{\xi_a} \\ 3 & \text{if } a \in \xi_a.Pls_{\xi_a} \\ 4 & \text{otherwise} \end{cases}$$

Proof. Pick an arbitrary RSI-consistent execution graph G' and its corresponding implementation graph G constructed as above. Let $\mathbf{hb}_0 \triangleq G.(\mathbf{po} \cup \mathbf{rf} \cup \mathbf{lo})$ and $\mathbf{hb}_{n+1} \triangleq \mathbf{hb}_0; \mathbf{hb}_n$, for all $n \in \mathbb{N}$. It is then straightforward to demonstrate that $G.\mathbf{hb} = \bigcup_{i \in \mathbb{N}} \mathbf{hb}_i$. We thus demonstrate instead that:

$$\begin{aligned}
& \forall n \in \mathbb{N}. \forall a, b, \xi_a, \xi_b. (a, b) \in G.\mathbf{hb}_n \Rightarrow \\
& \quad (\exists \xi. a, b \in G.E_\xi \Rightarrow (a, b) \in G.\mathbf{po}) \\
& \quad \wedge \left(\neg \exists \xi. a, b \in G.E_\xi \Rightarrow \right. \\
& \quad \quad \exists A, B. \emptyset \subset A \times B \subseteq G'.\mathbf{rsi-hb} \\
& \quad \quad \wedge (a \in G.\mathcal{NT} \Rightarrow A = \{a\}) \wedge (b \in G.\mathcal{NT} \Rightarrow B = \{b\}) \\
& \quad \quad \wedge [a \in G.E_{\xi_a} \Rightarrow \\
& \quad \quad \quad A = G'.\mathcal{T}_{\xi_a} \vee (\mathbf{stg}(a, \xi_a) \leq 2 \wedge A = G'.\mathcal{T}_{\xi_a} \cap \mathcal{R}_E) \\
& \quad \quad \quad \vee (\mathbf{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\mathbf{po}^?} d \wedge A = \{d\})] \\
& \quad \quad \wedge [b \in G.E_{\xi_b} \Rightarrow \\
& \quad \quad \quad B = G'.\mathcal{T}_{\xi_b} \vee (\mathbf{stg}(b, \xi_b) \geq 3 \wedge B = G'.\mathcal{T}_{\xi_b} \cap \mathcal{W})] \quad \left. \right)
\end{aligned}$$

We proceed by induction on n .

Base case $n = 0$

There are three cases to consider: 1) $(a, b) \in G.\text{po}$; or 2) $(a, b) \in G.\text{rf}$; or 3) $(a, b) \in G.\text{lo}$.

In case (1) there are five cases to consider: a) $\exists \xi. (a, b) \in G.E_\xi$; or b) $a, b \in G.\mathcal{NT}$; or c) $a \in G.\mathcal{NT}$ and $b \in G.E_{\xi_b}$; or d) $a \in G.E_{\xi_a}$ and $b \in G.\mathcal{NT}$; or e) $a \in G.E_{\xi_a}$, $b \in G.E_{\xi_b}$ and $\xi_a \neq \xi_b$. In case (1.a) we then have $(a, b) \in G.\text{po}_1$, as required. In case (1.b) from the construction of $G.\text{po}$ we have $(a, b) \in G'.\text{po}$, as required. In case (1.c) from the construction of $G.\text{po}$ we have $(\{a\} \times G'.\mathcal{T}_{\xi_b}) \in G'.\text{po}$, as required. In case (1.d) from the construction of $G.\text{po}$ we have $(G'.\mathcal{T}_{\xi_a} \times \{b\}) \in G'.\text{po}$, as required. In case (1.e) from the construction of $G.\text{po}$ we have $(G'.\mathcal{T}_{\xi_a} \times G'.\mathcal{T}_{\xi_b}) \in G'.\text{po}$, as required.

In case (2) there are five cases to consider: a) $\exists \xi. (a, b) \in G.E_\xi$; or b) $a, b \in G.\mathcal{NT}$; or c) $a \in G.\mathcal{NT}$ and $b \in G.E_{\xi_b}$; or d) $a \in G.E_{\xi_a}$ and $b \in G.\mathcal{NT}$; or e) $a \in G.E_{\xi_a}$, $b \in G.E_{\xi_b}$ and $\xi_a \neq \xi_b$. Case (2.a) holds vacuously as $(a, b) \in G.\text{rf}_1 = \emptyset$. In case (2.b) from the construction of $G.\text{rf}$ we have $(a, b) \in G'.\text{rf}$, as required. In case (2.c) from the construction of $G.\text{rf}$ we know that there exists $r \in G'.\mathcal{T}_{\xi_b}$ such that $(a, r) \in G'.\text{rf}$. As such we have $(\{w\} \times G'.\mathcal{T}_{\xi_b}) \subseteq G'.([\mathcal{NT}]; \text{rf}; \text{st}) \subseteq G'.\text{rsi-hb}$, as required. In case (2.d) from the construction of $G.\text{rf}$ we then have $(a, b) \in G'.\text{rf}$; that is a $\xrightarrow{G.\text{po}^?}$ a and $(a, b) \in G'.\text{rf}$, as required. In case (2.e) from the construction of $G.\text{rf}$ we know that there exists $r \in G'.\mathcal{T}_{\xi_b}$ such that $(a, r) \in G'.\text{rf}$. As such we have $(G'.\mathcal{T}_{\xi_a} \times G'.\mathcal{T}_{\xi_b}) \subseteq G'.\text{rf}_\top \subseteq G'.\text{rsi-hb}$, as required.

In case (3) there are five cases to consider: a) $\exists \xi. (a, b) \in G.E_\xi$; or b) $a, b \in G.\mathcal{NT}$; or c) $a \in G.\mathcal{NT}$ and $b \in G.E_{\xi_b}$; or d) $a \in G.E_{\xi_a}$ and $b \in G.\mathcal{NT}$; or e) $a \in G.E_{\xi_a}$, $b \in G.E_{\xi_b}$ and $\xi_a \neq \xi_b$. In case (3.a) from the construction of lo we have $(a, b) \in G.\text{po}$, as required. Cases (3.b-3.d) hold vacuously as there are no lo edge to or from non-transactional events. In case (3.e) from the construction of lo we know there exist x, c, d such that $c \in G'.\mathcal{T}_{\xi_a}$, $d \in G'.\mathcal{T}_{\xi_b}$, and either i) $(c, d) \in G'.\text{rf}$; or ii) $(c, d) \in G'.(\text{mo}; \text{rf})$; or iii) $(c, d) \in G'.\text{mo}$; or iv) $\text{loc}(c) = \text{loc}(d) = x$, $a = \xi_a.\text{rl}_x \vee a = \xi_a.\text{ru}_x$, $b = \xi_a.\text{pl}_x \vee b = \xi_a.\text{wu}_x$ and $c \in G'.\mathcal{R}_E$ and $(c, d) \in G'.\text{rb}$. In cases (3.e.i, 3.e.ii) we have $G'.\mathcal{T}_{\xi_a} \times G'.\mathcal{T}_{\xi_b} \subseteq G'.(\text{rf}_\top \cup (\text{mo}; \text{rf})_\top) \subseteq \text{rsi-rf} \subseteq \text{rsi-hb}$, as required. In case (3.e.iii) we have $G'.\mathcal{T}_{\xi_a} \times G'.\mathcal{T}_{\xi_b} \subseteq G'.\text{mo}_\top \subseteq \text{rsi-hb}$, as required. In case (3.e.iv) we then have $\text{stg}(a, \xi_a) \leq 2$, $\text{stg}(b, \xi_b) \geq 3$, and $G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E) \times G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}) \subseteq G'.\text{si-rb} \subseteq \text{rsi-hb}$, as required.

Inductive case $n = m+1$

$$\begin{aligned}
& \forall i \in \mathbb{N}. \forall a, b, \xi_a, \xi_b. i \leq m \wedge (a, b) \in G.\mathbf{hb}_i \Rightarrow \\
& \quad (\exists \xi. a, b \in G.E_\xi \Rightarrow (a, b) \in G.\mathbf{po}) \\
& \quad \wedge \left(\neg \exists \xi. a, b \in G.E_\xi \Rightarrow \right. \\
& \quad \quad \exists A, B. \emptyset \subset A \times B \subseteq G'.\mathbf{rsi-hb} \\
& \quad \quad \wedge (a \in G.\mathcal{NT} \Rightarrow A = \{a\}) \wedge (b \in G.\mathcal{NT} \Rightarrow B = \{b\}) \\
& \quad \quad \wedge [a \in G.E_{\xi_a} \Rightarrow \\
& \quad \quad \quad A = G'.\mathcal{T}_{\xi_a} \vee (\mathbf{stg}(a, \xi_a) \leq 2 \wedge A = G'.\mathcal{T}_{\xi_a} \cap \mathcal{R}_E) \\
& \quad \quad \quad \vee (\mathbf{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\mathbf{po}^?} d \wedge A = \{d\})] \\
& \quad \quad \wedge [b \in G.E_{\xi_b} \Rightarrow \\
& \quad \quad \quad B = G'.\mathcal{T}_{\xi_b} \vee (\mathbf{stg}(b, \xi_b) \geq 3 \wedge B = G'.\mathcal{T}_{\xi_b} \cap \mathcal{W})] \left. \right) \tag{I.H.}
\end{aligned}$$

Since $(a, b) \in \mathbf{hb}_n$, from the definition of \mathbf{hb}_n we know there exists c such that $(a, c) \in \mathbf{hb}_0$ and $(c, b) \in \mathbf{hb}_m$. There are then five cases to consider: 1) $\exists \xi. (a, b) \in G.E_\xi$; or 2) $a, b \in G.\mathcal{NT}$; or 3) $a \in G.\mathcal{NT}$ and $b \in G.E_{\xi_b}$; or 4) $a \in G.E_{\xi_a}$ and $b \in G.\mathcal{NT}$; or 5) $a \in G.E_{\xi_a}$, $b \in G.E_{\xi_b}$ and $\xi_a \neq \xi_b$.

Case 1

In case (1) pick arbitrary ξ such that $a, b \in G.E_\xi$. There are then three additional cases to consider: a) $c \in G.\mathcal{NT}$; b) $c \in G.E_\xi$; or c) there exists $\xi' \neq \xi$ such that $c \in G.E_{\xi'}$.

In case (1.a) from the proof of the base case and the (I.H.) we know there exist $A, B, C \neq \emptyset$ such that $C = \{c\}$, $A \times C \subseteq G'.\mathbf{rsi-hb}$, $C \times B \subseteq G'.\mathbf{rsi-hb}$ and thus $A \times B \subseteq G'.\mathbf{rsi-hb}$ and either: i) $A = G'.\mathcal{T}_\xi$ and $B = G'.\mathcal{T}_\xi$; or ii) $A = G'.\mathcal{T}_\xi$ and $B = (G'.(\mathcal{T}_\xi \cap \mathcal{W}))$; or iii) $\mathbf{stg}(a, \xi) \leq 2$, $A = G'.(\mathcal{T}_\xi \cap \mathcal{R}_E)$ and $B = G'.\mathcal{T}_\xi$; or iv) $\mathbf{stg}(a, \xi) \leq 2$, $A = G'.(\mathcal{T}_\xi \cap \mathcal{R}_E)$, $\mathbf{stg}(b, \xi) \geq 3$ and $B = G'.(\mathcal{T}_\xi \cap \mathcal{W})$; or v) $\mathbf{stg}(a, \xi) \leq 4$, $\exists d \in \xi.\mathcal{W}. a \xrightarrow{G.\mathbf{po}^?} d \wedge A = \{d\}$ and $B = G'.\mathcal{T}_\xi$; or vi) $\mathbf{stg}(a, \xi) \leq 4$, $\exists d \in \xi.\mathcal{W}. a \xrightarrow{G.\mathbf{po}^?} d \wedge A = \{d\}$ and $B = G'.(\mathcal{T}_\xi \cap \mathcal{W})$.

Case (i) cannot arise as we would have $A = B$ and thus $A \times A \subseteq G'.\mathbf{rsi-hb}$, contradicting the assumption that G' is RSI-consistent. Case (ii) cannot arise as we would have $(G'.(\mathcal{T}_\xi \cap \mathcal{W})) \times (G'.(\mathcal{T}_\xi \cap \mathcal{W})) \subseteq G'.\mathbf{rsi-hb}$, contradicting the assumption that G' is RSI-consistent. Case (iii) cannot arise as we would have $(G'.(\mathcal{T}_\xi \cap \mathcal{R}_E)) \times (G'.(\mathcal{T}_\xi \cap \mathcal{R}_E)) \subseteq G'.\mathbf{rsi-hb}$, contradicting the assumption that G' is RSI-consistent. In case (iv) since we have $\mathbf{stg}(a, \xi) \leq 2$ and $\mathbf{stg}(b, \xi) \geq 3$, from the definition of $\mathbf{stg}(\cdot, \cdot)$ and the construction of G we have $(a, b) \in G.\mathbf{po}$, as required. Cases (v-vi) cannot arise as we would have $\exists d \in \xi.\mathcal{W}. (d, d) \in G'.\mathbf{rsi-hb}$, contradicting the assumption that G' is RSI-consistent.

In case (1.b) from the proof of the base case we have $(a, c) \in G.\mathbf{po}$. On the other hand from (I.H.) we have $(c, b) \in G.\mathbf{po}$. As $G.\mathbf{po}$ is transitively closed, we have $(a, b) \in G.\mathbf{po}$, as required.

In case (1.c) from the proof of the base case (in cases 1.e, 2.e and 3.e) we have either A) $G'.\mathcal{T}_\xi \times G'.\mathcal{T}_{\xi'} \subseteq G'.\text{rsi-hb}$; or B) $\text{stg}(a, \xi) \leq 2$, $\text{stg}(c, \xi') \geq 3$, and $\emptyset \subset G'.(\mathcal{T}_\xi \cap \mathcal{R}_E) \times G'.(\mathcal{T}_{\xi'} \cap \mathcal{W}) \subseteq G'.\text{rsi-hb}$. On the other hand from (I.H.) we know there exist $C, B \neq \emptyset$ such that $C \times B \subseteq G'.\text{rsi-hb}$ and either: i) $C = G'.\mathcal{T}_{\xi'}$ and $B = G'.\mathcal{T}_\xi$; or ii) $C = G'.\mathcal{T}_{\xi'}$, $\text{stg}(b, \xi) \geq 3$ and $B = G'.(\mathcal{T}_\xi \cap \mathcal{W})$; or iii) $\text{stg}(c, \xi') \leq 2$, $C = G'.(\mathcal{T}_{\xi'} \cap \mathcal{R}_E)$ and $B = G'.\mathcal{T}_\xi$; or iv) $\text{stg}(c, \xi') \leq 2$, $C = G'.(\mathcal{T}_{\xi'} \cap \mathcal{R}_E)$, $\text{stg}(b, \xi) \geq 3$ and $B = G'.(\mathcal{T}_\xi \cap \mathcal{W})$; or v) $\text{stg}(c, \xi') \leq 4$, $\exists d \in \xi'.\mathcal{W}$. $c \xrightarrow{G.\text{po}^?} d \wedge C = \{d\}$ and $B = G'.\mathcal{T}_\xi$; or vi) $\text{stg}(c, \xi') \leq 4$, $\exists d \in \xi'.\mathcal{W}$. $c \xrightarrow{G.\text{po}^?} d \wedge C = \{d\}$, $\text{stg}(b, \xi) \geq 3$ and $B = G'.(\mathcal{T}_\xi \cap \mathcal{W})$.

Cases (A.i-A.vi) lead to a cycle in $G'.\text{rsi-hb}$, contradicting the assumption that G' is RSI-consistent. In cases (B.ii, B.iv, B.vi) we then have $\text{stg}(a, \xi) \leq 2$ and $\text{stg}(b, \xi) \geq 3$. Consequently from the definition of $\text{stg}(\cdot, \cdot)$ and the construction of G we have $(a, b) \in G.\text{po}$, as required. Case (B.i) leads to a cycle in $G'.\text{rsi-hb}$, contradicting the assumption that G' is RSI-consistent. In cases (B.iii) we have $\text{stg}(c, \xi') \geq 3$ and $\text{stg}(c, \xi) \leq 2$, leading to a contradiction. In case (B.v) we then know $\exists d \in \xi'.\mathcal{W}$. $\emptyset \subset G'.(\mathcal{T}_\xi \cap \mathcal{R}_E) \times \{d\} \subseteq G'.\text{rsi-hb} \wedge \{d\} \times G'.\mathcal{T}_\xi \subseteq G'.\text{rsi-hb}$. That is, we have $\emptyset \subset G'.(\mathcal{T}_\xi \cap \mathcal{R}_E) \times G'.(\mathcal{T}_\xi \cap \mathcal{R}_E) \subseteq G'.\text{rsi-hb}$, contradicting the assumption that G' is RSI-consistent.

Case 2

There are two additional cases to consider: a) $c \in G.\mathcal{NT}$; or b) there exists ξ such that $c \in G.E_\xi$.

In case (2.a) from the proof of the base case we have $(a, c) \in G'.\text{rsi-hb}$. On the other hand from (I.H.) we have $(c, b) \in G'.\text{rsi-hb}$. As $G'.\text{rsi-hb}$ is transitively closed, we have $(a, b) \in G'.\text{rsi-hb}$, as required.

In case (2.b) from the proof of the base case we know there exists $C_1 \neq \emptyset$ such that $\{a\} \times C_1 \subseteq G'.\text{rsi-hb}$ and either: A) $C_1 = G'.\mathcal{T}_\xi$; or B) $\text{stg}(c, \xi) \geq 3$ and $C_1 = G'.(\mathcal{T}_\xi \cap \mathcal{W})$. On the other hand, from (I.H.) we know there exists $C_2 \neq \emptyset$ such that $C_2 \times \{b\} \in G'.\text{rsi-hb}$ and either: i) $C_2 = G'.\mathcal{T}_\xi$; or ii) $\text{stg}(c, \xi) \leq 2$ and $C_2 = G'.(\mathcal{T}_\xi \cap \mathcal{R}_E)$; or iii) $\text{stg}(c, \xi) \leq 4$ and $\exists d \in \xi.\mathcal{W}$. $c \xrightarrow{G.\text{po}^?} d \wedge C_2 = \{d\}$.

In cases (A.i-A.iii) and (B.i, B.iii) from the transitivity of $G'.\text{rsi-hb}$ we have $(a, b) \in G'.\text{rsi-hb}$, as required. Case (B.ii) cannot arise as otherwise we would have $3 \leq \text{stg}(c, \xi) \leq 2$, leading to a contradiction.

Case 3

There are three additional cases to consider: a) $c \in G.\mathcal{NT}$; b) $c \in G.E_{\xi_b}$; or c) there exists $\xi_c \neq \xi_b$ such that $c \in G.E_{\xi_c}$.

In case (3.a) from the proof of the base case we have $(a, c) \in G'.\text{rsi-hb}$. On the other hand from (I.H.) we know there exists $B \neq \emptyset$ such that $\{c\} \times B \subseteq G'.\text{rsi-hb}$ and $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$. As $G'.\text{rsi-hb}$ is transitively closed we then know there exists $B \neq \emptyset$ such that $\{a\} \times B \subseteq G'.\text{rsi-hb}$ and $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$, as required.

In case (3.b) from the proof of the base case we know there exists $B \neq \emptyset$ such that $\{a\} \times B \subseteq G'.\text{rsi-hb}$ and $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(c, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$. On

the other hand, from (I.H.) we have $c \xrightarrow{G.\text{po}} b$ and thus from the definition of $\text{stg}(\cdot, \cdot)$ and the construction of G we have $\text{stg}(b, \xi_b) \geq \text{stg}(c, \xi_b)$. As such, we know there exists $B \neq \emptyset$ such that $\{a\} \times B \subseteq G'.\text{rsi-hb}$ and $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$, as required.

In case (3.c) from the proof of the base case we know there exists $C_1 \neq \emptyset$ such that $\{a\} \times C_1 \subseteq G'.\text{rsi-hb}$ and either A) $C_1 = G'.\mathcal{T}_{\xi_c}$; or B) $(\text{stg}(c, \xi_c) \geq 3 \wedge C_1 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{W}))$. On the other hand, from (I.H.) we know there exist $C_2, B \neq \emptyset$ such that $C_2 \times B \in G'.\text{rsi-hb}$ and either: i) $C_2 = G'.\mathcal{T}_{\xi_c}$ and $B = G'.\mathcal{T}_{\xi_b}$; or ii) $C_2 = G'.\mathcal{T}_{\xi_c}$ and $(\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$; or iii) $\text{stg}(c, \xi_c) \leq 2 \wedge C_2 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{R}_E)$ and $B = G'.\mathcal{T}_{\xi_b}$; or iv) $\text{stg}(c, \xi_c) \leq 2 \wedge C_2 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{R}_E)$ and $(\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$; or v) $\text{stg}(c, \xi_c) \leq 4 \wedge \exists d \in \xi_c.\mathcal{W}. c \xrightarrow{G.\text{po}^?} d \wedge C_2 = \{d\}$ and $B = G'.\mathcal{T}_{\xi_b}$; or vi) $\text{stg}(c, \xi_c) \leq 4 \wedge \exists d \in \xi_c.\mathcal{W}. c \xrightarrow{G.\text{po}^?} d \wedge C_2 = \{d\}$ and $(\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$.

In cases (A.i-A.vi) and (B.i, B.ii, B.v, B.vi) from the transitivity of $G'.\text{rsi-hb}$ we know there exists $B \neq \emptyset$ such that $\{a\} \times B \subseteq G'.\text{rsi-hb}$ and $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$, as required. Cases (B.iii, B.iv) cannot arise as we would otherwise have $3 \leq \text{stg}(c, \xi_c) \leq 2$, leading to a contradiction.

Case 4

There are three additional cases to consider: a) $c \in G.\mathcal{NT}$; b) $c \in G.E_{\xi_a}$; or c) there exists $\xi_c \neq \xi_a$ such that $c \in G.E_{\xi_c}$.

In case (4.a) from (I.H.) we have $(c, b) \in G'.\text{rsi-hb}$. On the other hand from the base case we know there exists $A \neq \emptyset$ such that $A \times \{c\} \subseteq G'.\text{rsi-hb}$ and $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee \text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{d\}$. As $G'.\text{rsi-hb}$ is transitively closed we then know there exists $A \neq \emptyset$ such that $A \times \{b\} \subseteq G'.\text{rsi-hb}$ and $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee \text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{d\}$, as required.

In case (4.b) from (I.H.) we know there exists $A \neq \emptyset$ such that $A \times \{b\} \subseteq G'.\text{rsi-hb}$ and $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(c, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee \text{stg}(c, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{d\}$. On the other hand, from the proof of the base case we have $a \xrightarrow{G.\text{po}} c$ and thus from the definition of $\text{stg}(\cdot, \cdot)$ and the construction of G we have $\text{stg}(a, \xi_a) \leq \text{stg}(c, \xi_a)$. As such, we know there exists $A \neq \emptyset$ such that $A \times \{b\} \subseteq G'.\text{rsi-hb}$ and $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee \text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{d\}$, as required.

In case (4.c) from the proof of the base case we know there exist $A, C_1 \neq \emptyset$ such that $A \times C_1 \subseteq G'.\text{rsi-hb}$ and either: i) $A = G'.\mathcal{T}_{\xi_a}$ and $C_1 = G'.\mathcal{T}_{\xi_c}$; or ii) $A = G'.\mathcal{T}_{\xi_a}$ and $(\text{stg}(c, \xi_c) \geq 3 \wedge C_1 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{W}))$; or iii) $\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)$ and $C_1 = G'.\mathcal{T}_{\xi_c}$; or iv) $\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)$ and $(\text{stg}(c, \xi_c) \geq 3 \wedge C_1 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{W}))$; or v) $\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{d\}$ and $C_1 = G'.\mathcal{T}_{\xi_c}$; or vi) $\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{d\}$ and $(\text{stg}(c, \xi_c) \geq 3 \wedge C_1 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{W}))$.

On the other hand, from (I.H.) we know there exists $C_2 \neq \emptyset$ such $C_2 \times \{b\} \subseteq G'.\text{rsi-hb}$ and either A) $C_2 = G'.\mathcal{T}_{\xi_c}$; or B) $(\text{stg}(c, \xi_c) \leq 2 \wedge C_2 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{R}_E))$; or C) $(\text{stg}(c, \xi_c) \leq 4 \wedge \exists e \in \xi_c.\mathcal{W}. c \xrightarrow{G.\text{po}^?} e \wedge C_2 = \{e\})$.

In cases (A.i-A.vi), (C.i-C.vi) and (B.i, B.iii, B.v) from the transitivity of $G'.\text{rsi-hb}$ we know there exists $A \neq \emptyset$ such that $A \times \{b\} \subseteq G'.\text{rsi-hb}$ and $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{d\})$, as required. Cases (B.ii, B.iv, B.vi) cannot arise as we would otherwise have $3 \leq \text{stg}(c, \xi_c) \leq 2$, leading to a contradiction.

Case 5

There are four additional cases to consider: a) $c \in G.\mathcal{NT}$; b) $c \in G.E_{\xi_a}$; or c) $c \in G.E_{\xi_b}$; or d) there exists ξ_c such that $\xi_c \neq \xi_a$, $\xi_c \neq \xi_b$ and $c \in G.E_{\xi_c}$.

In case (5.a) from the proof of the base case we know there exists $A \neq \emptyset$ such that $A \times \{c\} \subseteq G'.\text{rsi-hb}$ and $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{d\})$. On the other hand from (I.H.) we know there exists $B \neq \emptyset$ such that $\{c\} \times B \subseteq G'.\text{rsi-hb}$ and $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$. As such, since $G'.\text{rsi-hb}$ is transitive we know there exist $A, B \neq \emptyset$ such that $A \times B \subseteq G'.\text{rsi-hb}$; that $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{d\})$; and that $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$, as required.

In case (5.b) from (I.H.) we know there exist $A, B \neq \emptyset$ such that $A \times B \subseteq G'.\text{rsi-hb}$; that $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(c, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee (\text{stg}(c, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{d\})$; and that $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$. On the other hand from the proof of the base case we have $(a, c) \in G'.\text{po}$ and thus from the definition of $\text{stg}(\cdot, \cdot)$ and the construction of G we have $\text{stg}(a, \xi_a) \leq \text{stg}(c, \xi_a)$. As such we know there exist $A, B \neq \emptyset$ such that $A \times B \subseteq G'.\text{rsi-hb}$; that $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{d\})$ and that $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$, as required.

In case (5.c) from the proof of the base case we know there exist $A, B \neq \emptyset$ such that $A \times B \subseteq G'.\text{rsi-hb}$; that $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{d\})$, and that $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(c, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$. On the other hand from (I.H.) we have $(c, b) \in G'.\text{po}$ and thus from the definition of $\text{stg}(\cdot, \cdot)$ and the construction of G we have $\text{stg}(c, \xi_a) \leq \text{stg}(b, \xi_a)$. As such we know there exist $A, B \neq \emptyset$ such that $A \times B \subseteq G'.\text{rsi-hb}$; that $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{d\})$ and that $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$, as required.

In case (5.d) from the proof of the base case we know there exist $A, C_1 \neq \emptyset$ such that $A \times C_1 \subseteq G'.\text{rsi-hb}$; that $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap$

$\mathcal{R}_E)) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^2} d \wedge A = \{d\})$, and that either A) $C_1 = G'.\mathcal{T}_{\xi_c}$; or B) $(\text{stg}(c, \xi_c) \geq 3 \wedge C_1 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{W}))$.

On the other hand, from (I.H.) we know there exist $C_2, B \neq \emptyset$ such that $C_2 \times B \subseteq G'.\text{rsi-hb}$; that $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$; and that either: i) $C_2 = G'.\mathcal{T}_{\xi_c}$; or ii) $\text{stg}(c, \xi_c) \leq 2 \wedge C_2 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{R}_E)$; iii) $\text{stg}(c, \xi_c) \leq 4 \wedge \exists d \in \xi_c.\mathcal{W}. c \xrightarrow{G.\text{po}^2} d \wedge C_2 = \{d\}$.

In cases (A.i-A.iii) and (B.i, B.iii) from the transitivity of $G'.\text{rsi-hb}$ we know there exists $A, B \neq \emptyset$ such that $A \times B \subseteq G'.\text{rsi-hb}$; that $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^2} d \wedge A = \{d\})$; and that $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$, as required. Case (B.ii) cannot arise as we would otherwise have $3 \leq \text{stg}(c, \xi_c) \leq 2$, leading to a contradiction.

Theorem 12 (Completeness). For all RSI execution graphs G' and their counterpart implementation graphs G constructed as above,

$$\text{rsi-consistent}(G') \Rightarrow \text{RA-consistent}(G)$$

Proof. Pick an arbitrary RSI execution graph G' and its counterpart implementation graph G constructed as above and let us assume that $\text{rsi-consistent}(G')$ holds. From the definition of $\text{RA-consistent}(G)$ it then suffices to show:

1. $\text{irreflexive}(G.\text{hb})$
2. $\text{irreflexive}(G.\text{mo}; G.\text{hb})$
3. $\text{irreflexive}(G.\text{rb}; G.\text{hb})$

RTS. part 1

We proceed by contradiction. Let us assume that there exists a such that $(a, a) \in G.\text{hb}$. There are now two cases to consider: 1) $a \in G.\mathcal{NT}$; or 2) $\exists \xi. a \in G.E_\xi$. In case (1) from Lemma 16 we have $(a, a) \in G'.\text{rsi-hb}$, contradicting the assumption that G' is RSI-consistent. Similarly, in case (2) from Lemma 16 we have $(a, a) \in G.\text{po}$, leading to a contradiction as $G.\text{po}$ is acyclic by construction.

RTS. part 2

We proceed by contradiction. Let us assume that there exist a, b such that $(a, b) \in G.\text{hb}$ and $(b, a) \in G.\text{mo}$. From the construction of $G.\text{mo}$ we then know that $(b, a) \in G'.\text{mo}$. Let $\text{loc}(a) = \text{loc}(b) = \mathbf{x}$. There are then five cases to consider: 1) $\exists \xi. (a, b) \in G.E_\xi$; or 2) $a, b \in G.\mathcal{NT}$; or 3) $a \in G.\mathcal{NT}$ and $b \in G.E_{\xi_b}$; or 4) $a \in G.E_{\xi_a}$ and $b \in G.\mathcal{NT}$; or 5) $a \in G.E_{\xi_a}, b \in G.E_{\xi_b}$ and $\xi_a \neq \xi_b$.

In case (1) from Lemma 16 we have $(a, b) \in G.\text{po}$, and since $a, b \in G.\mathcal{W} = G'.\mathcal{W}$, from the construction of $G.\text{po}$ we have $(a, b) \in G'.\text{po}$. Moreover, since a, b are write events in the same transaction ξ , $(a, b) \in G'.\text{po}_I \cap \mathcal{W}^2 \subseteq G'.\text{rsi-hb}$. As such we have $a \xrightarrow{G'.\text{rsi-hb}} b \xrightarrow{G'.\text{mo}} a$, contradicting the assumption that G' is RSI-consistent.

In cases (2, 3) since $a, b \in G.\mathcal{W} = G'.\mathcal{W}$, from [Lemma 16](#) we have $(a, b) \in G'.\text{rsi-hb}$. As such we have $a \xrightarrow{G'.\text{rsi-hb}} b \xrightarrow{G'.\text{mo}} a$, contradicting the assumption that G' is RSI-consistent.

Similarly, in cases (4, 5) since $a, b \in G.\mathcal{W} = G'.\mathcal{W}$, from [Lemma 16](#) we have either i) $(a, b) \in G'.\text{rsi-hb}$; or ii) $\exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge (d, b) \in G'.\text{rsi-hb}$. On the other hand, since in case (ii) $a, d \in G.E_{\xi_a}$, $a \xrightarrow{G.\text{po}^?} d$ and a, d are both write events, we also have $a \xrightarrow{G'.\text{po}^?} d$. Moreover, since $G'.\text{po}_I \cap \mathcal{W}^2 \subseteq G'.\text{rsi-hb}$, we have $(a, d) \in G'.\text{rsi-hb}^?$. As such we have $a \xrightarrow{G'.\text{rsi-hb}^?} d \xrightarrow{G'.\text{rsi-hb}} b$ and from the transitivity of $G'.\text{rsi-hb}$ in both cases we have $(a, b) \in G'.\text{rsi-hb}$. Consequently, we have $a \xrightarrow{G'.\text{rsi-hb}} b \xrightarrow{G'.\text{mo}} a$, contradicting the assumption that G' is RSI-consistent.

RTS. part 3

We proceed by contradiction. Let us assume that there exists a, b such that $(a, b) \in G.\text{hb}$ and $(b, a) \in G.\text{rb}$. Let $\text{loc}(a) = \text{loc}(b) = \mathbf{x}$. There are then five cases to consider: 1) $\exists \xi. (a, b) \in G.E_\xi$; or 2) $a, b \in G.\mathcal{NT}$; or 3) $a \in G.\mathcal{NT}$ and $b \in G.E_{\xi_b}$; or 4) $a \in G.E_{\xi_a}$ and $b \in G.\mathcal{NT}$; or 5) $a \in G.E_{\xi_a}$, $b \in G.E_{\xi_b}$ and $\xi_a \neq \xi_b$.

In case (1) from [Lemma 16](#) we have $(a, b) \in G.\text{po}$. On the other hand since $(b, a) \in G.\text{rb}$, from the construction of G we know that $(b, a) \in G.\text{po}$. As such, since $G.\text{po}$ is transitively closed and $a \xrightarrow{G.\text{po}} b \xrightarrow{G.\text{po}} a$, we have $(a, a) \in G.\text{po}$, leading to a contradiction as G is acyclic by construction.

In case (2) from the construction of $G.\text{rb}$ we know that $(a, b) \in G'.\text{rb}$. On the other hand, from [Lemma 16](#) we have $(a, b) \in G'.\text{rsi-hb}$. As such we have $a \xrightarrow{G'.\text{rsi-hb}} b \xrightarrow{G'.\text{rb}} a$, contradicting the assumption that G' is RSI-consistent.

In case (3) from the construction of $G.\text{rb}$ we know that $b = \xi_b.\text{rs}_x \vee b = \xi_b.\text{vs}_x$. Consequently, from [Lemma 16](#) we have $\{a\} \times G'.\mathcal{T}_{\xi_b} \subseteq G'.\text{rsi-hb}$. On the other hand, from the construction of $G.\text{rb}$ we know there exists $r \in G'.\mathcal{T}_{\xi_b}$ such that $(r, a) \in G'.\text{rb}$. As such, we have $a \xrightarrow{G'.\text{rsi-hb}} r \xrightarrow{G'.\text{rb}} a$, contradicting the assumption that G' is RSI-consistent.

In case (4) from the construction of $G.\text{rb}$ we know that $(a, b) \in G'.\text{rb}$. Moreover, since $a \in G.\mathcal{W}$, from [Lemma 16](#) we have either i) $(a, b) \in G'.\text{rsi-hb}$; or ii) $\exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge (d, b) \in G'.\text{rsi-hb}$. As such, since a, d are both write events and $a \xrightarrow{G.\text{po}^?} d$, we also have $a \xrightarrow{G'.\text{po}^?} d$. Moreover, since $G'.\text{po}_I \cap \mathcal{W}^2 \subseteq G'.\text{rsi-hb}$, we have $(a, d) \in G'.\text{rsi-hb}^?$. As such we have $a \xrightarrow{G'.\text{rsi-hb}^?} d \xrightarrow{G'.\text{rsi-hb}} b$ and from the transitivity of $G'.\text{rsi-hb}$ in both cases (i, ii) we have $(a, b) \in G'.\text{rsi-hb}$. Consequently, we have $a \xrightarrow{G'.\text{rsi-hb}} b \xrightarrow{G'.\text{rb}} a$, contradicting the assumption that G' is RSI-consistent.

In case (5) from the construction of $G.\text{rb}$ we know that $b = \xi_b.\text{rs}_x \vee b = \xi_b.\text{vs}_x$. Consequently, from [Lemma 16](#) and since a is a write event, we have either i)

$G'.\mathcal{T}_{\xi_a} \times G'.\mathcal{T}_{\xi_b} \subseteq G'.\text{rsi-hb}$; or ii) $\exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge \{d\} \times G'.\mathcal{T}_{\xi_b} \subseteq G'.\text{rsi-hb}$.
 As such, since a, d are both write events and $a \xrightarrow{G.\text{po}^?} d$, we also have $a \xrightarrow{G'.\text{po}^?} d$.
 Moreover, since $G'.\text{po}_I \cap \mathcal{W}^2 \subseteq G'.\text{rsi-hb}$, we have $(a, d) \in G'.\text{rsi-hb}^?$. As such we
 have $a \xrightarrow{G'.\text{rsi-hb}^?} d \xrightarrow{G'.\text{rsi-hb}} b$ and from the transitivity of $G'.\text{rsi-hb}$ in both cases (i,
 ii) we have $\{a\} \times G'.\mathcal{T}_{\xi_b} \subseteq G'.\text{rsi-hb}$. On the other hand, from the construction of
 $G'.\text{rb}$ we know there exists $r \in G'.\mathcal{T}_{\xi_b}$ such that $(r, a) \in G'.\text{rb}$. Consequently, we
 have $a \xrightarrow{G'.\text{rsi-hb}} b \xrightarrow{G'.\text{rb}} a$, contradicting the assumption that G' is RSI-consistent.

F Soundness and Completeness of the Lazy RSI Implementation

Given an execution graph G of the lazy RSI implementation, let us assign a transaction identifier to each transaction executed by the program; and given a transaction ξ . Let $RS_\xi^0 = WS_\xi^0 = \emptyset$. Let us write $G.\mathcal{NT}$ for those events in $G.E$ that do not occur in a transaction. Observe that given a transaction ξ of the lazy RSI implementation, the trace of ξ , written θ_ξ , is of the form:

$$Fs \xrightarrow{\text{po|imm}} Is \xrightarrow{\text{po|imm}} Ts \xrightarrow{\text{po|imm}} VRs \xrightarrow{\text{po|imm}} RUs \xrightarrow{\text{po|imm}} PLs \xrightarrow{\text{po|imm}} Ws \xrightarrow{\text{po|imm}} Us$$

where:

- Fs denotes the sequence of events failing to obtain the necessary locks (i.e. those iterations that do not succeed in promoting the writer locks) or validate the snapshot;
- Is denotes the sequence of events initialising the values of **LS**, **RS** and **WS** with \emptyset , initialising **ws** with $[\]$ and initialising $\mathbf{s}[\mathbf{x}]$ with (\perp, \perp) for each location \mathbf{x} ;
- Ts denotes the sequence of events corresponding to the execution of (\mathbb{T}) and is of the form $t_1 \xrightarrow{\text{po|imm}} \dots \xrightarrow{\text{po|imm}} t_k$, where for all $m \in \{1 \dots k\}$:

$$t_m = \begin{cases} rd(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \xrightarrow{\text{po|imm}} lr_{\mathbf{x}_m} & \text{if } O_m = \mathbf{R}(\mathbf{x}_m, v_m) \\ wr(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \xrightarrow{\text{po|imm}} wws_{\mathbf{x}_m} & \text{if } O_m = \mathbf{W}(\mathbf{x}_m, v_m) \\ \xrightarrow{\text{po|imm}} \mathbf{R}(\mathbf{s}[\mathbf{x}_m], (v'_m, -)) \xrightarrow{\text{po|imm}} lw_{\mathbf{x}_m} \xrightarrow{\text{po|imm}} \mathbf{W}(\mathbf{ws}, ws_m) & \end{cases}$$

where O_m denotes the m th event in the trace of \mathbb{T} ; $lr_{\mathbf{x}_m} \triangleq \mathbf{R}(\mathbf{s}[\mathbf{x}_m], (-, v_m))$;

$$rd(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \triangleq \begin{cases} \mathbf{R}(\mathbf{s}[\mathbf{x}_m], (\perp, \perp)) & \text{if } \mathbf{x}_m \notin RS_{m-1} \cup WS_{m-1} \\ \xrightarrow{\text{po|imm}} fs_m & \\ \xrightarrow{\text{po|imm}} rl_{\mathbf{x}_m} & \\ \xrightarrow{\text{po|imm}} wrs_{\mathbf{x}_m} & \\ \xrightarrow{\text{po|imm}} rs_{\mathbf{x}_m} & \\ \xrightarrow{\text{po|imm}} ws_{\mathbf{x}_m} & \\ \emptyset & \text{otherwise} \end{cases}$$

fs_m denotes the sequence of events attempting (but failing) to acquire the read lock on \mathbf{x}_m , $rl_{\mathbf{x}_m} \triangleq \mathbf{RL}(\mathbf{x}_m)$, $wrs_{\mathbf{x}_m} \triangleq \mathbf{W}(\mathbf{RS}, RS_m)$, $rs_{\mathbf{x}_m} \triangleq \mathbf{R}(\mathbf{x}_m, v_{\mathbf{x}_m}^0)$, $ws_{\mathbf{x}_m} \triangleq \mathbf{W}(\mathbf{s}[\mathbf{x}_m], (v_{\mathbf{x}_m}^0, v_{\mathbf{x}_m}^0))$; and for all $m > 0$:

$$RS_{m+1} \triangleq \begin{cases} RS_m \cup \{\mathbf{x}_m\} & \text{if } O_m = \mathbf{R}(\mathbf{x}_m, -) \\ RS_m & \text{otherwise} \end{cases}$$

and

$$wr(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \triangleq \begin{cases} R(\mathbf{s}[\mathbf{x}_m], (\perp, \perp)) & \text{if } \mathbf{x}_m \notin RS_{m-1} \cup WS_{m-1} \\ \xrightarrow{\text{po|imm}} fs_m & \\ \xrightarrow{\text{po|imm}} rl_{\mathbf{x}_m} \emptyset & \text{otherwise} \end{cases}$$

$wws_{\mathbf{x}_m} \triangleq W(WS, WS_m)$; $lw_{\mathbf{x}_m} \triangleq W(\mathbf{s}[\mathbf{x}_m], (v'_m, v_m))$; fs_m and $rl_{\mathbf{x}_m}$ are as defined above; and for all $m > 0$:

$$WS_{m+1} \triangleq \begin{cases} WS_m \cup \{\mathbf{x}_m\} & \text{if } O_m = W(\mathbf{x}_m, -) \\ WS_m & \text{otherwise} \end{cases}$$

$$ws_{m+1} \triangleq \begin{cases} ws_m ++ [(\mathbf{x}_m, v_m)] & \text{if } O_m = W(\mathbf{x}_m, v_m) \\ ws_m & \text{otherwise} \end{cases}$$

Let $RS_\xi = RS_m$, $WS_\xi = WS_m$, and $ws_\xi = ws_m$; let $RS_\xi \cup WS_\xi$ be enumerated as $\{\mathbf{x}_1 \cdots \mathbf{x}_i\}$ for some i , and ws_ξ be enumerated as $\{(\mathbf{x}_1, v_1) \cdots (\mathbf{x}_j, v_j)\}$ for some j .

- VRs denotes the sequence of events validating the reads and is of the form $v_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \cdots \xrightarrow{\text{po|imm}} v_{\mathbf{x}_i}$, where for all $n \in \{1 \cdots i\}$:

$$v_{\mathbf{x}_n} = \begin{cases} R(\mathbf{s}[\mathbf{x}_n], (v_{\mathbf{x}_n}^0, -)) \xrightarrow{\text{po|imm}} vr_{\mathbf{x}_n} = R(\mathbf{x}_n, v_{\mathbf{x}_n}^0) & \text{if } \mathbf{x}_n \in RS_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- RUs denotes the sequence of events releasing the reader locks (when the given location is in the read set only) and is of the form $ru_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \cdots \xrightarrow{\text{po|imm}} ru_{\mathbf{x}_i}$, where for all $n \in \{1 \cdots i\}$:

$$ru_{\mathbf{x}_n} = \begin{cases} RU(\mathbf{x}_n) & \text{if } \mathbf{x}_n \notin WS_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- PLs denotes the sequence of events promoting the reader locks to writer ones (when the given location is in the write set), and is of the form $pl_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \cdots \xrightarrow{\text{po|imm}} pl_{\mathbf{x}_i}$, where for all $n \in \{1 \cdots i\}$:

$$pl_{\mathbf{x}_n} = \begin{cases} PL(\mathbf{x}_n) & \text{if } \mathbf{x}_n \in WS_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- Ws denotes the sequence of events committing the writes of (\mathbb{T}) and is of the form $c_{\mathbf{x}_1, v_1} \xrightarrow{\text{po|imm}} \cdots \xrightarrow{\text{po|imm}} c_{\mathbf{x}_j, v_j}$, where for all $n \in \{1 \cdots j\}$: $c_{\mathbf{x}_n, v_n} = W(\mathbf{x}_n, v_n)$

- Us denotes the sequence of events releasing the locks on the write set, and is of the form $wu_{\mathbf{x}_1} \xrightarrow{\text{po|imm}} \cdots \xrightarrow{\text{po|imm}} wu_{\mathbf{x}_i}$, where for all $n \in \{1 \cdots i\}$:

$$wu_{\mathbf{x}_n} = \begin{cases} WU(\mathbf{x}_n) & \text{if } \mathbf{x}_n \in WS_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

Given a transaction trace θ_ξ , we write e.g. $\xi.U_s$ to refer to its constituent U_s sub-trace and write $U_s.E$ for the set of events related by po in U_s . Similarly, we write $\xi.E$ for the set of events related by po in θ_ξ . Note that $G.E = \bigcup_{\xi \in \text{Tx}} \xi.E$.

Note that for each transaction ξ and each location \mathbf{x} , the $\xi.rl_{\mathbf{x}}$, $\xi.rs_{\mathbf{x}}$, $\xi.wr_{\mathbf{x}}$, $\xi.ru_{\mathbf{x}}$, $\xi.pl_{\mathbf{x}}$ and $\xi.wu_{\mathbf{x}}$ are uniquely identified when they exist. Let $\xi.w_{\mathbf{x}}$ the last (in po order) write to \mathbf{x} in W_s , when it exists.

For each location $\mathbf{x} \in \text{WS}_\xi$, let $fw_{\mathbf{x}}$ denote the maximal write (in po order within ξ) logging a write for \mathbf{x} in $\mathfrak{s}[\mathbf{x}]$. That is, when $\theta_\xi = t_1 \xrightarrow{\text{po}|_{\text{imm}}} \dots \xrightarrow{\text{po}|_{\text{imm}}} t_m$, let $fw_{\mathbf{x}} = \text{wmax}(\mathbf{x}, [t_1 \dots t_m])$, where

$$\text{wmax}(\mathbf{x}, L++[t]) \triangleq \begin{cases} \text{wmax}(\mathbf{x}, []) \text{ undefined} & \\ lw_{\mathbf{x}} & \text{if } t = wr(\mathbf{x}, -, -, -) \xrightarrow{\text{po}} lw_{\mathbf{x}} \xrightarrow{\text{po}} wws_{\mathbf{x}} \\ \text{wmax}(\mathbf{x}, L) & \text{otherwise} \end{cases}$$

F.1 Implementation Soundness

In order to establish the soundness of our implementation, it suffices to show that given an RA-consistent execution graph of the implementation G , we can construct a corresponding RSI-consistent execution graph G' with the same outcome.

Given a transaction $\xi \in \text{Tx}$ with $\text{RS}_\xi \cup \text{WS}_\xi = \{\mathbf{x}_1 \dots \mathbf{x}_i\}$ and trace θ_ξ as above with $T_s = t_1 \xrightarrow{\text{po}|_{\text{imm}}} \dots \xrightarrow{\text{po}|_{\text{imm}}} t_k$, we construct the corresponding RSI execution trace θ'_ξ as follows:

$$\theta'_\xi \triangleq t'_1 \xrightarrow{\text{po}|_{\text{imm}}} \dots \xrightarrow{\text{po}|_{\text{imm}}} t'_k$$

where for all $m \in \{1 \dots k\}$:

$$\begin{aligned} t'_m &= R(\mathbf{x}_m, v_m) \text{ when } t_m = rd(\mathbf{x}_m, v_m, S_m) \xrightarrow{\text{po}|_{\text{imm}}} lr_{\mathbf{x}_m} \\ t'_m &= W(\mathbf{x}_m, v_m) \text{ when } t_m = wr(\mathbf{x}_m, v_m, S_m) \xrightarrow{\text{po}|_{\text{imm}}} wws_{\mathbf{x}_m} \xrightarrow{\text{po}|_{\text{imm}}} \dots \end{aligned}$$

such that in the first case the identifier of t'_m is that of $lr_{\mathbf{x}_m}$; and in the second case the identifier of t'_m is that of $lw_{\mathbf{x}_m}$. Note that for each write operation w in θ'_ξ , there exists a *matching* write operation in $\theta_\xi.W_s$, denoted by $\text{mw}(w)$. That is,

$$\text{mw}(w) = w' \stackrel{\text{def}}{\iff} \exists i. \wedge (\theta'_\xi.E \cap \mathcal{W})|_i = w \wedge (\theta_\xi.W_s)|_i = w'$$

We then define:

$$\text{RF}_\xi \triangleq \left\{ (w, t'_j) \left| \begin{array}{l} t'_j \in \theta'_\xi.E \wedge \exists \mathbf{x}, v. t'_j = \mathbf{R}(\mathbf{x}, v) \wedge w = \mathbf{W}(\mathbf{x}, v) \\ \wedge (w \in \xi.E \Rightarrow w \xrightarrow{\text{po}} t'_j \wedge \\ \quad (\forall e \in \xi.E. w \xrightarrow{\text{po}} e \xrightarrow{\text{po}} t'_j \Rightarrow (\text{loc}(e) \neq \mathbf{x} \vee e \notin \mathcal{W}))) \\ \wedge (w \notin \xi.E \Rightarrow (\forall e \in \xi.E. (e \xrightarrow{\text{po}} t'_j \Rightarrow (\text{loc}(e) \neq \mathbf{x} \vee e \notin \mathcal{W})) \\ \quad \wedge ((\exists \xi'. (\xi'.w_{\mathbf{x}}, \xi'.rs_{\mathbf{x}}) \in G.\text{rf}) \wedge w = \xi'.fw_{\mathbf{x}}) \\ \quad \vee (w \in G.\mathcal{NT} \wedge (w, rs_{\mathbf{x}}) \in G.\text{rf}))) \end{array} \right. \right\}$$

$$\cup \left\{ (w, r) \left| \begin{array}{l} r \in G.\mathcal{NT} \wedge \\ ((w \in G.\mathcal{NT} \wedge (w, r) \in G.\text{rf}) \\ \vee (\exists w'. w' = \mathbf{mw}(w) \wedge (w', r) \in G.\text{rf})) \end{array} \right. \right\}$$

Similarly, we define:

$$\text{MO} \triangleq \left\{ (w_1, w_2) \mid \exists w'_1, w'_2. w'_1 = \mathbf{mw}(w_1) \wedge w'_2 = \mathbf{mw}(w_2) \wedge (w'_1, w'_2) \in G.\text{mo} \right\}$$

$$\cup \left\{ (w, w') \mid w, w' \in G.\mathcal{NT} \wedge (w, w') \in G.\text{mo} \right\}$$

$$\cup \left\{ (w, w') \mid w' \in G.\mathcal{NT} \wedge \exists w''. w'' = \mathbf{mw}(w) \wedge (w'', w') \in G.\text{mo} \right\}$$

$$\cup \left\{ (w', w) \mid w' \in G.\mathcal{NT} \wedge \exists w''. w'' = \mathbf{mw}(w) \wedge (w', w'') \in G.\text{mo} \right\}$$

We are now in a position to demonstrate the soundness of our implementation. Given an RA-consistent execution graph G of the implementation, we construct an RSI execution graph G' as follows and demonstrate that $\text{rsi-consistent}(G')$ holds.

– $G'.E = \bigcup_{\xi \in \text{Tx}} \theta'_\xi.E \cup G.\mathcal{NT}$, with the $\text{tx}(\cdot)$ function defined as:

$$\text{tx}(a) \triangleq \xi \quad \text{when } a \in \theta'_\xi \quad \text{tx}(a) \triangleq 0 \quad \text{when } a \in G.\mathcal{NT}$$

- $G'.\text{po} = G.\text{po}|_{G'.E}$
- $G'.\text{rf} = \bigcup_{\xi \in \text{Tx}} \text{RF}_\xi$
- $G'.\text{mo} = \text{MO}$
- $G'.\text{lo} = \emptyset$

Observe that the events of each θ'_ξ trace coincides with those of the equivalence class \mathcal{T}_ξ of G' . That is, $\theta'_\xi.E = \mathcal{T}_\xi$.

Lemma 17. *Given an RA-consistent execution graph G of the implementation and its corresponding RSI execution graph G' constructed as above, for all $a, b, \xi_a, \xi_b, \mathbf{x}$:*

$$\xi_a \neq \xi_b \wedge a \in \xi_a.E \wedge b \in \xi_b.E \wedge \text{loc}(a) = \text{loc}(b) = \mathbf{x} \Rightarrow$$

$$((a, b) \in G'.\text{rf} \Rightarrow \xi_a.wu_{\mathbf{x}} \xrightarrow{G.\text{hb}} \xi_b.rl_{\mathbf{x}}) \tag{34}$$

$$\wedge ((a, b) \in G'.\text{mo} \Rightarrow \xi_a.wu_{\mathbf{x}} \xrightarrow{G.\text{hb}} \xi_b.rl_{\mathbf{x}}) \tag{35}$$

$$\wedge ((a, b) \in G'.\text{rb} \Rightarrow (\mathbf{x} \in \text{WS}_{\xi_a} \wedge \xi_a.wu_{\mathbf{x}} \xrightarrow{G.\text{hb}} \xi_b.rl_{\mathbf{x}}) \vee (\mathbf{x} \notin \text{WS}_{\xi_a} \wedge \xi_a.ru_{\mathbf{x}} \xrightarrow{G.\text{hb}} \xi_b.pl_{\mathbf{x}})) \tag{36}$$

$$\wedge ((a, b) \in G'.(\text{mo}; \text{rf}) \Rightarrow \xi_a.wu_{\mathbf{x}} \xrightarrow{G.\text{hb}} \xi_b.rl_{\mathbf{x}}) \tag{37}$$

Proof. Pick an arbitrary RA-consistent execution graph G of the implementation and its corresponding RSI execution graph G' constructed as above. Pick an arbitrary $a, b, \xi_a, \xi_b, \mathbf{x}$ such that $\xi_a \neq \xi_b$, $a \in \xi_a.E$, $a \in \xi_a.E$, and $\text{loc}(a) = \text{loc}(b) = \mathbf{x}$.

RTS. (34)

Assume $(a, b) \in G'.\text{rf}$. From the definition of $G'.\text{rf}$ we then know $(\xi_a.w_x, \xi_b.rs_x) \in G.\text{rf}$. On the other hand, from [Lemma 2](#) we know that either i) $\mathbf{x} \in \text{WS}_{\xi_b}$ and $\xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x$; or ii) $\mathbf{x} \notin \text{WS}_{\xi_b}$ and $\xi_b.ru_x \xrightarrow{G.\text{hb}} \xi_a.pl_x$; or iii) $\xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x$. In case (i) we then have $\xi_a.w_x \xrightarrow{G.\text{rf}} \xi_b.rs_x \xrightarrow{G.\text{po}} \xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x \xrightarrow{G.\text{po}} \xi_a.w_x$. That is, we have $\xi_a.w_x \xrightarrow{G.\text{hb}^{\text{loc}}} \xi_a.w_x$, contradicting the assumption that G is RA-consistent. Similarly in case (ii) we have $\xi_a.w_x \xrightarrow{G.\text{rf}} \xi_b.rs_x \xrightarrow{G.\text{po}} \xi_b.ru_x \xrightarrow{G.\text{hb}} \xi_a.pl_x \xrightarrow{G.\text{po}} \xi_a.w_x$. That is, we have $\xi_a.w_x \xrightarrow{G.\text{hb}^{\text{loc}}} \xi_a.w_x$, contradicting the assumption that G is RA-consistent. In case (iii) the desired result holds trivially.

RTS. (35)

Assume $(a, b) \in G'.\text{mo}$. From the definition of $G'.\text{mo}$ we then know there exist $w_1 \in \xi_a.Ws$ and $w_2 \in \xi_b.Ws$ such that $(w_1, w_2) \in G.\text{mo}$ and $\text{loc}(w_1) = \text{loc}(w_2) = \mathbf{x}$. On the other hand, from [Lemma 2](#) we know that either i) $\xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x$; or ii) $\xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x$. In case (i) we then have $w_1 \xrightarrow{G.\text{mo}} w_2 \xrightarrow{G.\text{po}} \xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x \xrightarrow{G.\text{po}} w_1$. That is, we have $w_1 \xrightarrow{G.\text{mo}} w_2 \xrightarrow{G.\text{hb}^{\text{loc}}} w_1$, contradicting the assumption that G is RA-consistent. In case (ii) the desired result holds trivially.

RTS. (36)

Assume $(a, b) \in G'.\text{rb}$. From the definition of $G'.\text{rb}$ we then know that there exist $w \in \xi_b.Ws$ such that $\text{loc}(w) = \mathbf{x}$ and $(\xi_a.rs_x, w) \in G.\text{rb}$. From [Lemma 2](#) we then know that either i) $\xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x$; or ii) $\mathbf{x} \notin \text{WS}_{\xi_a}$ and $\xi_a.ru_x \xrightarrow{G.\text{hb}} \xi_a.pl_x$; or iii) $\mathbf{x} \in \text{WS}_{\xi_a}$ and $\xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x$. In case (i) we then have $w \xrightarrow{G.\text{po}} \xi_b.wu_x \xrightarrow{G.\text{hb}} \xi_a.rl_x \xrightarrow{G.\text{po}} \xi_a.rs_x \xrightarrow{G.\text{rb}} w$. That is, we have $w \xrightarrow{G.\text{hb}^{\text{loc}}} \xi_a.rs_x \xrightarrow{G.\text{rb}} w$, contradicting the assumption that G is RA-consistent. In cases (ii-iii) the desired result holds trivially.

Lemma 18. For all RA-consistent execution graphs G of the implementation and their counterpart RSI execution graphs G' constructed as above and $S \triangleq$

$G'.(\text{rsi-po} \cup \text{rsi-rf} \cup \text{mo}_\top)^+$:

$$\begin{aligned} & \forall \xi_a, \xi_b, a, b. \forall a, b. \\ & (a, b) \in S \Rightarrow (a, b \in G'.\mathcal{NT} \wedge (a, b) \in G.\mathbf{hb}) \\ & \quad \vee (a \in G'.\mathcal{T}_{\xi_a} \wedge b \in G'.\mathcal{T}_{\xi_b} \wedge \xi_a = \xi_b \wedge (a, b) \in G.\text{po}) \\ & \quad \vee \left(\begin{array}{l} a \in G'.\mathcal{T}_{\xi_a} \wedge b \in G'.\mathcal{T}_{\xi_b} \wedge \xi_a \neq \xi_b \\ \wedge \exists d \in G'.\mathcal{T}_{\xi_b}. \forall c \in G'.\mathcal{T}_{\xi_a}. (c, d) \in G.\mathbf{hb} \\ \wedge a \in \mathcal{W} \Rightarrow (\text{mw}(a), d) \in G.\mathbf{hb} \end{array} \right) \\ & \quad \vee \left(\begin{array}{l} a \in G'.\mathcal{NT} \wedge b \in G'.\mathcal{T}_{\xi_b} \\ \wedge \exists d \in G'.\mathcal{T}_{\xi_b}. (a, d) \in G.\mathbf{hb} \end{array} \right) \\ & \quad \vee \left(\begin{array}{l} a \in G'.\mathcal{T}_{\xi_a} \wedge b \in G'.\mathcal{NT} \\ \wedge \forall c \in G'.\mathcal{T}_{\xi_a}. (c, b) \in G.\mathbf{hb} \\ \wedge (\exists e \in \xi_a. \text{Ws. } (e, b) \in G.\mathbf{hb} \vee \forall c \in \xi_a. E. (c, b) \in G.\mathbf{hb}) \end{array} \right) \end{aligned}$$

Proof. Let $S_0 = G'.(\text{rsi-po} \cup \text{rsi-rf} \cup \text{mo}_\top)$, and $S_{n+1} = S_0; S_n$, for all $n \geq 0$. It is straightforward to demonstrate that $S = \bigcup_{i \in \mathbb{N}} S_i$. We thus demonstrate instead that:

$$\begin{aligned} & \forall i \in \mathbb{N}. \forall \xi_a, \xi_b, a, b. \forall a, b. \\ & (a, b) \in S_i \Rightarrow (a, b \in G'.\mathcal{NT} \wedge (a, b) \in G.\mathbf{hb}) \\ & \quad \vee (a \in G'.\mathcal{T}_{\xi_a} \wedge b \in G'.\mathcal{T}_{\xi_b} \wedge \xi_a = \xi_b \wedge (a, b) \in G.\text{po}) \\ & \quad \vee \left(\begin{array}{l} a \in G'.\mathcal{T}_{\xi_a} \wedge b \in G'.\mathcal{T}_{\xi_b} \wedge \xi_a \neq \xi_b \\ \wedge \exists d \in G'.\mathcal{T}_{\xi_b}. \forall c \in G'.\mathcal{T}_{\xi_a}. (c, d) \in G.\mathbf{hb} \\ \wedge (\exists e \in \xi_a. \text{Ws. } (e, d) \in G.\mathbf{hb} \\ \quad \vee \forall c \in \xi_a. E. (c, d) \in G.\mathbf{hb}) \end{array} \right) \\ & \quad \vee \left(\begin{array}{l} a \in G'.\mathcal{NT} \wedge b \in G'.\mathcal{T}_{\xi_b} \\ \wedge \exists d \in G'.\mathcal{T}_{\xi_b}. (a, d) \in G.\mathbf{hb} \end{array} \right) \\ & \quad \vee \left(\begin{array}{l} a \in G'.\mathcal{T}_{\xi_a} \wedge b \in G'.\mathcal{NT} \\ \wedge \forall c \in G'.\mathcal{T}_{\xi_a}. (c, b) \in G.\mathbf{hb} \\ \wedge (\exists e \in \xi_a. \text{Ws. } (e, b) \in G.\mathbf{hb} \\ \quad \vee \forall c \in \xi_a. E. (c, b) \in G.\mathbf{hb}) \end{array} \right) \end{aligned}$$

We proceed by induction on i .

Base case $i = 0$

Pick arbitrary ξ_a, ξ_b, a, b such that $\xi_a \neq \xi_b$ and $(a, b) \in S_0$. There are now four cases to consider: A) $a \in G'.\mathcal{T}_{\xi_a}, b \in G'.\mathcal{T}_{\xi_b}$; or B) $a, b \in G'.\mathcal{NT}$; or C) $a \in G'.\mathcal{NT}, b \in G'.\mathcal{T}_{\xi_b}$; or D) $a \in G'.\mathcal{T}_{\xi_a}, b \in G'.\mathcal{NT}$.

In case (A) there are three additional cases to consider: 1) $(a, b) \in G'.\text{rsi-po}$; or 2) $(a, b) \in G'.\text{rsi-rf}$; or 3) $(a, b) \in G'.\text{mo}_\top$.

In case (A.1), pick an arbitrary $c \in G'.\mathcal{T}_{\xi_a}$. From the definition of $G'.\text{po}_\top$ we have $(c, b) \in G.\text{po} \subseteq G.\mathbf{hb}$, as required. Now assume that $a \in \mathcal{W}$. From the definition of $G'.\text{po}$ we then have $(\text{mw}(a), d) \in G.\text{po} \subseteq G.\mathbf{hb}$, as required.

In case (A.2), we then know there exists $w \in G'.\mathcal{T}_{\xi_a}$ and $r \in G'.\mathcal{T}_{\xi_b}$ such that $(w, r) \in G'.\text{rf}$. Let $\text{loc}(w) = \text{loc}(r) = \mathbf{x}$. From [Lemma 17](#) we then have $\xi_a.wu_{\mathbf{x}} \xrightarrow{G.\mathbf{hb}} \xi_b.rl_{\mathbf{x}}$. Pick an arbitrary $c \in G'.\mathcal{T}_{\xi_a}$. As such we have

$c \xrightarrow{G.\text{po}} \xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x \xrightarrow{G.\text{po}} \xi_b.rs_x$. That is, we have $(c, \xi_b.rs_x) \in G.\text{hb}$, as required. Now assume that $a \in \mathcal{W}$. We then have $(\text{mw}(a), \xi_a.wu_x) \in G.\text{po}$. As such, from the transitivity of $G.\text{hb}$ we also have $(\text{mw}(a), \xi_b.rs_x) \in G.\text{hb}$, as required.

In case (A.3), we then know there exists $w \in G'.\mathcal{T}_{\xi_a}$ and $w' \in G'.\mathcal{T}_{\xi_b}$ such that $(w, w') \in G'.\text{mo}$. Let $\text{loc}(w) = \text{loc}(w') = \mathbf{x}$. From Lemma 17 we then have $\xi_a.wu_x \xrightarrow{G.\text{hb}} \xi_b.rl_x$. Pick an arbitrary $c \in G'.\mathcal{T}_{\xi_a}$. As we also have $\xi_b.rl_x \xrightarrow{G.\text{po}} w_2$, $w_1 \xrightarrow{G.\text{po}} \xi_a.wu_x$, and $c \xrightarrow{G.\text{po}} \xi_a.wu_x$, we then have $(c, w_2) \in G.\text{hb}$ and $(w_1, w_2) \in G.\text{hb}$, as required. Now assume that $a \in \mathcal{W}$. We then have $(\text{mw}(a), \xi_a.wu_x) \in G.\text{po}$. As such, from the transitivity of $G.\text{hb}$ we also have $(\text{mw}(a), w_2) \in G.\text{hb}$, as required.

In case (B) there are three additional cases to consider: 1) $(a, b) \in G'.\text{rsi-po}$; or 2) $(a, b) \in G'.\text{rsi-rf}$; or 3) $(a, b) \in G'.\text{mo}_\top$.

In case (B.1) we then have $(a, b) \in G'.\text{po}$ and from the definition of $G'.\text{po}$ we have $(a, b) \in G.\text{po} \subseteq G.\text{hb}$, as required. In case (B.2) we then have $(a, b) \in G'.\text{rf}$ and from the definition of $G'.\text{rf}$ we have $(a, b) \in G.\text{rf} \subseteq G.\text{hb}$, as required. Case (B.3) does not apply as $a, b \in G'.\mathcal{NT}$.

In case (C) there are three additional cases to consider: 1) $(a, b) \in G'.\text{rsi-po}$; or 2) $(a, b) \in G'.\text{rsi-rf}$; or 3) $(a, b) \in G'.\text{mo}_\top$.

In case (C.1) we then have $(a, b) \in G'.\text{po}$ and from the definition of $G'.\text{po}$ we have $(a, b) \in G.\text{po} \subseteq G.\text{hb}$, as required. In case (C.2) we know there exists $r \in G'.\mathcal{T}_{\xi_b}$ such that $(a, r) \in G'.\text{rf}$. Let $\text{loc}(a) = \text{loc}(r) = \mathbf{x}$. From the definition of $G'.\text{rf}$ we then have $(a, \xi_b.rs_x) \in G.\text{rf} \subseteq G.\text{hb}$, as required. Case (C.3) does not apply as $a \in G'.\mathcal{NT}$.

In case (D) there are three additional cases to consider: 1) $(a, b) \in G'.\text{rsi-po}$; or 2) $(a, b) \in G'.\text{rsi-rf}$; or 3) $(a, b) \in G'.\text{mo}_\top$.

In case (D.1) we then have $(a, b) \in G'.\text{po}$. Pick an arbitrary $c \in G'.\mathcal{T}_{\xi_a}$. From the definition of $G'.\text{po}$ we have $(c, b) \in G'.\text{po}$. As such, from the definition of $G'.\text{po}$ we have $(c, b) \in G.\text{po} \subseteq G.\text{hb}$, as required. Now assume that $a \in \mathcal{W}$. We then have $(\text{mw}(a), b) \in G.\text{po} \subseteq G.\text{hb}$, as required.

In case (D.2) we then have $(a, b) \in G'.\text{rf}$ and from the definition of $G'.\text{rf}$ we know that $\text{mw}(a) \in \xi_a.Ws$ and $(\text{mw}(a), b) \in G.\text{rf} \subseteq G.\text{hb}$, as required. Pick an arbitrary $c \in G'.\mathcal{T}_{\xi_a}$. We then know $(c, \text{mw}(a)) \in G.\text{po}$. We then have $c \xrightarrow{G.\text{po}} \text{mw}(a) \xrightarrow{G.\text{rf}} b$. As such, from the definition of $G.\text{hb}$ we have $(c, b) \in G.\text{hb}$, as required. Case (D.3) does not apply as $b \in G'.\mathcal{NT}$.

Inductive case $i=n+1$

$\forall j \in \mathbb{N}. \forall \xi_a, \xi_b, a, b. \forall a, b.$

$$\begin{aligned}
(a, b) \in S_j \wedge j \leq n \Rightarrow & (a, b \in G'.\mathcal{NT} \wedge (a, b) \in G.\mathbf{hb}) \\
& \vee (a \in G'.\mathcal{T}_{\xi_a} \wedge b \in G'.\mathcal{T}_{\xi_b} \wedge \xi_a = \xi_b \wedge (a, b) \in G.\mathbf{po}) \\
& \vee \left(\begin{array}{l} a \in G'.\mathcal{T}_{\xi_a} \wedge b \in G'.\mathcal{T}_{\xi_b} \wedge \xi_a \neq \xi_b \\ \exists d \in G'.\mathcal{T}_{\xi_b}. \forall c \in G'.\mathcal{T}_{\xi_a}. (c, d) \in G.\mathbf{hb} \\ \wedge (\exists e \in \xi_a. \text{Ws. } (e, d) \in G.\mathbf{hb} \\ \quad \forall c \in \xi_a. E. (c, d) \in G.\mathbf{hb}) \end{array} \right) \\
& \vee \left(\begin{array}{l} a \in G'.\mathcal{NT} \wedge b \in G'.\mathcal{T}_{\xi_b} \\ \wedge \exists d \in G'.\mathcal{T}_{\xi_b}. (a, d) \in G.\mathbf{hb} \end{array} \right) \\
& \vee \left(\begin{array}{l} a \in G'.\mathcal{T}_{\xi_a} \wedge b \in G'.\mathcal{NT} \\ \wedge \forall c \in G'.\mathcal{T}_{\xi_a}. (c, b) \in G.\mathbf{hb} \\ \wedge (\exists e \in \xi_a. \text{Ws. } (e, b) \in G.\mathbf{hb} \\ \quad \forall c \in \xi_a. E. (c, b) \in G.\mathbf{hb}) \end{array} \right)
\end{aligned} \tag{I.H.}$$

Pick arbitrary $\xi_a, \xi_b, a \in G'.\mathcal{T}_{\xi_a}, b \in G'.\mathcal{T}_{\xi_b}$ such that $(a, b) \in S_i$. From the definition of S_i we then know there exist e, ξ_e such that $e \in G'.\mathcal{T}_{\xi_e}, (a, e) \in S_0$ and $(e, b) \in S_n$. The desired result then follows from the inductive hypothesis and case analysis on a, b and c .

Theorem 13 (Soundness). For all execution graphs G of the implementation and their counterpart RSI execution graphs G' constructed as above,

$$\text{RA-consistent}(G) \Rightarrow \text{rsi-consistent}(G')$$

Proof. Pick an arbitrary execution graph G of the implementation such that $\text{RA-consistent}(G)$ holds, and its associated RSI execution graph G' constructed as described above.

RTS. $\text{irreflexive}(G'.\mathbf{rsi-hb})$

We proceed by contradiction. Let us assume $\neg \text{irreflexive}(G'.\mathbf{rsi-hb})$. Let $S = G'.(\mathbf{rsi-po} \cup \mathbf{rsi-rf} \cup \mathbf{mo}_T)^+$. There are now two cases to consider: either there is an $\mathbf{rsi-hb}$ cycle without a $\mathbf{si-rb}$ edge; or there is a cycle with one or more $\mathbf{si-rb}$ edges. That is, either 1) there exists a such that $(a, a) \in S$; or 2) there exist $a_1, b_1, \dots, a_n, b_n$ such that $a_1 \xrightarrow{G'.\mathbf{si-rb}} b_1 \xrightarrow{S} a_2 \xrightarrow{G'.\mathbf{si-rb}} b_2 \xrightarrow{S} \dots \xrightarrow{S} a_n \xrightarrow{G'.\mathbf{si-rb}} b_n \xrightarrow{S} a_1$.

In case (1) we then know that either i) there exists ξ such that $a \in G'.\mathcal{T}_{\xi}$; or ii) $a \in G'.\mathcal{NT}$. In case (1.i) from [Lemma 18](#) we then have $(a, a) \in G.\mathbf{po}$, contradicting the assumption that G is RA-consistent. Similarly, in case (1.ii) from [Lemma 18](#) we know that $(a, a) \in G.\mathbf{hb}$, contradicting the assumption that G is RA-consistent.

In case (2), for an arbitrary $i \in \{1 \dots n\}$, let $j = i+1$ when $i \neq n$; and $j = 1$ when $i = n$. As $a_i \xrightarrow{G'.\mathbf{si-rb}} b_i$, we know there exists ξ_{a_i}, ξ_{b_i} such that $a_i \in \mathcal{T}_{\xi_{a_i}}, b_i \in \mathcal{T}_{\xi_{b_i}}$, and that there exist $r_i \in \mathcal{T}_{\xi_{a_i}} \cap \mathcal{R}_E$ and $w_i \in \mathcal{T}_{\xi_{b_i}} \cap \mathcal{W}$

such that $(r_i, w_i) \in G'.\mathbf{rb}$. Let $\text{loc}(r_i) = \text{loc}(w_i) = \mathbf{x}_i$. From [Lemma 17](#) we then know that either i) $\xi_{a_i}.ru_{\mathbf{x}_i} \xrightarrow{G.\mathbf{hb}} \xi_{b_i}.pl_{\mathbf{x}_i}$; or ii) $\xi_{a_i}.wu_{\mathbf{x}_i} \xrightarrow{G.\mathbf{hb}} \xi_{b_i}.rl_{\mathbf{x}_i}$. Note that for all $w' \in \xi_{b_i}.Ws$, we know $\xi_{b_i}.rl_{\mathbf{x}_i} \xrightarrow{G.\mathbf{po}} w'$ and $\xi_{b_i}.pl_{\mathbf{x}_i} \xrightarrow{G.\mathbf{po}} w'$. As such, as $(b_i, a_j) \in S$, from [Lemma 18](#) and since $G.\mathbf{hb}$ is transitively closed, we know there exists $d_j \in \mathcal{T}_{\xi_{a_j}}$ such that either $\xi_{a_i}.ru_{\mathbf{x}_i} \xrightarrow{G.\mathbf{hb}} d_j$, or $\xi_{a_i}.wu_{\mathbf{x}_i} \xrightarrow{G.\mathbf{hb}} d_j$. That is, $\xi_{a_i}.u_{\mathbf{x}_i} \xrightarrow{G.\mathbf{hb}} d_j$, where either $u_{\mathbf{x}_i} = ru_{\mathbf{x}_i}$ or $u_{\mathbf{x}_i} = wu_{\mathbf{x}_i}$. On the other hand, observe that for all $d_i \in \mathcal{T}_{\xi_{a_i}}$ we have $d_i \xrightarrow{G.\mathbf{po}} \xi_{a_i}.u_{\mathbf{x}_i}$ i.e. $d_i \xrightarrow{G.\mathbf{hb}} \xi_{a_i}.u_{\mathbf{x}_i}$. As such, we have $d_j \xrightarrow{G.\mathbf{hb}} \xi_{a_j}.u_{\mathbf{x}_j}$. As we also have $\xi_{a_i}.u_{\mathbf{x}_i} \xrightarrow{G.\mathbf{hb}} d_j$ and $G.\mathbf{hb}$ is transitively closed, we have $\xi_{a_i}.u_{\mathbf{x}_i} \xrightarrow{G.\mathbf{hb}} \xi_{a_j}.u_{\mathbf{x}_j}$. We then have $\xi_{a_1}.u_{\mathbf{x}_1} \xrightarrow{G.\mathbf{hb}} \xi_{a_2}.u_{\mathbf{x}_2} \xrightarrow{G.\mathbf{hb}} \dots \xrightarrow{G.\mathbf{hb}} \xi_{a_n}.u_{\mathbf{x}_n} \xrightarrow{G.\mathbf{hb}} \xi_{a_1}.u_{\mathbf{x}_1}$. That is, $\xi_{a_1}.u_{\mathbf{x}_1} \xrightarrow{G.\mathbf{hb}} \xi_{a_1}.u_{\mathbf{x}_1}$, contradicting the assumption that G is RA-consistent.

RTS. $\mathbf{rf}_I \cup \mathbf{mo}_I \cup \mathbf{rb}_I \subseteq \mathbf{po}$

Follows immediately from the construction of G' .

RTS. $\text{irreflexive}(G'.(\mathbf{rsi}\text{-}\mathbf{hb}; \mathbf{mo}))$

We proceed by contradiction. Let us assume $\neg \text{irreflexive}(G'.(\mathbf{rsi}\text{-}\mathbf{hb}; \mathbf{mo}))$.

That is, there exists a, b such that $(a, b) \in \mathbf{rsi}\text{-}\mathbf{hb}$ and $(b, a) \in \mathbf{mo}$. There are now five cases to consider: 1) there exists ξ such that $a, b \in \mathcal{T}_\xi$; or 2) there exists ξ_a, ξ_b such that $\xi_a \neq \xi_b$, $a \in \mathcal{T}_{\xi_a}$ and $b \in \mathcal{T}_{\xi_b}$; or 3) there exists ξ such that $a \in \mathcal{T}_\xi$ and $b \in G'.\mathcal{NT}$; or 4) there exists ξ such that $b \in \mathcal{T}_\xi$ and $a \in G'.\mathcal{NT}$; or 5) $a, b \in G'.\mathcal{NT}$.

In case (1) we then have $(b, a) \in \mathbf{mo}_I \subseteq \mathbf{po}_I$ (from the proof of the previous part) and since a, b are both write events, we have $(b, a) \subseteq \mathbf{rsi}\text{-}\mathbf{hb}$. We then have $a \xrightarrow{\mathbf{rsi}\text{-}\mathbf{hb}} b \xrightarrow{\mathbf{rsi}\text{-}\mathbf{hb}} a$, contradicting our proof above that $\mathbf{rsi}\text{-}\mathbf{hb}$ is irreflexive.

In case (2) we then have $(b, a) \in \mathbf{mo}_T \subseteq \mathbf{rsi}\text{-}\mathbf{hb}$. We then have $a \xrightarrow{\mathbf{rsi}\text{-}\mathbf{hb}} b \xrightarrow{\mathbf{rsi}\text{-}\mathbf{hb}} a$, contradicting our proof above that $\mathbf{rsi}\text{-}\mathbf{hb}$ is irreflexive.

In case (3) from [Lemma 18](#) we know $(\mathbf{mw}(a), b) \in G.\mathbf{hb}$. From the definition of $G'.\mathbf{mo}$ we also have $(b, \mathbf{mw}(a)) \in G.\mathbf{mo}$. We then have $\mathbf{mw}(a) \xrightarrow{G.\mathbf{hb}} b \xrightarrow{G.\mathbf{mo}} \mathbf{mw}(a)$, contradicting the assumption that G is RA-consistent.

In case (4) from [Lemma 18](#) we know there exists $d \in \mathcal{T}_\xi$ such that $(a, d) \in G.\mathbf{hb}$. From the definition of $G'.\mathbf{mo}$ we also have $(\mathbf{mw}(b), a) \in G.\mathbf{mo}$. Moreover, from the construction of G' we know $(d, \mathbf{mw}(b)) \in G.\mathbf{po}$. We then have $a \xrightarrow{G.\mathbf{hb}} d \xrightarrow{G.\mathbf{po}} \mathbf{mw}(b) \xrightarrow{G.\mathbf{mo}} a$, contradicting the assumption that G is RA-consistent.

In case (5) from [Lemma 18](#) we know $(a, b) \in G.\mathbf{hb}$. From the definition of $G'.\mathbf{mo}$ we also have $(b, a) \in G.\mathbf{mo}$. We then have $a \xrightarrow{G.\mathbf{hb}} b \xrightarrow{G.\mathbf{mo}} a$, contradicting the assumption that G is RA-consistent.

RTS. $\text{irreflexive}(G'.(\mathbf{rsi}\text{-}\mathbf{hb}; \mathbf{rb}))$

We proceed by contradiction. Let us assume $\neg \text{irreflexive}(G'.(\mathbf{rsi}\text{-}\mathbf{hb}; \mathbf{rb}))$.

That is, there exists a, b such that $(a, b) \in \mathbf{rsi}\text{-}\mathbf{hb}$ and $(b, a) \in \mathbf{rb}$. There are now

five cases to consider: 1) there exists ξ such that $a, b \in \mathcal{T}_\xi$; or 2) there exists ξ_a, ξ_b such that $\xi_a \neq \xi_b$, $a \in \mathcal{T}_{\xi_a}$ and $b \in \mathcal{T}_{\xi_b}$; or 3) there exists ξ such that $a \in \mathcal{T}_\xi$ and $b \in G'.\mathcal{NT}$; or 4) there exists ξ such that $b \in \mathcal{T}_\xi$ and $a \in G'.\mathcal{NT}$; or 5) $a, b \in G'.\mathcal{NT}$.

In case (1) we then have $(b, a) \in \mathbf{rb}_I \subseteq \mathbf{po}_I$ (from the proof of the earlier part) and thus $(b, a) \in G.\mathbf{po}$. Moreover, from [Lemma 18](#) we know $(a, b) \in G.\mathbf{po}$. We then have $a \xrightarrow{G.\mathbf{po}} b \xrightarrow{G.\mathbf{po}} a$, contradicting the assumption that G is RA-consistent.

In case (2) we then have $(b, a) \in \mathbf{si}\text{-}\mathbf{rb} \cup \mathbf{mo}_T \subseteq \mathbf{rsi}\text{-}\mathbf{hb}$. We then have $a \xrightarrow{\mathbf{rsi}\text{-}\mathbf{hb}} b \xrightarrow{\mathbf{rsi}\text{-}\mathbf{hb}} a$, contradicting our proof above that $\mathbf{rsi}\text{-}\mathbf{hb}$ is irreflexive.

In case (3) from [Lemma 18](#) we know $(\mathbf{mw}(a), b) \in G.\mathbf{hb}$. From the definition of $G'.\mathbf{rb}$ we also have $(b, \mathbf{mw}(a)) \in G.\mathbf{rb}$. We then have $\mathbf{mw}(a) \xrightarrow{G.\mathbf{hb}} b \xrightarrow{G.\mathbf{rb}} \mathbf{mw}(a)$, contradicting the assumption that G is RA-consistent.

In case (4) from [Lemma 18](#) we know there exists $d \in \mathcal{T}_\xi$ such that $(a, d) \in G.\mathbf{hb}$. Let $\mathbf{loc}(a) = \mathbf{loc}(b) = \mathbf{x}$. From the definition of $G'.\mathbf{rb}$ and our race-freedom stipulation of non-transactional writes with the same transaction we also have $(\xi_b.\mathbf{vr}_x, a) \in G.\mathbf{rb}$. Moreover, from the construction of G' we know $(d, \xi_b.\mathbf{vr}_x, b) \in G.\mathbf{po}$. We then have $a \xrightarrow{G.\mathbf{hb}} d \xrightarrow{G.\mathbf{po}} \xi_b.\mathbf{vr}_x \xrightarrow{G.\mathbf{rb}} a$, contradicting the assumption that G is RA-consistent.

In case (5) from [Lemma 18](#) we know $(a, b) \in G.\mathbf{hb}$. From the definition of $G'.\mathbf{rb}$ we also have $(b, a) \in G.\mathbf{rb}$. We then have $a \xrightarrow{G.\mathbf{hb}} b \xrightarrow{G.\mathbf{rb}} a$, contradicting the assumption that G is RA-consistent.

F.2 Implementation Completeness

In order to establish the completeness of our implementation, it suffices to show that given an RSI-consistent execution graph $G' = (E, \mathbf{po}, \mathbf{rf}, \mathbf{mo}, \mathbf{lo})$, we can construct a corresponding RA-consistent execution graph G of the implementation.

Note that the execution trace for each transaction $\mathcal{T}_\xi \in G'.\mathcal{T}/\mathbf{st}$ is of the form $\theta'_\xi = t'_1 \xrightarrow{\mathbf{po}|_{\text{imm}}} \dots \xrightarrow{\mathbf{po}|_{\text{imm}}} t'_k$ for some k , where each t'_i is a read or write event. As such, we have $G'.\mathcal{T} = \bigcup_{\mathcal{T}_\xi \in G'.\mathcal{T}/\mathbf{st}} \mathcal{T}_\xi = \theta'_\xi.E$. For each transaction ξ , we construct the implementation trace θ_ξ as follows:

$$Fs \xrightarrow{\mathbf{po}|_{\text{imm}}} Is \xrightarrow{\mathbf{po}|_{\text{imm}}} Ts \xrightarrow{\mathbf{po}|_{\text{imm}}} VRs \xrightarrow{\mathbf{po}|_{\text{imm}}} RUs \xrightarrow{\mathbf{po}|_{\text{imm}}} PLs \xrightarrow{\mathbf{po}|_{\text{imm}}} Ws \xrightarrow{\mathbf{po}|_{\text{imm}}} Us$$

where:

- Fs denotes the sequence of events failing to obtain the necessary locks (i.e. those iterations that do not succeed in promoting the writer locks) or validate the snapshot;
- Is denotes the sequence of events initialising the values of \mathbf{LS} , \mathbf{RS} and \mathbf{WS} with \emptyset , initialising \mathbf{ws} with \square and initialising $\mathbf{s}[\mathbf{x}]$ with (\perp, \perp) for each location \mathbf{x} ;

- T_s denotes the sequence of events corresponding to the execution of (\mathbb{T}) and is of the form $t_1 \xrightarrow{\text{pol}_{\text{imm}}} \dots \xrightarrow{\text{pol}_{\text{imm}}} t_k$, where for all $m \in \{1 \dots k\}$:

$$t_m = \begin{cases} rd(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \xrightarrow{\text{pol}_{\text{imm}}} lr_{\mathbf{x}_m} & \text{if } O_m = R(\mathbf{x}_m, v_m) \\ wr(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \xrightarrow{\text{pol}_{\text{imm}}} wws_{\mathbf{x}_m} & \text{if } O_m = W(\mathbf{x}_m, v_m) \\ \xrightarrow{\text{pol}_{\text{imm}}} R(\mathbf{s}[\mathbf{x}_m], (v'_m, -)) \xrightarrow{\text{pol}_{\text{imm}}} lw_{\mathbf{x}_m} \xrightarrow{\text{pol}_{\text{imm}}} W(\mathbf{ws}, ws_m) & \end{cases}$$

where O_m denotes the m th event in the trace of \mathbb{T} ; $lr_{\mathbf{x}_m} \triangleq R(\mathbf{s}[\mathbf{x}_m], (-, v_m))$;

$$rd(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \triangleq \begin{cases} R(\mathbf{s}[\mathbf{x}_m], (\perp, \perp)) & \text{if } \mathbf{x}_m \notin RS_{m-1} \cup WS_{m-1} \\ \xrightarrow{\text{pol}_{\text{imm}}} fs_m \\ \xrightarrow{\text{pol}_{\text{imm}}} rl_{\mathbf{x}_m} \\ \xrightarrow{\text{pol}_{\text{imm}}} wrs_{\mathbf{x}_m} \\ \xrightarrow{\text{pol}_{\text{imm}}} rs_{\mathbf{x}_m} \\ \xrightarrow{\text{pol}_{\text{imm}}} ws_{\mathbf{x}_m} \\ \emptyset & \text{otherwise} \end{cases}$$

fs_m denotes the sequence of events attempting (but failing) to acquire the read lock on \mathbf{x}_m , $rl_{\mathbf{x}_m} \triangleq RL(\mathbf{x}_m)$, $wrs_{\mathbf{x}_m} \triangleq W(RS, RS_m)$, $rs_{\mathbf{x}_m} \triangleq R(\mathbf{x}_m, v_{\mathbf{x}_m}^0)$, $ws_{\mathbf{x}_m} \triangleq W(\mathbf{s}[\mathbf{x}_m], (v_{\mathbf{x}_m}^0, v_{\mathbf{x}_m}^0))$; and for all $m > 0$:

$$RS_{m+1} \triangleq \begin{cases} RS_m \cup \{\mathbf{x}_m\} & \text{if } O_m = R(\mathbf{x}_m, -) \\ RS_m & \text{otherwise} \end{cases}$$

and

$$wr(\mathbf{x}_m, v_m, RS_{m-1}, WS_{m-1}) \triangleq \begin{cases} R(\mathbf{s}[\mathbf{x}_m], (\perp, \perp)) & \text{if } \mathbf{x}_m \notin RS_{m-1} \cup WS_{m-1} \\ \xrightarrow{\text{pol}_{\text{imm}}} fs_m \\ \xrightarrow{\text{pol}_{\text{imm}}} rl_{\mathbf{x}_m} \emptyset & \text{otherwise} \end{cases}$$

$wws_{\mathbf{x}_m} \triangleq W(\mathbf{ws}, WS_m)$; $lw_{\mathbf{x}_m} \triangleq W(\mathbf{s}[\mathbf{x}_m], (v'_m, v_m))$; fs_m and $rl_{\mathbf{x}_m}$ are as defined above; and for all $m > 0$:

$$WS_{m+1} \triangleq \begin{cases} WS_m \cup \{\mathbf{x}_m\} & \text{if } O_m = W(\mathbf{x}_m, -) \\ WS_m & \text{otherwise} \end{cases}$$

$$ws_{m+1} \triangleq \begin{cases} ws_m + +[(\mathbf{x}_m, v_m)] & \text{if } O_m = W(\mathbf{x}_m, v_m) \\ ws_m & \text{otherwise} \end{cases}$$

Let $RS_\xi = RS_m$, $WS_\xi = WS_m$, and $ws_\xi = ws_m$; let $RS_\xi \cup WS_\xi$ be enumerated as $\{\mathbf{x}_1 \dots \mathbf{x}_i\}$ for some i , and ws_ξ be enumerated as $\{(\mathbf{x}_1, v_1) \dots (\mathbf{x}_j, v_j)\}$ for some j .

- VRs denotes the sequence of events validating the reads and is of the form $v_{x_1} \xrightarrow{\text{po}|_{\text{imm}}} \dots \xrightarrow{\text{po}|_{\text{imm}}} v_{x_i}$, where for all $n \in \{1 \dots i\}$:

$$v_{x_n} = \begin{cases} \mathbf{R}(\mathbf{s}[\mathbf{x}_n], (v_{x_n}^0, -)) \xrightarrow{\text{po}|_{\text{imm}}} vr_{x_n} = \mathbf{R}(\mathbf{x}_n, v_{x_n}^0) & \text{if } \mathbf{x}_n \in \mathbf{RS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- RUs denotes the sequence of events releasing the reader locks (when the given location is in the read set only) and is of the form $ru_{x_1} \xrightarrow{\text{po}|_{\text{imm}}} \dots \xrightarrow{\text{po}|_{\text{imm}}} ru_{x_i}$, where for all $n \in \{1 \dots i\}$:

$$ru_{x_n} = \begin{cases} \mathbf{RU}(\mathbf{x}1_n) & \text{if } \mathbf{x}_n \notin \mathbf{WS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- PLs denotes the sequence of events promoting the reader locks to writer ones (when the given location is in the write set), and is of the form $pl_{x_1} \xrightarrow{\text{po}|_{\text{imm}}} \dots \xrightarrow{\text{po}|_{\text{imm}}} pl_{x_i}$, where for all $n \in \{1 \dots i\}$:

$$pl_{x_n} = \begin{cases} \mathbf{PL}(\mathbf{x}1_n) & \text{if } \mathbf{x}_n \in \mathbf{WS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

- Ws denotes the sequence of events committing the writes of (\mathbb{T}) and is of the form $c_{x_1, v_1} \xrightarrow{\text{po}|_{\text{imm}}} \dots \xrightarrow{\text{po}|_{\text{imm}}} c_{x_j, v_j}$, where for all $n \in \{1 \dots j\}$: $c_{x_n, v_n} = \mathbf{W}(\mathbf{x}_n, v_n)$
- Us denotes the sequence of events releasing the locks on the write set, and is of the form $wu_{x_1} \xrightarrow{\text{po}|_{\text{imm}}} \dots \xrightarrow{\text{po}|_{\text{imm}}} wu_{x_i}$, where for all $n \in \{1 \dots i\}$:

$$wu_{x_n} = \begin{cases} \mathbf{WU}(\mathbf{x}1_n) & \text{if } \mathbf{x}_n \in \mathbf{WS}_\xi \\ \emptyset & \text{otherwise} \end{cases}$$

Given a transaction trace θ_ξ , we write e.g. $\xi.Us$ to refer to its constituent Us sub-trace and write $Us.E$ for the set of events related by po in Us . Similarly, we write $\xi.E$ for the set of events related by po in θ_ξ . Note that $G.E = \bigcup_{\xi \in \mathbf{TX}} \xi.E$.

Note that for each transaction ξ and each location \mathbf{x} , the $\xi.rl_{\mathbf{x}}$, $\xi.rs_{\mathbf{x}}$, $\xi.wr_{\mathbf{x}}$, $\xi.ru_{\mathbf{x}}$, $\xi.pl_{\mathbf{x}}$ and $\xi.wu_{\mathbf{x}}$ are uniquely identified when they exist. Let $\xi.w_{\mathbf{x}}$ the last (in po order) write to \mathbf{x} in Ws , when it exists.

Note that for each write operation w in θ'_ξ , there exists a *matching* write operation in $\theta_\xi.Ws$, denoted by $\mathbf{mw}(w)$ is a one-to-one function such that: That is,

$$\mathbf{mw}(w) = w' \stackrel{\text{def}}{\iff} \exists i. \wedge (\theta'_\xi.E \cap \mathcal{W})|_i = w \wedge (\theta_\xi.Ws)|_i = w'$$

We then define:

$$\text{RF} \triangleq \left\{ (w, r) \left| \begin{array}{l} r \in G'.\mathcal{NT} \\ \wedge ((w \in G.\mathcal{NT} \wedge (w, r) \in G'.\text{rf}) \\ \vee (\exists \xi'. \exists w' \in \xi'. E \cap \mathcal{W}. (w', r) \in G'.\text{rf} \wedge w = \text{mw}(w'))) \end{array} \right. \right\} \\ \cup \left\{ \begin{array}{l} (w, \xi.\text{rs}_x), \\ (w, \xi.\text{vr}_x) \end{array} \left| \begin{array}{l} (w \in G.\mathcal{NT} \wedge (w, \xi.\text{rs}_x) \in G'.\text{rf}) \\ \vee (\exists \xi'. \exists w' \in \xi'. E \cap \mathcal{W}. (w', \xi.\text{rs}_x) \in G'.\text{rf} \wedge w = \text{mw}(w')) \end{array} \right. \right\}$$

Similarly, we define:

$$\text{MO} \triangleq \left\{ (w_1, w_2) \mid \exists w'_1, w'_2. w_1 = \text{mw}(w'_1) \wedge w_2 = \text{mw}(w'_2) \wedge (w'_1, w'_2) \in G'.\text{mo} \right\} \\ \cup \left\{ (w, w') \mid w, w' \in G.\mathcal{NT} \wedge (w, w') \in G'.\text{mo} \right\} \\ \cup \left\{ (w, w') \mid w' \in G.\mathcal{NT} \wedge \exists w''. w = \text{mw}(w'') \wedge (w'', w') \in G'.\text{mo} \right\} \\ \cup \left\{ (w', w) \mid w' \in G.\mathcal{NT} \wedge \exists w''. w = \text{mw}(w'') \wedge (w', w'') \in G'.\text{mo} \right\}$$

For each location x we then define:

$$\text{LO}_x \triangleq \left\{ (\xi.\text{rl}_x, \xi.\text{pl}_x), (\xi.\text{rl}_x, \xi.\text{wu}_x), (\xi.\text{pl}_x, \xi.\text{wu}_x) \mid G'.\mathcal{T}_\xi \cap \mathcal{W}_x \neq \emptyset \right\} \\ \cup \left\{ \begin{array}{l} (\xi.\text{rl}_x, \xi'.\text{pl}_x), (\xi.\text{rl}_x, \xi'.\text{wu}_x), \\ (\xi.\text{ru}_x, \xi'.\text{pl}_x), (\xi.\text{ru}_x, \xi'.\text{wu}_x) \end{array} \left| \begin{array}{l} \xi \neq \xi' \wedge \exists a, b, x. \\ a \in G'.\mathcal{T}_\xi \wedge b \in G'.\mathcal{T}'_{\xi'} \\ \wedge \text{loc}(a) = \text{loc}(b) = x \\ \wedge a \in G'.\mathcal{R}_E \wedge (a, b) \in G'.\text{rb} \end{array} \right. \right\} \\ \cup \left\{ \begin{array}{l} (\xi.\text{rl}_x, \xi'.\text{pl}_x), (\xi.\text{rl}_x, \xi'.\text{wu}_x), \\ (\xi.\text{pl}_x, \xi'.\text{rl}_x), (\xi.\text{pl}_x, \xi'.\text{pl}_x), (\xi.\text{pl}_x, \xi'.\text{wu}_x), \\ (\xi.\text{wu}_x, \xi'.\text{rl}_x), (\xi.\text{wu}_x, \xi'.\text{pl}_x), (\xi.\text{wu}_x, \xi'.\text{wu}_x) \end{array} \left| \begin{array}{l} \xi \neq \xi' \wedge \exists a, b, x. \\ a \in G'.\mathcal{T}_\xi \wedge b \in G'.\mathcal{T}'_{\xi'} \\ \wedge \text{loc}(a) = \text{loc}(b) = x \\ \wedge (a, b) \in G'.\text{mo} \end{array} \right. \right\} \\ \cup \left\{ \begin{array}{l} (\xi.\text{pl}_x, \xi'.\text{rl}_x), (\xi.\text{pl}_x, \xi'.\text{ru}_x), \\ (\xi.\text{wu}_x, \xi'.\text{rl}_x), (\xi.\text{wu}_x, \xi'.\text{ru}_x) \end{array} \left| \begin{array}{l} \xi \neq \xi' \wedge \exists a, b, x. \\ a \in G'.\mathcal{T}_\xi \wedge b \in G'.\mathcal{T}'_{\xi'} \\ \wedge G'.\mathcal{T}'_{\xi'} \cap \mathcal{W}_x = \emptyset \\ \wedge \text{loc}(a) = \text{loc}(b) = x \\ \wedge (a, b) \in G'.(\text{mo}^?; \text{rf}) \end{array} \right. \right\}$$

Note that each LO_x satisfies the conditions in [Def. 4](#).

We are now in a position to demonstrate the completeness of our implementation. Given an RSI-consistent execution graph G' , we construct an execution graph G of the implementation as follows and demonstrate that $\text{RA-consistent}(G)$ holds.

- $G.E = \bigcup_{\mathcal{T}_\xi \in G'.\mathcal{T}/\text{st}} \theta_\xi.E \cup G'.\mathcal{NT}$. Observe that $G'.E \subseteq G.E$.
- $G.\text{po}$ is defined as $G'.\text{po}$ extended by the po for the additional events of G , given by each θ_ξ trace defined above. Note that $G.\text{po}$ does not introduce additional orderings between events of $G'.E$. That is, $\forall a, b \in G'.E. (a, b) \in G'.\text{po} \Leftrightarrow (a, b) \in G.\text{po}$.
- $G.\text{rf} = \text{RF}$.
- $G.\text{mo} = \text{MO}$.
- $G.\text{lo} = \bigcup_{x \in \text{Locs}} \text{LO}_x$, with LO_x as defined above.

Notation In what follows, given an RSI implementation graph G as constructed above we write $G.\mathcal{NT}$ for the non-transactional events of G , i.e. $G.\mathcal{NT} \triangleq \{a \mid \neg \exists \xi. a \in G.E_\xi\}$. Moreover, as before, given a relation $r \subseteq G.E \times G.E$, we override the r_I notation and write r_I for $\{(a, b) \in r \mid \exists \xi. a, b \in \theta_\xi.E\}$.

Lemma 19. *Given an RSI-consistent execution graph G' and its corresponding implementation graph G constructed as above, for all a, b, ξ_a, ξ_b :*

$$\begin{aligned}
& (a, b) \in G.\mathbf{hb} \Rightarrow \\
& \quad (\exists \xi. a, b \in G.E_\xi \Rightarrow (a, b) \in G.\mathbf{po}) \\
& \quad \wedge \left(\neg \exists \xi. a, b \in G.E_\xi \Rightarrow \right. \\
& \quad \quad \exists A, B. \emptyset \subset A \times B \subseteq G'.\mathbf{rsi-hb} \\
& \quad \quad \wedge (a \in G.\mathcal{NT} \Rightarrow A = \{a\}) \wedge (b \in G.\mathcal{NT} \Rightarrow B = \{b\}) \\
& \quad \quad \wedge [a \in G.E_{\xi_a} \Rightarrow \\
& \quad \quad \quad A = G'.\mathcal{T}_{\xi_a} \vee (\mathbf{stg}(a, \xi_a) \leq 2 \wedge A = G'.\mathcal{T}_{\xi_a} \cap \mathcal{R}_E) \\
& \quad \quad \quad \vee (\mathbf{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\mathbf{po}^?} d \wedge A = \{\mathbf{mw}^{-1}(d)\})] \\
& \quad \quad \wedge [b \in G.E_{\xi_b} \Rightarrow \\
& \quad \quad \quad B = G'.\mathcal{T}_{\xi_b} \vee (\mathbf{stg}(b, \xi_b) \geq 3 \wedge B = G'.\mathcal{T}_{\xi_b} \cap \mathcal{W})] \quad \left. \right)
\end{aligned}$$

where

$$\mathbf{stg}(a, \xi_a) \triangleq \begin{cases} 1 & \text{if } a \notin \xi_a.RU_{\xi_a} \cup \xi_a.PL_{\xi_a} \cup \xi_a.W_{\xi_a} \cup \xi_a.U_{\xi_a} \\ 2 & \text{if } a \in \xi_a.RU_{\xi_a} \\ 3 & \text{if } a \in \xi_a.PL_{\xi_a} \\ 4 & \text{otherwise} \end{cases}$$

Proof. Pick an arbitrary RSI-consistent execution graph G' and its corresponding implementation graph G constructed as above. Let $\mathbf{hb}_0 \triangleq G.(\mathbf{po} \cup \mathbf{rf} \cup \mathbf{lo})$ and $\mathbf{hb}_{n+1} \triangleq \mathbf{hb}_0; \mathbf{hb}_n$, for all $n \in \mathbb{N}$. It is then straightforward to demonstrate that $G.\mathbf{hb} = \bigcup_{i \in \mathbb{N}} \mathbf{hb}_i$. We thus demonstrate instead that:

$$\begin{aligned}
& \forall n \in \mathbb{N}. \forall a, b, \xi_a, \xi_b. (a, b) \in G.\mathbf{hb}_n \Rightarrow \\
& \quad (\exists \xi. a, b \in G.E_\xi \Rightarrow (a, b) \in G.\mathbf{po}) \\
& \quad \wedge \left(\neg \exists \xi. a, b \in G.E_\xi \Rightarrow \right. \\
& \quad \quad \exists A, B. \emptyset \subset A \times B \subseteq G'.\mathbf{rsi-hb} \\
& \quad \quad \wedge (a \in G.\mathcal{NT} \Rightarrow A = \{a\}) \wedge (b \in G.\mathcal{NT} \Rightarrow B = \{b\}) \\
& \quad \quad \wedge [a \in G.E_{\xi_a} \Rightarrow \\
& \quad \quad \quad A = G'.\mathcal{T}_{\xi_a} \vee (\mathbf{stg}(a, \xi_a) \leq 2 \wedge A = G'.\mathcal{T}_{\xi_a} \cap \mathcal{R}_E) \\
& \quad \quad \quad \vee (\mathbf{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\mathbf{po}^?} d \wedge A = \{\mathbf{mw}^{-1}(d)\})] \\
& \quad \quad \wedge [b \in G.E_{\xi_b} \Rightarrow \\
& \quad \quad \quad B = G'.\mathcal{T}_{\xi_b} \vee (\mathbf{stg}(b, \xi_b) \geq 3 \wedge B = G'.\mathcal{T}_{\xi_b} \cap \mathcal{W})] \quad \left. \right)
\end{aligned}$$

We proceed by induction on n .

Base case $n = 0$

There are three cases to consider: 1) $(a, b) \in G.\text{po}$; or 2) $(a, b) \in G.\text{rf}$; or 3) $(a, b) \in G.\text{lo}$.

In case (1) there are five cases to consider: a) $\exists \xi. (a, b) \in G.E_\xi$; or b) $a, b \in G.\mathcal{NT}$; or c) $a \in G.\mathcal{NT}$ and $b \in G.E_{\xi_b}$; or d) $a \in G.E_{\xi_a}$ and $b \in G.\mathcal{NT}$; or e) $a \in G.E_{\xi_a}$, $b \in G.E_{\xi_b}$ and $\xi_a \neq \xi_b$. In case (1.a) we then have $(a, b) \in G.\text{po}_1$, as required. In case (1.b) from the construction of $G.\text{po}$ we have $(a, b) \in G'.\text{po}$, as required. In case (1.c) from the construction of $G.\text{po}$ we have $(\{a\} \times G'.\mathcal{T}_{\xi_b}) \in G'.\text{po}$, as required. In case (1.d) from the construction of $G.\text{po}$ we have $(G'.\mathcal{T}_{\xi_a} \times \{b\}) \in G'.\text{po}$, as required. In case (1.e) from the construction of $G.\text{po}$ we have $(G'.\mathcal{T}_{\xi_a} \times G'.\mathcal{T}_{\xi_b}) \in G'.\text{po}$, as required.

In case (2) there are five cases to consider: a) $\exists \xi. (a, b) \in G.E_\xi$; or b) $a, b \in G.\mathcal{NT}$; or c) $a \in G.\mathcal{NT}$ and $b \in G.E_{\xi_b}$; or d) $a \in G.E_{\xi_a}$ and $b \in G.\mathcal{NT}$; or e) $a \in G.E_{\xi_a}$, $b \in G.E_{\xi_b}$ and $\xi_a \neq \xi_b$. Case (2.a) holds vacuously as $(a, b) \in G.\text{rf}_1 = \emptyset$. In case (2.b) from the construction of $G.\text{rf}$ we have $(a, b) \in G'.\text{rf}$, as required. In case (2.c) from the construction of $G.\text{rf}$ we know that there exists $r \in G'.\mathcal{T}_{\xi_b}$ such that $(a, r) \in G'.\text{rf}$. As such we have $(\{w\} \times G'.\mathcal{T}_{\xi_b}) \subseteq G'.([\mathcal{NT}]; \text{rf}; \text{st}) \subseteq G'.\text{rsi-hb}$, as required. In case (2.d) from the construction of $G.\text{rf}$ we then have $(\text{mw}^{-1}(a), b) \in G'.\text{rf}$; that is a $G \xrightarrow{\text{po}^?}$ a and $(\text{mw}^{-1}(a), b) \in G'.\text{rf}$, as required. In case (2.e) from the construction of $G.\text{rf}$ we know that there exists $w \in G'.\mathcal{T}_{\xi_a}$, $r \in G'.\mathcal{T}_{\xi_b}$ such that $(w, r) \in G'.\text{rf}$. As such we have $(G'.\mathcal{T}_{\xi_a} \times G'.\mathcal{T}_{\xi_b}) \subseteq G'.\text{rf}_\top \subseteq G'.\text{rsi-hb}$, as required.

In case (3) there are five cases to consider: a) $\exists \xi. (a, b) \in G.E_\xi$; or b) $a, b \in G.\mathcal{NT}$; or c) $a \in G.\mathcal{NT}$ and $b \in G.E_{\xi_b}$; or d) $a \in G.E_{\xi_a}$ and $b \in G.\mathcal{NT}$; or e) $a \in G.E_{\xi_a}$, $b \in G.E_{\xi_b}$ and $\xi_a \neq \xi_b$. In case (3.a) from the construction of lo we have $(a, b) \in G.\text{po}$, as required. Cases (3.b-3.d) hold vacuously as there are no lo edge to or from non-transactional events. In case (3.e) from the construction of lo we know there exist x, c, d such that $c \in G'.\mathcal{T}_{\xi_a}$, $d \in G'.\mathcal{T}_{\xi_b}$, and either i) $(c, d) \in G'.\text{rf}$; or ii) $(c, d) \in G'.(\text{mo}; \text{rf})$; or iii) $(c, d) \in G'.\text{mo}$; or iv) $\text{loc}(c) = \text{loc}(d) = x$, $a = \xi_a.\text{rl}_x \vee a = \xi_a.\text{ru}_x$, $b = \xi_a.\text{pl}_x \vee b = \xi_a.\text{wu}_x$ and $c \in G'.\mathcal{R}_E$ and $(c, d) \in G'.\text{rb}$. In cases (3.e.i, 3.e.ii) we have $G'.\mathcal{T}_{\xi_a} \times G'.\mathcal{T}_{\xi_b} \subseteq G'.(\text{rf}_\top \cup (\text{mo}; \text{rf})_\top) \subseteq \text{rsi-rf} \subseteq \text{rsi-hb}$, as required. In case (3.e.iii) we have $G'.\mathcal{T}_{\xi_a} \times G'.\mathcal{T}_{\xi_b} \subseteq G'.\text{mo}_\top \subseteq \text{rsi-hb}$, as required. In case (3.e.iv) we then have $\text{stg}(a, \xi_a) \leq 2$, $\text{stg}(b, \xi_b) \geq 3$, and $G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E) \times G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}) \subseteq G'.\text{si-rb} \subseteq \text{rsi-hb}$, as required.

Inductive case $n = m+1$

$$\begin{aligned}
& \forall i \in \mathbb{N}. \forall a, b, \xi_a, \xi_b. i \leq m \wedge (a, b) \in G.\mathbf{hb}_i \Rightarrow \\
& \quad (\exists \xi. a, b \in G.E_\xi \Rightarrow (a, b) \in G.\mathbf{po}) \\
& \quad \wedge \left(\neg \exists \xi. a, b \in G.E_\xi \Rightarrow \right. \\
& \quad \quad \exists A, B. \emptyset \subset A \times B \subseteq G'.\mathbf{rsi-hb} \\
& \quad \quad \wedge (a \in G.\mathcal{NT} \Rightarrow A = \{a\}) \wedge (b \in G.\mathcal{NT} \Rightarrow B = \{b\}) \\
& \quad \quad \wedge [a \in G.E_{\xi_a} \Rightarrow \\
& \quad \quad \quad A = G'.\mathcal{T}_{\xi_a} \vee (\mathbf{stg}(a, \xi_a) \leq 2 \wedge A = G'.\mathcal{T}_{\xi_a} \cap \mathcal{R}_E) \\
& \quad \quad \quad \vee (\mathbf{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\mathbf{po}^?} d \wedge A = \{\mathbf{mw}^{-1}(d)\})] \\
& \quad \quad \wedge [b \in G.E_{\xi_b} \Rightarrow \\
& \quad \quad \quad B = G'.\mathcal{T}_{\xi_b} \vee (\mathbf{stg}(b, \xi_b) \geq 3 \wedge B = G'.\mathcal{T}_{\xi_b} \cap \mathcal{W})] \quad \left. \right) \\
& \hspace{15em} \text{(I.H.)}
\end{aligned}$$

Since $(a, b) \in \mathbf{hb}_n$, from the definition of \mathbf{hb}_n we know there exists c such that $(a, c) \in \mathbf{hb}_0$ and $(c, b) \in \mathbf{hb}_m$. There are then five cases to consider: 1) $\exists \xi. (a, b) \in G.E_\xi$; or 2) $a, b \in G.\mathcal{NT}$; or 3) $a \in G.\mathcal{NT}$ and $b \in G.E_{\xi_b}$; or 4) $a \in G.E_{\xi_a}$ and $b \in G.\mathcal{NT}$; or 5) $a \in G.E_{\xi_a}$, $b \in G.E_{\xi_b}$ and $\xi_a \neq \xi_b$.

Case 1

In case (1) pick arbitrary ξ such that $a, b \in G.E_\xi$. There are then three additional cases to consider: a) $c \in G.\mathcal{NT}$; b) $c \in G.E_\xi$; or c) there exists $\xi' \neq \xi$ such that $c \in G.E_{\xi'}$.

In case (1.a) from the proof of the base case and the (I.H.) we know there exist $A, B, C \neq \emptyset$ such that $C = \{c\}$, $A \times C \subseteq G'.\mathbf{rsi-hb}$, $C \times B \subseteq G'.\mathbf{rsi-hb}$ and thus $A \times B \subseteq G'.\mathbf{rsi-hb}$ and either: i) $A = G'.\mathcal{T}_\xi$ and $B = G'.\mathcal{T}_\xi$; or ii) $A = G'.\mathcal{T}_\xi$ and $B = (G'.(\mathcal{T}_\xi \cap \mathcal{W}))$; or iii) $\mathbf{stg}(a, \xi) \leq 2$, $A = G'.(\mathcal{T}_\xi \cap \mathcal{R}_E)$ and $B = G'.\mathcal{T}_\xi$; or iv) $\mathbf{stg}(a, \xi) \leq 2$, $A = G'.(\mathcal{T}_\xi \cap \mathcal{R}_E)$, $\mathbf{stg}(b, \xi) \geq 3$ and $B = G'.(\mathcal{T}_\xi \cap \mathcal{W})$; or v) $\mathbf{stg}(a, \xi) \leq 4$, $\exists d \in \xi.\mathcal{W}. a \xrightarrow{G.\mathbf{po}^?} d \wedge A = \{\mathbf{mw}^{-1}(d)\}$ and $B = G'.\mathcal{T}_\xi$; or vi) $\mathbf{stg}(a, \xi) \leq 4$, $\exists d \in \xi.\mathcal{W}. a \xrightarrow{G.\mathbf{po}^?} d \wedge A = \{\mathbf{mw}^{-1}(d)\}$ and $B = G'.(\mathcal{T}_\xi \cap \mathcal{W})$.

Case (i) cannot arise as we would have $A = B$ and thus $A \times A \subseteq G'.\mathbf{rsi-hb}$, contradicting the assumption that G' is RSI-consistent. Case (ii) cannot arise as we would have $(G'.(\mathcal{T}_\xi \cap \mathcal{W})) \times (G'.(\mathcal{T}_\xi \cap \mathcal{W})) \subseteq G'.\mathbf{rsi-hb}$, contradicting the assumption that G' is RSI-consistent. Case (iii) cannot arise as we would have $(G'.(\mathcal{T}_\xi \cap \mathcal{R}_E)) \times (G'.(\mathcal{T}_\xi \cap \mathcal{R}_E)) \subseteq G'.\mathbf{rsi-hb}$, contradicting the assumption that G' is RSI-consistent. In case (iv) since we have $\mathbf{stg}(a, \xi) \leq 2$ and $\mathbf{stg}(b, \xi) \geq 3$, from the definition of $\mathbf{stg}(\cdot, \cdot)$ and the construction of G we have $(a, b) \in G.\mathbf{po}$, as required. Cases (v-vi) cannot arise as we would have $\exists d \in \xi.\mathcal{W}. (\mathbf{mw}^{-1}(d), \mathbf{mw}^{-1}(d)) \in G'.\mathbf{rsi-hb}$, contradicting the assumption that G' is RSI-consistent.

In case (1.b) from the proof of the base case we have $(a, c) \in G.\mathbf{po}$. On the other hand from (I.H.) we have $(c, b) \in G.\mathbf{po}$. As $G.\mathbf{po}$ is transitively closed, we

have $(a, b) \in G.\text{po}$, as required.

In case (1.c) from the proof of the base case (in cases 1.e, 2.e and 3.e) we have either A) $G'.\mathcal{T}_\xi \times G'.\mathcal{T}_{\xi'} \subseteq G'.\text{rsi-hb}$; or B) $\text{stg}(a, \xi) \leq 2$, $\text{stg}(c, \xi') \geq 3$, and $\emptyset \subset G'.(\mathcal{T}_\xi \cap \mathcal{R}_E) \times G'.(\mathcal{T}_{\xi'} \cap \mathcal{W}) \subseteq G'.\text{rsi-hb}$. On the other hand from (I.H.) we know there exist $C, B \neq \emptyset$ such that $C \times B \subseteq G'.\text{rsi-hb}$ and either: i) $C = G'.\mathcal{T}_{\xi'}$ and $B = G'.\mathcal{T}_\xi$; or ii) $C = G'.\mathcal{T}_{\xi'}$, $\text{stg}(b, \xi) \geq 3$ and $B = G'.(\mathcal{T}_\xi \cap \mathcal{W})$; or iii) $\text{stg}(c, \xi') \leq 2$, $C = G'.(\mathcal{T}_{\xi'} \cap \mathcal{R}_E)$ and $B = G'.\mathcal{T}_\xi$; or iv) $\text{stg}(c, \xi') \leq 2$, $C = G'.(\mathcal{T}_{\xi'} \cap \mathcal{R}_E)$, $\text{stg}(b, \xi) \geq 3$ and $B = G'.(\mathcal{T}_\xi \cap \mathcal{W})$; or v) $\text{stg}(c, \xi') \leq 4$, $\exists d \in \xi'.\mathcal{W}$. $c \xrightarrow{G.\text{po}^?} d \wedge C = \{\text{mw}^{-1}(d)\}$ and $B = G'.\mathcal{T}_\xi$; or vi) $\text{stg}(c, \xi') \leq 4$, $\exists d \in \xi'.\mathcal{W}$. $c \xrightarrow{G.\text{po}^?} d \wedge C = \{\text{mw}^{-1}(d)\}$, $\text{stg}(b, \xi) \geq 3$ and $B = G'.(\mathcal{T}_\xi \cap \mathcal{W})$.

Cases (A.i-A.vi) lead to a cycle in $G'.\text{rsi-hb}$, contradicting the assumption that G' is RSI-consistent. In cases (B.ii, B.iv, B.vi) we then have $\text{stg}(a, \xi) \leq 2$ and $\text{stg}(b, \xi) \geq 3$. Consequently from the definition of $\text{stg}(\cdot, \cdot)$ and the construction of G we have $(a, b) \in G.\text{po}$, as required. Case (B.i) leads to a cycle in $G'.\text{rsi-hb}$, contradicting the assumption that G' is RSI-consistent. In cases (B.iii) we have $\text{stg}(c, \xi') \geq 3$ and $\text{stg}(c, \xi') \leq 2$, leading to a contradiction. In case (B.v) we then know $\exists d \in \xi'.\mathcal{W}$. $\emptyset \subset G'.(\mathcal{T}_\xi \cap \mathcal{R}_E) \times \{\text{mw}^{-1}(d)\} \subseteq G'.\text{rsi-hb} \wedge \{\text{mw}^{-1}(d)\} \times G'.\mathcal{T}_\xi \subseteq G'.\text{rsi-hb}$. That is, we have $\emptyset \subset G'.(\mathcal{T}_\xi \cap \mathcal{R}_E) \times G'.(\mathcal{T}_\xi \cap \mathcal{R}_E) \subseteq G'.\text{rsi-hb}$, contradicting the assumption that G' is RSI-consistent.

Case 2

There are two additional cases to consider: a) $c \in G.\mathcal{NT}$; or b) there exists ξ such that $c \in G.E_\xi$.

In case (2.a) from the proof of the base case we have $(a, c) \in G'.\text{rsi-hb}$. On the other hand from (I.H.) we have $(c, b) \in G'.\text{rsi-hb}$. As $G'.\text{rsi-hb}$ is transitively closed, we have $(a, b) \in G'.\text{rsi-hb}$, as required.

In case (2.b) from the proof of the base case we know there exists $C_1 \neq \emptyset$ such that $\{a\} \times C_1 \subseteq G'.\text{rsi-hb}$ and either: A) $C_1 = G'.\mathcal{T}_\xi$; or B) $\text{stg}(c, \xi) \geq 3$ and $C_1 = G'.(\mathcal{T}_\xi \cap \mathcal{W})$. On the other hand, from (I.H.) we know there exists $C_2 \neq \emptyset$ such that $C_2 \times \{b\} \in G'.\text{rsi-hb}$ and either: i) $C_2 = G'.\mathcal{T}_\xi$; or ii) $\text{stg}(c, \xi) \leq 2$ and $C_2 = G'.(\mathcal{T}_\xi \cap \mathcal{R}_E)$; or iii) $\text{stg}(c, \xi) \leq 4$ and $\exists d \in \xi.\mathcal{W}$. $c \xrightarrow{G.\text{po}^?} d \wedge C_2 = \{\text{mw}^{-1}(d)\}$.

In cases (A.i-A.iii) and (B.i, B.iii) from the transitivity of $G'.\text{rsi-hb}$ we have $(a, b) \in G'.\text{rsi-hb}$, as required. Case (B.ii) cannot arise as otherwise we would have $3 \leq \text{stg}(c, \xi) \leq 2$, leading to a contradiction.

Case 3

There are three additional cases to consider: a) $c \in G.\mathcal{NT}$; b) $c \in G.E_{\xi_b}$; or c) there exists $\xi_c \neq \xi_b$ such that $c \in G.E_{\xi_c}$.

In case (3.a) from the proof of the base case we have $(a, c) \in G'.\text{rsi-hb}$. On the other hand from (I.H.) we know there exists $B \neq \emptyset$ such that $\{c\} \times B \subseteq G'.\text{rsi-hb}$ and $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$. As $G'.\text{rsi-hb}$ is transitively closed we then know there exists $B \neq \emptyset$ such that $\{a\} \times B \subseteq G'.\text{rsi-hb}$ and $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$, as required.

In case (3.b) from the proof of the base case we know there exists $B \neq \emptyset$ such that $\{a\} \times B \subseteq G'.\text{rsi-hb}$ and $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(c, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$. On the other hand, from (I.H.) we have $c \xrightarrow{G.\text{po}} b$ and thus from the definition of $\text{stg}(\cdot, \cdot)$ and the construction of G we have $\text{stg}(b, \xi_b) \geq \text{stg}(c, \xi_b)$. As such, we know there exists $B \neq \emptyset$ such that $\{a\} \times B \subseteq G'.\text{rsi-hb}$ and $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$, as required.

In case (3.c) from the proof of the base case we know there exists $C_1 \neq \emptyset$ such that $\{a\} \times C_1 \subseteq G'.\text{rsi-hb}$ and either A) $C_1 = G'.\mathcal{T}_{\xi_c}$; or B) $(\text{stg}(c, \xi_c) \geq 3 \wedge C_1 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{W}))$. On the other hand, from (I.H.) we know there exist $C_2, B \neq \emptyset$ such that $C_2 \times B \in G'.\text{rsi-hb}$ and either: i) $C_2 = G'.\mathcal{T}_{\xi_c}$ and $B = G'.\mathcal{T}_{\xi_b}$; or ii) $C_2 = G'.\mathcal{T}_{\xi_c}$ and $(\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$; or iii) $\text{stg}(c, \xi_c) \leq 2 \wedge C_2 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{R}_E)$ and $B = G'.\mathcal{T}_{\xi_b}$; or iv) $\text{stg}(c, \xi_c) \leq 2 \wedge C_2 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{R}_E)$ and $(\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$; or v) $\text{stg}(c, \xi_c) \leq 4 \wedge \exists d \in \xi_c.\mathcal{W}. c \xrightarrow{G.\text{po}^?} d \wedge C_2 = \{\text{mw}^{-1}(d)\}$ and $B = G'.\mathcal{T}_{\xi_b}$; or vi) $\text{stg}(c, \xi_c) \leq 4 \wedge \exists d \in \xi_c.\mathcal{W}. c \xrightarrow{G.\text{po}^?} d \wedge C_2 = \{\text{mw}^{-1}(d)\}$ and $(\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$.

In cases (A.i-A.vi) and (B.i, B.ii, B.v, B.vi) from the transitivity of $G'.\text{rsi-hb}$ we know there exists $B \neq \emptyset$ such that $\{a\} \times B \subseteq G'.\text{rsi-hb}$ and $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$, as required. Cases (B.iii, B.iv) cannot arise as we would otherwise have $3 \leq \text{stg}(c, \xi_c) \leq 2$, leading to a contradiction.

Case 4

There are three additional cases to consider: a) $c \in G.\mathcal{NT}$; b) $c \in G.E_{\xi_a}$; or c) there exists $\xi_c \neq \xi_a$ such that $c \in G.E_{\xi_c}$.

In case (4.a) from (I.H.) we have $(c, b) \in G'.\text{rsi-hb}$. On the other hand from the base case we know there exists $A \neq \emptyset$ such that $A \times \{c\} \subseteq G'.\text{rsi-hb}$ and $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee \text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{\text{mw}^{-1}(d)\}$. As $G'.\text{rsi-hb}$ is transitively closed we then know we know there exists $A \neq \emptyset$ such that $A \times \{b\} \subseteq G'.\text{rsi-hb}$ and $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee \text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{\text{mw}^{-1}(d)\}$, as required.

In case (4.b) from (I.H.) we know there exists $A \neq \emptyset$ such that $A \times \{b\} \subseteq G'.\text{rsi-hb}$ and $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(c, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee \text{stg}(c, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{\text{mw}^{-1}(d)\}$. On the other hand, from the proof of the base case we have $a \xrightarrow{G.\text{po}} c$ and thus from the definition of $\text{stg}(\cdot, \cdot)$ and the construction of G we have $\text{stg}(a, \xi_a) \leq \text{stg}(c, \xi_a)$. As such, we know there exists $A \neq \emptyset$ such that $A \times \{b\} \subseteq G'.\text{rsi-hb}$ and $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee \text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{\text{mw}^{-1}(d)\}$, as required.

In case (4.c) from the proof of the base case we know there exist $A, C_1 \neq \emptyset$ such that $A \times C_1 \subseteq G'.\text{rsi-hb}$ and either: i) $A = G'.\mathcal{T}_{\xi_a}$ and $C_1 = G'.\mathcal{T}_{\xi_c}$; or ii) $A = G'.\mathcal{T}_{\xi_a}$ and $(\text{stg}(c, \xi_c) \geq 3 \wedge C_1 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{W}))$; or iii) $\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)$ and $C_1 = G'.\mathcal{T}_{\xi_c}$; or iv) $\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)$ and

($\text{stg}(c, \xi_c) \geq 3 \wedge C_1 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{W})$); or *v*) $\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{\text{mw}^{-1}(d)\}$ and $C_1 = G'.\mathcal{T}_{\xi_c}$; or *vi*) $\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{\text{mw}^{-1}(d)\}$ and ($\text{stg}(c, \xi_c) \geq 3 \wedge C_1 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{W})$).

On the other hand, from (I.H.) we know there exists $C_2 \neq \emptyset$ such $C_2 \times \{b\} \subseteq G'.\text{rsi-hb}$ and either *A*) $C_2 = G'.\mathcal{T}_{\xi_c}$; or *B*) ($\text{stg}(c, \xi_c) \leq 2 \wedge C_2 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{R}_E)$); or *C*) ($\text{stg}(c, \xi_c) \leq 4 \wedge \exists e \in \xi_c.\mathcal{W}. c \xrightarrow{G.\text{po}^?} e \wedge C_2 = \{\text{mw}^{-1}(e)\}$).

In cases (A.i-A.vi), (C.i-C.vi) and (B.i, B.iii, B.v) from the transitivity of $G'.\text{rsi-hb}$ we know there exists $A \neq \emptyset$ such that $A \times \{b\} \subseteq G'.\text{rsi-hb}$ and $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{d\})$, as required. Cases (B.ii, B.iv, B.vi) cannot arise as we would otherwise have $3 \leq \text{stg}(c, \xi_c) \leq 2$, leading to a contradiction.

Case 5

There are four additional cases to consider: *a*) $c \in G.\mathcal{NT}$; *b*) $c \in G.E_{\xi_a}$; or *c*) $c \in G.E_{\xi_b}$; or *d*) there exists ξ_c such that $\xi_c \neq \xi_a$, $\xi_c \neq \xi_b$ and $c \in G.E_{\xi_c}$.

In case (5.a) from the proof of the base case we know there exists $A \neq \emptyset$ such that $A \times \{c\} \subseteq G'.\text{rsi-hb}$ and $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{\text{mw}^{-1}(d)\})$. On the other hand from (I.H.) we know there exists $B \neq \emptyset$ such that $\{c\} \times B \subseteq G'.\text{rsi-hb}$ and $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$. As such, since $G'.\text{rsi-hb}$ is transitive we know there exist $A, B \neq \emptyset$ such that $A \times B \subseteq G'.\text{rsi-hb}$; that $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{\text{mw}^{-1}(d)\})$; and that $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$, as required.

In case (5.b) from (I.H.) we know there exist $A, B \neq \emptyset$ such that $A \times B \subseteq G'.\text{rsi-hb}$; that $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(c, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee (\text{stg}(c, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{\text{mw}^{-1}(d)\})$; and that $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$. On the other hand from the proof of the base case we have $(a, c) \in G'.\text{po}$ and thus from the definition of $\text{stg}(\cdot, \cdot)$ and the construction of G we have $\text{stg}(a, \xi_a) \leq \text{stg}(c, \xi_a)$. As such we know there exist $A, B \neq \emptyset$ such that $A \times B \subseteq G'.\text{rsi-hb}$; that $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{\text{mw}^{-1}(d)\})$ and that $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$, as required.

In case (5.c) from the proof of the base case we know there exist $A, B \neq \emptyset$ such that $A \times B \subseteq G'.\text{rsi-hb}$; that $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E)) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{\text{mw}^{-1}(d)\})$, and that $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(c, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$. On the other hand from (I.H.) we have $(c, b) \in G'.\text{po}$ and thus from the definition of $\text{stg}(\cdot, \cdot)$ and the construction of G we have $\text{stg}(c, \xi_a) \leq \text{stg}(b, \xi_a)$. As such we know there exist $A, B \neq \emptyset$ such that $A \times B \subseteq G'.\text{rsi-hb}$; that $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A =$

$G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{\text{mw}^{-1}(d)\})$ and that $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$, as required.

In case (5.d) from the proof of the base case we know there exist $A, C_1 \neq \emptyset$ such that $A \times C_1 \subseteq G'.\text{rsi-hb}$; that $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{\text{mw}^{-1}(d)\})$, and that either $A) C_1 = G'.\mathcal{T}_{\xi_c}$; or $B) (\text{stg}(c, \xi_c) \geq 3 \wedge C_1 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{W}))$.

On the other hand, from (I.H.) we know there exist $C_2, B \neq \emptyset$ such that $C_2 \times B \subseteq G'.\text{rsi-hb}$; that $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$; and that either: i) $C_2 = G'.\mathcal{T}_{\xi_c}$; or ii) $\text{stg}(c, \xi_c) \leq 2 \wedge C_2 = G'.(\mathcal{T}_{\xi_c} \cap \mathcal{R}_E)$; iii) $\text{stg}(c, \xi_c) \leq 4 \wedge \exists e \in \xi_c.\mathcal{W}. c \xrightarrow{G.\text{po}^?} e \wedge C_2 = \{\text{mw}^{-1}(e)\}$.

In cases (A.i-A.iii) and (B.i, B.iii) from the transitivity of $G'.\text{rsi-hb}$ we know there exists $A, B \neq \emptyset$ such that $A \times B \subseteq G'.\text{rsi-hb}$; that $A = G'.\mathcal{T}_{\xi_a} \vee (\text{stg}(a, \xi_a) \leq 2 \wedge A = G'.(\mathcal{T}_{\xi_a} \cap \mathcal{R}_E) \vee (\text{stg}(a, \xi_a) \leq 4 \wedge \exists d \in \xi_a.\mathcal{W}. a \xrightarrow{G.\text{po}^?} d \wedge A = \{\text{mw}^{-1}(d)\})$; and that $B = G'.\mathcal{T}_{\xi_b} \vee (\text{stg}(b, \xi_b) \geq 3 \wedge B = G'.(\mathcal{T}_{\xi_b} \cap \mathcal{W}))$, as required. Case (B.ii) cannot arise as we would otherwise have $3 \leq \text{stg}(c, \xi_c) \leq 2$, leading to a contradiction.

Theorem 14 (Completeness). For all RSI execution graphs G' and their counterpart implementation graphs G constructed as above,

$$\text{rsi-consistent}(G') \Rightarrow \text{RA-consistent}(G)$$

Proof. Pick an arbitrary RSI execution graph G' and its counterpart implementation graph G constructed as above and let us assume that $\text{rsi-consistent}(G')$ holds. From the definition of $\text{RA-consistent}(G)$ it then suffices to show:

1. $\text{irreflexive}(G.\text{hb})$
2. $\text{irreflexive}(G.\text{mo}; G.\text{hb})$
3. $\text{irreflexive}(G.\text{rb}; G.\text{hb})$

RTS. part 1

We proceed by contradiction. Let us assume that there exists a such that $(a, a) \in G.\text{hb}$. There are now two cases to consider: 1) $a \in G.\mathcal{NT}$; or 2) $\exists \xi. a \in G.E_\xi$. In case (1) from Lemma 19 we have $(a, a) \in G'.\text{rsi-hb}$, contradicting the assumption that G' is RSI-consistent. Similarly, in case (2) from Lemma 19 we have $(a, a) \in G.\text{po}$, leading to a contradiction as $G.\text{po}$ is acyclic by construction.

RTS. part 2

We proceed by contradiction. Let us assume that there exist a, b such that $(a, b) \in G.\text{hb}$ and $(b, a) \in G.\text{mo}$. Let $\text{loc}(a) = \text{loc}(b) = \mathbf{x}$. There are then five cases to consider: 1) $\exists \xi. (a, b) \in G.E_\xi$; or 2) $a, b \in G.\mathcal{NT}$; or 3) $a \in G.\mathcal{NT}$ and $b \in G.E_{\xi_b}$; or 4) $a \in G.E_{\xi_a}$ and $b \in G.\mathcal{NT}$; or 5) $a \in G.E_{\xi_a}$, $b \in G.E_{\xi_b}$ and $\xi_a \neq \xi_b$.

In case (1) from the construction of $G.\text{mo}$ we have $(\text{mw}^{-1}(b), \text{mw}^{-1}(a)) \in G'.\text{mo}$. Moreover, from Lemma 19 we have $(a, b) \in G.\text{po}$, and thus from the

construction of $G.\text{po}$ we have $(\text{mw}^{-1}(a), \text{mw}^{-1}(b)) \in G'.\text{po}$. Moreover, since $\text{mw}^{-1}(a), \text{mw}^{-1}(b)$ are write events in the same transaction ξ , we have $(\text{mw}^{-1}(a), \text{mw}^{-1}(b)) \in G'.\text{po}_I \cap \mathcal{W}^2 \subseteq G'.\text{rsi-hb}$. As such we have $\text{mw}^{-1}(a) \xrightarrow{G'.\text{rsi-hb}} \text{mw}^{-1}(b) \xrightarrow{G'.\text{mo}} \text{mw}^{-1}(a)$, contradicting the assumption that G' is RSI-consistent.

In case (2) from the construction of $G.\text{mo}$ we have $(b, a) \in G'.\text{mo}$. Moreover, from [Lemma 19](#) we have $(a, b) \in G'.\text{rsi-hb}$. As such we have $a \xrightarrow{G'.\text{rsi-hb}} b \xrightarrow{G'.\text{mo}} a$, contradicting the assumption that G' is RSI-consistent.

In case (3) from the construction of $G.\text{mo}$ we have $(\text{mw}^{-1}(b), a) \in G'.\text{mo}$. Moreover, since $b \in G.\mathcal{W}$ and thus $\text{mw}^{-1}(b) \in G'.\mathcal{W}$, from [Lemma 19](#) we have $(a, b) \in G'.\text{rsi-hb}$. As such we have $a \xrightarrow{G'.\text{rsi-hb}} \text{mw}^{-1}(b) \xrightarrow{G'.\text{mo}} a$, contradicting the assumption that G' is RSI-consistent.

In cases (4, 5) from the construction of G we know that $\text{mw}^{-1}(a) \xrightarrow{G.\text{po}} a \xrightarrow{G.\text{hb}} b$, and thus from the transitivity of $G.\text{hb}$ we have $(\text{mw}^{-1}(a), b) \in G.\text{hb}$. As $\text{stg}(\text{mw}^{-1}(a), \xi_a) = 1$, from [Lemma 19](#) we have $(\text{mw}^{-1}(a), b) \in G'.\text{rsi-hb}$. On the other hand, from the construction of $G.\text{mo}$ we have $(b, \text{mw}^{-1}(a)) \in G'.\text{mo}$. Consequently, we have $\text{mw}^{-1}(a) \xrightarrow{G'.\text{rsi-hb}} b \xrightarrow{G'.\text{mo}} \text{mw}^{-1}(a)$, contradicting the assumption that G' is RSI-consistent.

RTS. part 3

We proceed by contradiction. Let us assume that there exists a, b such that $(a, b) \in G.\text{hb}$ and $(b, a) \in G.\text{rb}$. Let $\text{loc}(a) = \text{loc}(b) = \mathbf{x}$. There are then five cases to consider: 1) $\exists \xi. (a, b) \in G.E_\xi$; or 2) $a, b \in G.\mathcal{NT}$; or 3) $a \in G.\mathcal{NT}$ and $b \in G.E_{\xi_b}$; or 4) $a \in G.E_{\xi_a}$ and $b \in G.\mathcal{NT}$; or 5) $a \in G.E_{\xi_a}$, $b \in G.E_{\xi_b}$ and $\xi_a \neq \xi_b$.

Case (1) cannot arise as from the definition of $G.\text{rb}$ we know a is a write event in $\xi.Ws$ while b is a read event in $\xi.Ts$ and no po edge exists between the events of $\xi.Ts$ and $\xi.Ws$.

In case (1) from [Lemma 19](#) we have $(a, b) \in G.\text{po}$. This however leads to a contradiction as from the definition of $G.\text{rb}$ we know a is a write event in $\xi.Ws$ while b is a read event in $\xi.Ts$ and no po edge exists between the events of $\xi.Ws$ and $\xi.Ts$.

In case (2) from the construction of $G.\text{rb}$ we know that $(a, b) \in G'.\text{rb}$. On the other hand, from [Lemma 19](#) we have $(a, b) \in G'.\text{rsi-hb}$. As such we have $a \xrightarrow{G'.\text{rsi-hb}} b \xrightarrow{G'.\text{rb}} a$, contradicting the assumption that G' is RSI-consistent.

In case (3), as b is a read event and thus $\text{stg}(b, \xi_b) = 1$, from [Lemma 19](#) we have $(a, b) \in G'.\text{rsi-hb}$. On the other hand, from the construction of $G.\text{rb}$ we have $(b, a) \in G'.\text{rb}$. Consequently, we have $a \xrightarrow{G'.\text{rsi-hb}} b \xrightarrow{G'.\text{rb}} a$, contradicting the assumption that G' is RSI-consistent.

In case (4) since a is a write, from the construction of G we know that $\text{mw}^{-1}(a) \xrightarrow{G.\text{po}} a \xrightarrow{G.\text{hb}} b$, and thus from the transitivity of $G.\text{hb}$ we have $(\text{mw}^{-1}(a), b) \in G.\text{hb}$. As $\text{stg}(\text{mw}^{-1}(a), \xi_a) = 1$, from [Lemma 19](#) we have $(\text{mw}^{-1}(a), b) \in G'.\text{rsi-hb}$. On the other hand, from the construction of $G.\text{rb}$ we have $(b, \text{mw}^{-1}(a)) \in G'.\text{rb}$.

Consequently, we have $\text{mw}^{-1}(a) \xrightarrow{G'.\text{rsi-hb}} b \xrightarrow{G'.\text{rb}} \text{mw}^{-1}(a)$, contradicting the assumption that G' is RSI-consistent.

Similarly, in case (5) since a is a write event, from the construction of G we know that $\text{mw}^{-1}(a) \xrightarrow{G.\text{po}} a \xrightarrow{G.\text{hb}} b$, and thus from the transitivity of $G.\text{hb}$ we have $(\text{mw}^{-1}(a), b) \in G.\text{hb}$. As $\text{stg}(\text{mw}^{-1}(a), \xi_a) = \text{stg}(b, \xi_b) = 1$, from [Lemma 19](#) we have $(\text{mw}^{-1}(a), b) \in G'.\text{rsi-hb}$. On the other hand, from the construction of $G.\text{rb}$ we have $(b, \text{mw}^{-1}(a)) \in G'.\text{rb}$. Consequently, we have $\text{mw}^{-1}(a) \xrightarrow{G'.\text{rsi-hb}} b \xrightarrow{G'.\text{rb}} \text{mw}^{-1}(a)$, contradicting the assumption that G' is RSI-consistent.