On Parallel Snapshot Isolation and Release/Acquire Consistency
(Technical Appendix)

Azalea Raad[1], Ori Lahav[2], and Viktor Vafeiadis[1]

[1] MPI-SWS
[2] Tel Aviv University

1 Parallel Snapshot Isolation (PSI)

Lemma 1. *For all PSI execution graphs* $\Gamma = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathcal{T})$*:*

$$\mathsf{rf}_{\mathrm{I}} \cup \mathsf{mo}_{\mathrm{I}} \cup \mathsf{rb}_{\mathrm{I}} \subseteq \mathsf{po} \wedge \mathsf{irreflexive}((\mathsf{po}_{\mathsf{T}} \cup \mathsf{rf}_{\mathsf{T}} \cup \mathsf{mo}_{\mathsf{T}})^+; \mathsf{rb}_{\mathsf{T}}{}^?)$$
$$\iff$$
$$\mathsf{acyclic}(\mathsf{psi\text{-}hb}_{loc} \cup \mathsf{mo} \cup \mathsf{rb}) \quad \textit{where} \quad \mathsf{psi\text{-}hb} \triangleq (\mathsf{po} \cup \mathsf{rf} \cup \mathsf{rf}_{\mathsf{T}} \cup \mathsf{mo}_{\mathsf{T}})^+$$

Proof (the $\Leftarrow$ *direction).* Pick an arbitrary $\Gamma = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathcal{T})$ such that $\mathsf{acyclic}(\mathsf{psi\text{-}hb}_{loc} \cup \mathsf{mo} \cup \mathsf{rb})$ holds. It suffices to show that 1) $\mathsf{irreflexive}((\mathsf{po}_{\mathsf{T}} \cup \mathsf{rf}_{\mathsf{T}} \cup \mathsf{mo}_{\mathsf{T}})^+)$; 2) $\mathsf{irreflexive}((\mathsf{po}_{\mathsf{T}} \cup \mathsf{rf}_{\mathsf{T}} \cup \mathsf{mo}_{\mathsf{T}})^+; \mathsf{rb}_{\mathsf{T}})$; 3) $\mathsf{rf}_{\mathrm{I}} \subseteq \mathsf{po}$; 4) $\mathsf{mo}_{\mathrm{I}} \subseteq \mathsf{po}$; and 5) $\mathsf{rb}_{\mathrm{I}} \subseteq \mathsf{po}$.

RTS. (1)
We proceed by contradiction. Pick an arbitrary a such that $(a, a) \in (\mathsf{po}_{\mathsf{T}} \cup \mathsf{rf}_{\mathsf{T}} \cup \mathsf{mo}_{\mathsf{T}})^+$. From the definition of po_{T} we then have $(a, a) \in (\mathsf{po} \cup \mathsf{rf}_{\mathsf{T}} \cup \mathsf{mo}_{\mathsf{T}})^+$ and thus $(a, a) \in \mathsf{psi\text{-}hb}$, contradicting the assumption that $\mathsf{acyclic}(\mathsf{psi\text{-}hb}_{loc} \cup \mathsf{mo} \cup \mathsf{rb})$ holds.

RTS. (2)
We proceed by contradiction. Pick arbitrary a, b such that $(a, b) \in (\mathsf{po}_{\mathsf{T}} \cup \mathsf{rf}_{\mathsf{T}} \cup \mathsf{mo}_{\mathsf{T}})^+$ and $(b, a) \in \mathsf{rb}_{\mathsf{T}}$. From the definition of rb_{T} we then know that $[a]_{\mathsf{st}} \neq [b]_{\mathsf{st}}$ and that there exist c, d such that $[a]_{\mathsf{st}} = [c]_{\mathsf{st}}$, $[b]_{\mathsf{st}} = [d]_{\mathsf{st}}$ and that $(d, c) \in \mathsf{rb}$. On the other hand, from the proof of part (1) we have $(a, b) \in \mathsf{psi\text{-}hb}$. As such, from the auxiliary Lemma 2 in §1.3 we have $[a]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq (\mathsf{po}_{\mathsf{T}} \cup \mathsf{rf}_{\mathsf{T}} \cup \mathsf{mo}_{\mathsf{T}})^+$ and thus $[a]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq \mathsf{psi\text{-}hb}$. In particular we have $(c, d) \in \mathsf{psi\text{-}hb}$. As such we have $c \xrightarrow{\mathsf{psi\text{-}hb}} d \xrightarrow{\mathsf{rb}} c$, contradicting the assumption that $\mathsf{acyclic}(\mathsf{psi\text{-}hb}_{loc} \cup \mathsf{mo} \cup \mathsf{rb})$ holds.

RTS. (3)
We proceed by contradiction. Pick arbitrary a, b such that $(a, b) \in \mathsf{rf}_{\mathrm{I}}$ and $(a, b) \notin \mathsf{po}$. As a and b are in the same transaction class (from the definition of rf_{I}), we know that they are related by po. As $(a, b) \notin \mathsf{po}$ we then know $(b, a) \in \mathsf{po}$. As such we have $a \xrightarrow{\mathsf{rf}} b \xrightarrow{\mathsf{po}} a$, contradicting the assumption that $\mathsf{acyclic}(\mathsf{psi\text{-}hb}_{loc} \cup \mathsf{mo} \cup \mathsf{rb})$ holds. The proof of parts (4-5) is analogous and omitted here.

Proof (the ⇒ direction). Pick an arbitrary $\Gamma = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathcal{T})$ such that $\mathsf{rf}_\mathrm{I} \cup \mathsf{mo}_\mathrm{I} \cup \mathsf{rb}_\mathrm{I} \subseteq \mathsf{po} \land \mathsf{irreflexive}((\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+; \mathsf{rb}_\mathsf{T}{}^?)$ holds. It suffices to show that 1) $\mathsf{irreflexive}(\mathsf{psi\text{-}hb}_{loc})$; 2) $\mathsf{irreflexive}(\mathsf{psi\text{-}hb}_{loc}; \mathsf{mo})$; and 3) $\mathsf{irreflexive}(\mathsf{psi\text{-}hb}_{loc}; \mathsf{rb})$.

RTS. (1)
We proceed by contradiction. Pick arbitrary a such that $(a, a) \in \mathsf{psi\text{-}hb}_{loc}$. From the auxiliary Lemma 3in §1.3 below we then have $(a, a) \in \mathsf{po}$, contradicting the assumption that po is a strict total order on the events of each thread.

RTS. (2)
We proceed by contradiction. Pick arbitrary a, b such that $(a, b) \in \mathsf{psi\text{-}hb}_{loc}$ and $(b, a) \in \mathsf{mo}$. There are now two cases to consider: 1) $[a]_\mathsf{st} = [b]_\mathsf{st}$; or 2) $[a]_\mathsf{st} \neq [b]_\mathsf{st}$. In case (1) from the auxiliary Lemma 3 in §1.3 we then have $(a, b) \in \mathsf{po}$. On the other hand, we have $(b, a) \in \mathsf{mo}$ and thus $(b, a) \in \mathsf{mo}_\mathrm{I} \subseteq \mathsf{po}$. As such we have $a \stackrel{\mathsf{po}}{\to} b \stackrel{\mathsf{po}}{\to} a$, contradicting the assumption that po is a strict total order on the events of each thread.

In case (2) from the auxiliary Lemma 2 in §1.3 we have $[a]_\mathsf{st} \times [b]_\mathsf{st} \subseteq (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+$. On the other hand as we have $(b, a) \in \mathsf{mo}$ and $[a]_\mathsf{st} \neq [b]_\mathsf{st}$ we have $(b, a) \in \mathsf{mo}_\mathsf{T}$ and thus $(b, a) \in (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+$ By the definition of transitive closures we thus have $(a, a) \in (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+$, contradicting the assumption that $\mathsf{irreflexive}((\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+)$ holds.

RTS. (3)
We proceed by contradiction. Pick arbitrary a, b such that $(a, b) \in \mathsf{psi\text{-}hb}_{loc}$ and $(b, a) \in \mathsf{rb}$. There are now two cases to consider: 1) $[a]_\mathsf{st} = [b]_\mathsf{st}$; or 2) $[a]_\mathsf{st} \neq [b]_\mathsf{st}$. In case (1) from the auxiliary Lemma 3 in §1.3 we then have $(a, b)\mathsf{po}$. On the other hand, we have $(b, a) \in \mathsf{rb}$ and thus $(b, a) \in \mathsf{rb}_\mathrm{I} \subseteq \mathsf{po}$. As such we have $a \stackrel{\mathsf{po}}{\to} b \stackrel{\mathsf{po}}{\to} a$, contradicting the assumption that po is a strict total order on the events of each thread.

In case (2) from the auxiliary Lemma 2 in §1.3 we have $[a]_\mathsf{st} \times [b]_\mathsf{st} \subseteq (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+$. On the other hand as we have $(b, a) \in \mathsf{rb}$ and $[a]_\mathsf{st} \neq [b]_\mathsf{st}$ we have $(b, a) \in \mathsf{rb}_\mathsf{T}$. We thus have $(a, a) \in (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+; \mathsf{rb}_\mathsf{T}$, contradicting the assumption that $\mathsf{irreflexive}((\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+; \mathsf{rb}_\mathsf{T})$ holds. □

1.1 PSI Implementation Soundness

Our PSI implementation in Fig. 1 is *sound*: for each consistent implementation graph G, a corresponding specification graph Γ can be constructed with the same program outcome such that $\mathsf{psi\text{-}consistent}(\Gamma)$ holds.

Constructing Consistent Specification Graphs Observe that given an execution of our implementation with t transactions, the trace of each transaction $i \in \{1 \cdots t\}$ is of the form $\theta_i = Ls_i \stackrel{\mathsf{po}}{\to} FS_i \stackrel{\mathsf{po}}{\to} S_i \stackrel{\mathsf{po}}{\to} Ts_i \stackrel{\mathsf{po}}{\to} Us_i$, where Ls_i, FS_i, S_i, Ts_i and Us_i respectively denote the sequence of events acquiring the version locks, attempting but failing to obtain a valid snapshot, recording a valid

```
0. for (x ∈ WS) lock vx;
1. for (x ∈ RS) {
2.   a := vx;
3.   if (is-odd(a) && x ∉ WS) continue;
4.   if (x ∉ WS) v[x] := a;
5.   s[x] := x;  }
6. for (x ∈ RS)
7.   if (¬valid(x)) goto line 1;
8. ⟦T⟧;
9. for (x ∈ WS) unlock vx;
```

```
lock vx ≜
  retry: v[x] := vx;
         if (is-odd(v[x]))
           goto retry;
         if (!CAS(vx,v[x],v[x]+1))
           goto retry;

unlock vx ≜ vx := v[x] + 2

    valid(x) ≜ vx == v[x]
validRPSI(x) ≜ vx == v[x]  &&  x == s[x]

      ⟦a := x⟧ ≜ a := s[x]
      ⟦x := a⟧ ≜ x := a; s[x] := a
     ⟦S1 ; S2⟧ ≜ ⟦S1⟧; ⟦S2⟧
⟦while(e) S⟧ ≜ while(e) ⟦S⟧
        ...  and so on  ...
```

Fig. 1: PSI implementation of transaction T given RS, WS; the RPSI implementation (§2) is obtained by replacing `valid` on line 7 with $\texttt{valid}_{\text{RPSI}}$.

snapshot, performing the transactional operations, and releasing the version locks. In particular, we have:

- Ls_i denotes the sequence of events acquiring the locks and is of the form $FL_1 \overset{\mathsf{po}|_{\text{imm}}}{\rightarrow} L_1 \overset{\mathsf{po}|_{\text{imm}}}{\rightarrow} \cdots \overset{\mathsf{po}|_{\text{imm}}}{\rightarrow} FL_j \overset{\mathsf{po}|_{\text{imm}}}{\rightarrow} L_j$ with
$$FL_n{=}\texttt{R}(\texttt{sy}_n, wa'_n)^* \qquad L_n{=}\texttt{U}(\texttt{sy}_n, wa_n, wa_n{+}1)$$
such that $wa'_n \bmod 2 = 1$ and $wa_n \bmod 2 = 0$;
- Us_i denotes the sequence of events releasing the locks for the write set. That is, the events in Us correspond to the execution of the last line of the implementation in Fig. 1, and is of the form $U_1 \overset{\mathsf{po}|_{\text{imm}}}{\rightarrow} \cdots \overset{\mathsf{po}|_{\text{imm}}}{\rightarrow} U_j$ with
$$U_n = \texttt{W}(\texttt{sy}_n, wa_n + 2)$$

Given a transaction ξ we write e.g. $\xi.U^{\texttt{x}}$ to denote the event in Us_i above releasing the version lock on `x`.

For each transactional trace θ_i of our implementation, we thus construct a corresponding trace of the specification as $\theta'_i = B_i \overset{\mathsf{po}}{\rightarrow} Ts'_i \overset{\mathsf{po}}{\rightarrow} E_i$, where B_i and E_i denote the transaction begin and end events ($\texttt{lab}(B_i){=}\texttt{B}$ and $\texttt{lab}(E_i){=}\texttt{E}$). When Ts_i is of the form $t_1 \overset{\mathsf{po}}{\rightarrow} \cdots \overset{\mathsf{po}}{\rightarrow} t_n$, we construct Ts'_i as $t'_1 \overset{\mathsf{po}}{\rightarrow} \cdots \overset{\mathsf{po}}{\rightarrow} t'_n$ with each t'_j defined either as $t'_j \triangleq \texttt{R}(\texttt{x}, v)$ when $t_j = \texttt{R}(\texttt{s[x]}, v)$ (i.e. the corresponding implementation event is a read event); or as $t'_j \triangleq \texttt{W}(\texttt{x}, v)$ when $t_j{=}\texttt{W}(\texttt{x}, v) \overset{\mathsf{po}}{\rightarrow} \texttt{W}(\texttt{s[x]}, v)$.

For each specification trace θ'_i we construct the 'reads-from' relation as:

$$\mathsf{RF}_i \triangleq \left\{ (w, t'_j) \;\middle|\; \begin{array}{l} t'_j \in Ts'_i \wedge \exists \mathtt{x}, v.\ t'_j {=} \mathtt{R}(\mathtt{x}, v) \wedge w {=} \mathtt{W}(\mathtt{x}, v) \\ \wedge (w \in Ts'_i \Rightarrow w \xrightarrow{\mathsf{po}} t'_j \wedge \\ \qquad\qquad (\forall e \in Ts'_i.\ w \xrightarrow{\mathsf{po}} e \xrightarrow{\mathsf{po}} t'_j \Rightarrow (\mathtt{loc}(e) {\neq} \mathtt{x} \vee e {\notin} \mathcal{W}))) \\ \wedge (w \notin Ts'_i \Rightarrow (\forall e {\in} Ts'_i.\ (e \xrightarrow{\mathsf{po}} t'_j \Rightarrow (\mathtt{loc}(e) \neq \mathtt{x} \vee e \notin \mathcal{W})) \\ \qquad\qquad \wedge \exists r' \in S_i.\ \mathtt{loc}(r') {=} \mathtt{x} \wedge (w, r') \in G.\mathsf{rf}) \end{array} \right\}$$

That is, we construct our graph such that each read event t'_j from location $\mathtt{x}$ in Ts'_i either i) is preceded by a write event w to $\mathtt{x}$ in Ts'_i without an intermediate write in between them and thus 'reads-from' w (lines two and three); or ii) is not preceded by a write event in Ts'_i and thus 'reads-from' the write event w from which the initial snapshot read r' in S_i obtained the value of $\mathtt{x}$ (last two lines).

Given a consistent implementation graph $G = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo})$, we construct a consistent specification graph $\Gamma = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathcal{T})$ such that:

- $\Gamma.E \triangleq \bigcup_{i \in \{1 \cdots t\}} \theta'_i.E$ – the events of $\Gamma.E$ is the union of events in each transaction trace θ'_i of the specification constructed as above;
- $\Gamma.\mathsf{po} \triangleq G.\mathsf{po}|_{\Gamma.E}$ – the $\Gamma.\mathsf{po}$ is that of $G.\mathsf{po}$ limited to the events in $\Gamma.E$;
- $\Gamma.\mathsf{rf} \triangleq \bigcup_{i \in \{1 \cdots t\}} \mathsf{RF}_i$ – the $\Gamma.\mathsf{rf}$ is the union of RF_i relations defined above;
- $\Gamma.\mathsf{mo} \triangleq G.\mathsf{mo}|_{\Gamma.E}$ – the $\Gamma.\mathsf{mo}$ is that of $G.\mathsf{mo}$ limited to the events in $\Gamma.E$;
- $\Gamma.\mathcal{T} \triangleq \Gamma.E$, where for each $e \in \Gamma.\mathcal{T}$, we define $\mathtt{tx}(e) = i$ when $e \in \theta'_i$.

Theorem 1 (Soundness). *For all RA-consistent implementation graphs G of the implementation in Fig. 1, there exists a PSI-consistent specification graph Γ of the corresponding transactional program that has the same program outcome.*

Proof. Pick an arbitrary G such that $\mathsf{RA\text{-}consistent}(G)$, and its associated Γ constructed as described above. It then suffices to show 1) $\mathsf{irreflexive}(\Gamma.\mathsf{psi\text{-}hb}_{loc})$; 2) $\mathsf{irreflexive}(\Gamma.\mathsf{psi\text{-}hb}_{loc}; \Gamma.\mathsf{mo})$; and 3) $\mathsf{irreflexive}(\Gamma.\mathsf{psi\text{-}hb}_{loc}; \Gamma.\mathsf{rb})$.

RTS. $\mathsf{irreflexive}(\Gamma.\mathsf{psi\text{-}hb}_{loc})$
We proceed by contradiction. Let us assume there exists a such that $(a, a) \in \Gamma.\mathsf{psi\text{-}hb}_{loc}$. From the auxiliary Lemma 5.5 in §1.3 we have $(a, a) \in G.\mathsf{hb}$, contradicting our assumption that G is consistent.

RTS. $\mathsf{irreflexive}(\Gamma.\mathsf{psi\text{-}hb}_{loc}; \Gamma.\mathsf{mo})$
We proceed by contradiction. Let us assume there exist a, b such that $(a, b) \in \Gamma.\mathsf{psi\text{-}hb}_{loc}$ and $(b, a) \in \Gamma.\mathsf{mo}$. From the auxiliary Lemma 5.5 in §1.3 we have $(a, b) \in G.\mathsf{hb}$. Similarly, from auxiliary Lemma 5.3 in §1.3 we have $(b, a) \in G.\mathsf{hb}$. As $G.\mathsf{hb}$ is transitively closed, we have $(a, a) \in G.\mathsf{hb}$, contradicting our assumption that G is consistent.

RTS. $\mathsf{irreflexive}(\Gamma.\mathsf{psi\text{-}hb}_{loc}; \Gamma.\mathsf{rb})$
We proceed by contradiction. Let us assume there exist w, r such that $(w, r) \in \Gamma.\mathsf{psi\text{-}hb}_{loc}$ and $(r, w) \in \Gamma.\mathsf{rb}$. There are then two cases to consider: 1) $[w]_{\mathsf{st}} = [r]_{\mathsf{st}}$; or 2) $[w]_{\mathsf{st}} \neq [r]_{\mathsf{st}}$.

In case (1), let $\texttt{loc}(w) = \texttt{loc}(r) = \texttt{x}$. From the definition of $\Gamma.\textsf{rb}$, $\Gamma.\textsf{rf}$ and $\Gamma.\textsf{mo}$ we know there exists w' such that $[w']_{\textsf{st}} = [r]_{\textsf{st}}$, $(w', r) \in \Gamma.\textsf{rf}$, $(w', r) \in \Gamma.\textsf{po}$, $(w', w) \in \Gamma.\textsf{mo}$, $(w', w) \in G.\textsf{mo}$ and for all e if $w' \stackrel{\Gamma.\textsf{po}}{\to} e \stackrel{\Gamma.\textsf{po}}{\to} r$ then $\texttt{loc}(e) \neq \texttt{x} \lor e \notin \mathcal{W}$. Now since $[w']_{\textsf{st}} = [w]_{\textsf{st}}$ we know that either a) $(w, w') \in \Gamma.\textsf{po}$; or b) $(w', w) \in \Gamma.\textsf{po}$.

In case (1.a) from the definition of $\Gamma.\textsf{po}$ we have $(w, w') \in G.\textsf{po}$. As such we have $w \stackrel{G.\textsf{po}}{\to} w' \stackrel{G.\textsf{mo}}{\to} w$, contradicting the assumption that $G.\textsf{hb}$ is consistent.

In case (1.b) since $[w]_{\textsf{st}} = [r]_{\textsf{st}}$ we know that either i) $(w, r) \in \Gamma.\textsf{po}$; or 2) $(r, w) \in \Gamma.\textsf{po}$. Moreover, since we know for all e if $w' \stackrel{\Gamma.\textsf{po}}{\to} e \stackrel{\Gamma.\textsf{po}}{\to} r$ then $\texttt{loc}(e) \neq \texttt{x} \lor e \notin \mathcal{W}$, and we have $w \in \mathcal{W}$ and $\texttt{loc}(w) = \texttt{x}$, we thus know $(r, w) \in \Gamma.\textsf{po}$. From the definition of $\Gamma.\textsf{po}$ we then have $(r, w) \in G.\textsf{po} \subseteq G.\textsf{hb}$ We thus have $w \stackrel{G.\textsf{hb}}{\to} r \stackrel{G.\textsf{hb}}{\to} w$, contradicting the assumption that G is consistent.

In case (2) we then know, there exists ξ_1, ξ_2 such that $\xi_1 \neq \xi_2$, $r \in \xi_1$, $w \in \xi_2$. Let $w = \texttt{W}(\texttt{x}, -)$ and $r = \texttt{R}(\texttt{x}, v)$. From the construction of Γ we know $\xi_1.L_{\texttt{vx}} \stackrel{\textsf{po}}{\to} w \stackrel{\textsf{po}}{\to} \xi_1.U_{\texttt{vx}}$ and that there exist $rv_1, rx, rv_2 \in G.E$ and d such that $rv_1 = \texttt{R}(\texttt{vx}, d)$, $rx = \texttt{R}(\texttt{x}, v)$, $rv_2 = \texttt{R}(\texttt{vx}, d)$, $(rx, w) \in G.\textsf{rb}$ and $rv_1 \stackrel{G.\textsf{po}}{\to} rx \stackrel{G.\textsf{po}}{\to} rv_2 \stackrel{G.\textsf{po}}{\to} r = \texttt{R}(\texttt{s[x]}, v))$.

We will first demonstrate that $(rv_2, \xi_1.L_{\texttt{vx}}) \in G.\textsf{rb}$. There are two cases to consider: A) either $\texttt{x} \in \texttt{WS}_{\xi_2}$; or $\texttt{x} \notin \texttt{WS}_{\xi_2}$.

In the former case (A), from the auxiliary Lemma 4.2 in §1.3 we then know that either i) $\xi_2.U_{\texttt{vx}} \stackrel{G.\textsf{hb}}{\to} \xi_1.L_{\texttt{vx}}$; or ii) $\xi_1.U_{\texttt{vx}} \stackrel{G.\textsf{hb}}{\to} \xi_2.L_{\texttt{vx}}$. In the former case (A.i), we then have $\xi_2.U_{\texttt{vx}} \stackrel{G.\textsf{mo}}{\to} \xi_1.L_{\texttt{vx}}$ (since otherwise we would have a cycle $\xi_2.U_{\texttt{vx}} \stackrel{G.\textsf{hb}}{\to} \xi_1.L_{\texttt{vx}} \stackrel{G.\textsf{mo}}{\to} \xi_2.U_{\texttt{vx}}$, contradicting our assumption that G is consistent). As such we have $(rv_2, \xi_1.L_{\texttt{vx}}) \in G.\textsf{rb}$. In the latter case (A.ii) we then have $w \stackrel{G.\textsf{po}}{\to} \xi_1.U_{\texttt{vx}} \stackrel{G.\textsf{hb}}{\to} \xi_1.L_{\texttt{vx}} \stackrel{G.\textsf{po}}{\to} rx \stackrel{G.\textsf{rb}}{\to} w$. That is, since we have $G.\textsf{po} \subseteq G.\textsf{hb}$ and $G.\textsf{hb}$ is transitively closed, we have $w \stackrel{G.\textsf{hb}}{\to} rx \stackrel{G.\textsf{rb}}{\to} w$, contradicting the assumption that G is consistent. So in case (A) we know $(rv_2, \xi_1.L_{\texttt{vx}}) \in G.\textsf{rb}$.

In the latter case (B) we then know d (in $rv_1 = \texttt{R}(\texttt{vx}, d)$) is even. As such, from our implementation we know there exists ξ_3 such that $\texttt{x} \in \texttt{WS}_{\xi_3}$, $\xi_3.L_{\texttt{vx}} \stackrel{\textsf{po}}{\to} \xi_3.U_{\texttt{vx}}$, and that $(\xi_3.U_{\texttt{vx}}, rv_1) \in G.\textsf{rf}$. Since the values written to $\texttt{vx}$ are unique (Lemma 4.3 in §1.3) and $\texttt{val}_{\texttt{r}}(rv_1) = \texttt{val}_{\texttt{r}}(rv_2) = d$, we also have $(\xi_3.U_{\texttt{vx}}, rv_2) \in G.\textsf{rf}$. On the other hand, from Lemma 4.2 in §1.3 we have either i) $\xi_1.U_{\texttt{vx}} \stackrel{G.\textsf{hb}}{\to} \xi_3.L_{\texttt{vx}}$, or ii) $\xi_3.U_{\texttt{vx}} \stackrel{G.\textsf{hb}}{\to} \xi_1.L_{\texttt{vx}}$.

In the former case (B.i) we then have $rx \stackrel{G.\textsf{rb}}{\to} w \stackrel{G.\textsf{po}}{\to} \xi_1.U_{\texttt{vx}} \stackrel{G.\textsf{hb}}{\to} \xi_3.L_{\texttt{vx}} \stackrel{G.\textsf{po}}{\to} \xi_3.U_{\texttt{vx}} \stackrel{G.\textsf{rf}}{\to} rv_1 \stackrel{G.\textsf{po}}{\to} rx$. As $G.\textsf{po}, G.\textsf{rf} \subseteq G.\textsf{hb}$ and $G.\textsf{hb}$ is transitively closed, we then have $rx \stackrel{G.\textsf{rb}}{\to} w \stackrel{G.\textsf{hb}}{\to} rx$, contradicting the assumption that G is consistent.

We thus know that the only possible case in (B) is that of (B.ii) and we have $\xi_3.U_{\texttt{vx}} \stackrel{G.\textsf{hb}}{\to} \xi_1.L_{\texttt{vx}}$. Consequently we have $\xi_3.U_{\texttt{vx}} \stackrel{G.\textsf{mo}}{\to} \xi_1.L_{\texttt{vx}}$ (since otherwise we would have a cycle $\xi_3.U_{\texttt{vx}} \stackrel{G.\textsf{hb}}{\to} \xi_1.L_{\texttt{vx}} \stackrel{G.\textsf{mo}}{\to} \xi_3.U_{\texttt{vx}}$, contradicting our assumption that G is consistent). As such, we have $(rv_2, \xi_1.L_{\texttt{vx}}) \in G.\textsf{rb}$.

In both cases (A) and (B) we have $(rv_2, \xi_1.L_{\texttt{vx}}) \in G.\textsf{rb}$. On the other hand, recall that we have $(w, r) \in G.\textsf{hb}$. It is straightforward to demonstrate that $G.\textsf{hb} = G.\textsf{po}^+ \cup (G.\textsf{po} \cup G.\textsf{rf})^*; (G.\textsf{rf} \setminus G.\textsf{po}); G.\textsf{po}^*$. There are then two cases to consider: 1) $(w, r) \in G.\textsf{po}^+$; or 2) $(w, r) \in (G.\textsf{po} \cup G.\textsf{rf})^*; (G.\textsf{rf} \setminus G.\textsf{po}); G.\textsf{po}^*$.

In case (1), since we have $(w, r) \in G.\textsf{po}^+$ and w and rx belong to two distinct transactions ($\xi_1 \neq \xi_2$), we also have $(w, rx) \in G.\textsf{po}^+$. As such, we have $w \overset{G.\textsf{po}^+}{\rightarrow} rx \overset{G.\textsf{rb}}{\rightarrow} w$, contradicting the assumption that G is consistent.

In case (2), we know there exist d, e such that $(w, d) \in (G.\textsf{po} \cup G.\textsf{rf})^*$, $(d, e) \in (G.\textsf{rf} \setminus G.\textsf{po})$ and $(e, r) \in G.\textsf{po}^*$. Either a) $e \notin \xi_2$ or b) $e \in \xi_2$.

In case (2.a) since we have $(e, r) \in G.\textsf{po}^*$ and $r \in \xi_2$, we know that $(e, r) \in G.\textsf{po}$. On the other hand, since $e \notin \xi_2$ and $rx \in \xi_2$, we have $(e, rx) \in G.\textsf{po}$. As such, we have $w \overset{(G.\textsf{po} \cup G.\textsf{rf})^*}{\rightarrow} d \overset{G.\textsf{rf}}{\rightarrow} e \overset{G.\textsf{po}}{\rightarrow} rx \overset{G.\textsf{rb}}{\rightarrow} w$. That is, since $G.\textsf{rf} \subseteq G.\textsf{hb}$, $G.\textsf{po} \subseteq G.\textsf{hb}$ and $G.\textsf{hb}$ is transitively closed, we have $w \overset{G.\textsf{hb}}{\rightarrow} rx \overset{G.\textsf{rb}}{\rightarrow} w$, contradicting the assumption that G is consistent.

In the latter case (2.b) there are two additional cases to consider: either i) $e \overset{G.\textsf{po}}{\rightarrow} rv_2$, or ii) $rv_2 \overset{G.\textsf{po}}{\rightarrow} e$. In case (2.b.i) we then have $rv_2 \overset{G.\textsf{rb}}{\rightarrow} \xi_1.L_{\texttt{vx}} \overset{G.\textsf{po}}{\rightarrow} w \overset{(G.\textsf{po} \cup G.\textsf{rf})^*}{\rightarrow} d \overset{G.\textsf{rf}}{\rightarrow} e \overset{G.\textsf{po}}{\rightarrow} rv_2$. That is, since $G.\textsf{rf} \subseteq G.\textsf{hb}$, $G.\textsf{po} \subseteq G.\textsf{hb}$ and $G.\textsf{hb}$ is transitively closed, we have $rv_2 \overset{G.\textsf{rb}}{\rightarrow} \xi_1.L_{\texttt{vx}} \overset{G.\textsf{hb}}{\rightarrow} rv_2$, contradicting the assumption that G is consistent.

In case (2.b.ii), since $e \in \xi_2$ and e is a read event, we know that $\texttt{loc}(e) = \texttt{vy}$ for some location lock $\texttt{vy}$, where $\texttt{y}$ is in the read set of ξ_2. Let $e = \texttt{R}(\texttt{vy}, v')$. From our implementation we then know that there exists $e' = \texttt{R}(\texttt{vy}, v')$ such that $e' \overset{G.\textsf{po}}{\rightarrow} rv_2$. On the other hand, since $(d, e) \in G.\textsf{rf}$ and from the auxiliary Lemma 4.3 in §1.3 we know that the writes to $\texttt{vy}$ have unique values, we have $(d, e') \in G.\textsf{rf}$. As such we have $rv_2 \overset{G.\textsf{rb}}{\rightarrow} \xi_1.L_{\texttt{vx}} \overset{G.\textsf{po}}{\rightarrow} w \overset{(G.\textsf{po} \cup G.\textsf{rf})^*}{\rightarrow} d \overset{G.\textsf{rf}}{\rightarrow} e' \overset{G.\textsf{po}}{\rightarrow} rv_2$. That is, since $G.\textsf{rf} \subseteq G.\textsf{hb}$, $G.\textsf{po} \subseteq G.\textsf{hb}$ and $G.\textsf{hb}$ is transitively closed, we have $rv_2 \overset{G.\textsf{rb}}{\rightarrow} \xi_1.L_{\texttt{vx}} \overset{G.\textsf{hb}}{\rightarrow} rv_2$, contradicting the assumption that G is consistent. □

1.2 Implementation Completeness

The PSI implementation in Fig. 1 is *complete*: for each consistent specification graph Γ a corresponding implementation graph Γ can be constructed with the same program outcome such that $\textsf{RA-consistent}(G)$ holds.

Constructing Consistent Implementation Graphs In order to construct an execution graph of the implementation G from the specification Γ, we follow similar steps as those in the soundness construction, in reverse order. More concretely, given each trace θ'_i of the specification, we construct an analogous trace of the implementation by inserting the appropriate events for acquiring and inspecting the version locks, as well as obtaining a snapshot. For each transaction class $\mathcal{T}_i \in \mathcal{T}/\textsf{st}$, we must first determine its read and write sets and subsequently decide the order in which the version locks are acquired (for locations in the write set) and inspected (for locations in the read set). This then enables us to construct the 'reads-from' and 'modification-order' relations for the events associated with version locks.

Given a consistent execution graph of the specification $\Gamma = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathcal{T})$, and a transaction class $\mathcal{T}_i \in \Gamma.\mathcal{T}/\mathsf{st}$, we write $\mathtt{WS}_{\mathcal{T}_i}$ for the set of locations written to by $\mathcal{T}_i$. That is, $\mathtt{WS}_{\mathcal{T}_i} \triangleq \bigcup_{e \in \mathcal{T}_i \cap \mathcal{W}} \mathtt{loc}(e)$. Similarly, we write $\mathtt{RS}_{\mathcal{T}_i}$ for the set of locations read from by $\mathcal{T}_i$, *prior to* being written to by $\mathcal{T}_i$. For each location $\mathtt{x}$ read from by $\mathcal{T}_i$, we additionally record the first read event in $\mathcal{T}_i$ that retrieved the value of $\mathtt{x}$. That is,

$$\mathtt{RS}_{\mathcal{T}_i} \triangleq \left\{ (\mathtt{x}, r) \,\middle|\, r \in \mathcal{T}_i \cap \mathcal{R}_{\mathtt{x}} \wedge \neg\exists e \in \mathcal{T}_i \cap E_{\mathtt{x}}.\, e \xrightarrow{\mathsf{po}} r \right\}$$

Note that transaction $\mathcal{T}_i$ may contain several read events reading from $\mathtt{x}$, prior to subsequently updating it. However, the internal-read-consistency property ensures that all such read events read from the same write event. As such, as part of the read set of $\mathcal{T}_i$ we record the first such read event (in program-order).

Determining the ordering of lock events hinges on the following observation. Given a consistent execution graph of the specification $\Gamma = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathcal{T})$, let for each location $\mathtt{x}$ the total order mo be given as: $w_1 \xrightarrow{\mathsf{mo}|_{\mathrm{imm}}} \cdots \xrightarrow{\mathsf{mo}|_{\mathrm{imm}}} w_{n_{\mathtt{x}}}$. Observe that this order can be broken into adjacent segments where the events of each segment belong to the *same* transaction. That is, given the transaction classes $\Gamma.\mathcal{T}/\mathsf{st}$, the order above is of the following form where $\mathcal{T}_1, \cdots, \mathcal{T}_m \in \Gamma.\mathcal{T}/\mathsf{st}$ and for each such $\mathcal{T}_i$ we have $\mathtt{x} \in \mathtt{WS}_{\mathcal{T}_i}$ and $w_{(i,1)} \cdots w_{(i,n_i)} \in \mathcal{T}_i$:

$$\underbrace{w_{(1,1)} \xrightarrow{\mathsf{mo}|_{\mathrm{imm}}} \cdots \xrightarrow{\mathsf{mo}|_{\mathrm{imm}}} w_{(1,n_1)}}_{\mathcal{T}_1} \xrightarrow{\mathsf{mo}|_{\mathrm{imm}}} \cdots \xrightarrow{\mathsf{mo}|_{\mathrm{imm}}} \underbrace{w_{(m,1)} \xrightarrow{\mathsf{mo}|_{\mathrm{imm}}} \cdots \xrightarrow{\mathsf{mo}|_{\mathrm{imm}}} w_{(m,n_m)}}_{\mathcal{T}_m}$$

Were this not the case and we had $w_1 \xrightarrow{\mathsf{mo}} w \xrightarrow{\mathsf{mo}} w_2$ such that $w_1, w_2 \in \mathcal{T}_i$ and $w \in \mathcal{T}_j \neq \mathcal{T}_i$, we would consequently have $w_1 \xrightarrow{\mathsf{mo}_{\mathsf{T}}} w \xrightarrow{\mathsf{mo}_{\mathsf{T}}} w_1$, contradicting the assumption that Γ is consistent. Given the above order, let us then define $\Gamma.\mathsf{MO}_{\mathtt{x}} = [\mathcal{T}_1 \cdots \mathcal{T}_m]$. We write $\Gamma.\mathsf{MO}_{\mathtt{x}}|_i$ for the i^{th} item of $\Gamma.\mathsf{MO}_{\mathtt{x}}$. As we describe shortly, we use $\Gamma.\mathsf{MO}_{\mathtt{x}}$ to determine the order of lock events.

Note that the execution trace for each transaction $\mathcal{T}_i \in \Gamma.\mathcal{T}/\mathsf{st}$ is of the form $\theta'_i = B_i \xrightarrow{\mathsf{po}} Ts'_i \xrightarrow{\mathsf{po}} E_i$, where B_i is a transaction-begin ($\mathtt{B}$) event, E_i is a transaction-end ($\mathtt{E}$) event, and $Ts'_i = t'_1 \xrightarrow{\mathsf{po}} \cdots \xrightarrow{\mathsf{po}} t'_n$ for some n, where each t'_j is either a read or a write event. As such, we have $\Gamma.E = \Gamma.\mathcal{T} = \bigcup_{\mathcal{T}_i \in \Gamma.\mathcal{T}/\mathsf{st}} \mathcal{T}_i = \theta'_i.E$.

For each trace θ'_i of the specification, we construct a corresponding trace of our implementation θ_i as follows. Let $\mathtt{RS}_{\mathcal{T}_i} = \{(\mathtt{x}_1, r_1) \cdots (\mathtt{x}_p, r_p)\}$ and $\mathtt{WS}_{\mathcal{T}_i} = \{\mathtt{y}_1 \cdots \mathtt{y}_q\}$. We then construct $\theta_i = Ls_i \xrightarrow{\mathsf{po}} S_i \xrightarrow{\mathsf{po}} Ts_i \xrightarrow{\mathsf{po}} Us_i$, where

- $Ls_i = L_i^{\mathtt{y}_1} \xrightarrow{\mathsf{po}} \cdots \xrightarrow{\mathsf{po}} L_i^{\mathtt{y}_q}$ and $U_i = U_i^{\mathtt{y}_1} \xrightarrow{\mathsf{po}} \cdots \xrightarrow{\mathsf{po}} U_i^{\mathtt{y}_q}$ denote the sequence of events acquiring and releasing the version locks, respectively. Each $L_i^{\mathtt{y}_j}$ and $U_i^{\mathtt{y}_j}$ are defined as follows, the first event $L_i^{\mathtt{y}_1}$ has the same identifier as that of B_i, the last event $U_i^{\mathtt{y}_q}$ has the same identifier as that of E_i, and the identifiers of the remaining events are picked fresh:

$$L_i^{\mathtt{y}_j} = \mathtt{U}(\mathtt{vy}_j, 2a, 2a{+}1) \quad U_i^{\mathtt{y}_j} = \mathtt{W}(\mathtt{vy}_j, 2a{+}2) \text{ where } \mathsf{MO}_{\mathtt{y}_j}\big|_a = \mathcal{T}_i$$

We then define the mo relation for version locks such that if transaction $\mathcal{T}_i$ writes to y immediately after $\mathcal{T}_j$ (i.e. $\mathcal{T}_i$ is $\mathsf{MO}_\mathtt{y}$-ordered immediately after $\mathcal{T}_j$), then $\mathcal{T}_i$ acquires the vy version lock immediately after $\mathcal{T}_j$ has released it. On the other hand, if $\mathcal{T}_i$ is the first transaction to write to y, then it acquires vy immediately after the event initialising the value of vy, written $init_{\mathtt{vy}}$. Moreover, each vy release event of $\mathcal{T}_i$ is mo-ordered immediately after the corresponding vy acquisition event in $\mathcal{T}_i$:

$$\mathsf{IMO}_i \triangleq \bigcup_{\mathtt{y} \in \mathtt{WS}_{\mathcal{T}_i}} \left\{ \begin{matrix}(L_i^\mathtt{y}, U_i^\mathtt{y}),\\ (w, L_i^\mathtt{y})\end{matrix} \;\middle|\; \begin{matrix}(\Gamma.\mathsf{MO}_\mathtt{x}|_0 = \mathcal{T}_i \Rightarrow w = init_{\mathtt{vy}}) \wedge \\ (\exists \mathcal{T}_j, a > 0.\ \Gamma.\mathsf{MO}_\mathtt{y}\big|_a = \mathcal{T}_i \wedge \Gamma.\mathsf{MO}_\mathtt{y}\big|_{a-1} = \mathcal{T}_j \\ \Rightarrow w = U_j^\mathtt{y})\end{matrix} \right\}$$

This partial mo order on lock events of $\mathcal{T}_i$ also determines the rf relation for its lock acquisition events: $\mathsf{IRF}_i^1 \triangleq \bigcup_{\mathtt{y} \in \mathtt{WS}_{\mathcal{T}_i}} \{(w, L_i^\mathtt{y}) \mid (w, L_i^\mathtt{y}) \in \mathsf{IMO}_i\}$.

- $S_i = tr_i^{\mathtt{x}_1} \stackrel{\mathsf{po}}{\rightarrow} \cdots \stackrel{\mathsf{po}}{\rightarrow} tr_i^{\mathtt{x}_p} \stackrel{\mathsf{po}}{\rightarrow} vr_i^{\mathtt{x}_1} \stackrel{\mathsf{po}}{\rightarrow} \cdots \stackrel{\mathsf{po}}{\rightarrow} vr_i^{\mathtt{x}_p}$ denotes the sequence of events obtaining a tentative snapshot ($tr_i^{\mathtt{x}_j}$) and subsequently validating it ($vr_i^{\mathtt{x}_j}$). Each $tr_i^{\mathtt{x}_j}$ sequence is defined as $ir_i^{\mathtt{x}_j} \stackrel{\mathsf{po}}{\rightarrow} r_i^{\mathtt{x}_j} \stackrel{\mathsf{po}}{\rightarrow} s_i^{\mathtt{x}_j}$ (reading the version lock $\mathtt{vx}_j$, reading $\mathtt{x}_j$ and recoding it in s), with $ir_i^{\mathtt{x}_j}$, $r_i^{\mathtt{x}_j}$, $s_i^{\mathtt{x}_j}$ and $vr_i^{\mathtt{x}_j}$ events defined as follows (with fresh identifiers). We then define the rf relation for each of these read events in S_i.
  For each $(\mathtt{x}, r) \in \mathtt{RS}_{\mathcal{T}_i}$, when r (i.e. the read event in the specification class $\mathcal{T}_i$ that reads the value of x) reads from an event w in the specification graph ($(w, r) \in \Gamma.\mathsf{rf}$), we add $(w, r_i^\mathtt{x})$ to the rf relation of G (the first line of IRF_i^2 below). For version locks, if transaction $\mathcal{T}_i$ also writes to $\mathtt{x}_j$, then $ir_i^{\mathtt{x}_j}$ and $vr_i^{\mathtt{x}_j}$ events (reading and validating the value of version lock $\mathtt{vx}_j$), read from the lock event in $\mathcal{T}_i$ that acquired $\mathtt{vx}_j$, namely $L_i^{\mathtt{x}_j}$. On the other hand, if transaction $\mathcal{T}_i$ does not write to $\mathtt{x}_j$ and it reads the value of $\mathtt{x}_j$ written by $\mathcal{T}_j$, then $ir_i^{\mathtt{x}_j}$ and $vr_i^{\mathtt{x}_j}$ read the value written to $\mathtt{vx}_j$ by $\mathcal{T}_j$ when releasing it ($U_j^\mathtt{x}$). Lastly, if $\mathcal{T}_i$ does not write to $\mathtt{x}_j$ and it reads the value of $\mathtt{x}_j$ written by the initial write, $mathitinit_\mathtt{x}$, then $ir_i^{\mathtt{x}_j}$ and $vr_i^{\mathtt{x}_j}$ read the value written to $\mathtt{vx}_j$ by the initial write to vx, $init_{\mathtt{vx}}$.

$$\mathsf{IRF}_i^2 \triangleq \bigcup_{(\mathtt{x}, r) \in \mathtt{RS}_{\mathcal{T}_i}} \left\{ \begin{matrix}(w, r_i^\mathtt{x}),\\ (w', ir_i^\mathtt{x}),\\ (w', vr_i^\mathtt{x})\end{matrix} \;\middle|\; \begin{matrix}(w, r) \in \Gamma.\mathsf{rf} \\ \wedge (\mathtt{x} \in \mathtt{WS}_{\mathcal{T}_i} \Rightarrow w' = L_i^\mathtt{x}) \\ \wedge (\mathtt{x} \notin \mathtt{WS}_{\mathcal{T}_i} \wedge \exists \mathcal{T}_j.\ w \in \mathcal{T}_j \Rightarrow w' = U_j^\mathtt{x}) \\ \wedge (\mathtt{x} \notin \mathtt{WS}_{\mathcal{T}_i} \wedge w = init_\mathtt{x} \Rightarrow w' = init_{\mathtt{vx}})\end{matrix} \right\}$$

$$r_i^{\mathtt{x}_j} = \mathtt{R}(\mathtt{x}_j, v) \quad s_i^{\mathtt{x}_j} = \mathtt{W}(\mathtt{s[x}_j\mathtt{]}, v) \quad \text{s.t. } \exists w.\ (w, r_i^{\mathtt{x}_j}) \in \mathsf{IRF}_i^2 \wedge \mathtt{val_w}(w) = v$$

$$ir_i^{\mathtt{x}_j} = vr_i^{\mathtt{x}_j} = \mathtt{R}(\mathtt{vx}_j, v) \quad \text{s.t. } \exists w.\ (w, ir_i^{\mathtt{x}_j}) \in \mathsf{IRF}_i^2 \wedge \mathtt{val_w}(w) = v$$

- $Ts_i = t_1 \stackrel{\mathsf{po}}{\rightarrow} \cdots \stackrel{\mathsf{po}}{\rightarrow} t_n$ (when $Ts'_i = t'_1 \stackrel{\mathsf{po}}{\rightarrow} \cdots \stackrel{\mathsf{po}}{\rightarrow} t'_n$), with t_j defined as follows:

$$t_j = \mathtt{R}(\mathtt{s[x]}, v) \text{ when } t'_j = \mathtt{R}(\mathtt{x}, v)$$

$$t_j = \mathtt{W}(\mathtt{x}, v) \stackrel{\mathsf{po}|_{\mathrm{imm}}}{\rightarrow} \mathtt{W}(\mathtt{s[x]}, v) \text{ when } t'_j = \mathtt{W}(\mathtt{x}, v)$$

When t'_j is a read event, the t_j has the same identifier as that of t'_j. When t'_j is a write event, the first event in t_j has the same identifier as that of t_j and the identifier of the second event is picked fresh.

We are now in a position to construct our implementation graph. Given a consistent execution graph Γ of the specification, we construct an execution graph $G = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo})$ of the implementation as follows.

- $G.E = \bigcup_{\mathcal{T}_i \in \Gamma.\mathcal{T}/\mathsf{st}} \theta_i.E$. Observe that $G.E$ is an extension of $\Gamma.E$: $\Gamma.E \subseteq G.E$.
- $G.\mathsf{po}$ is defined as $\Gamma.\mathsf{po}$ extended by the po for the additional events of G, given by the θ_i traces defined above.
- $G.\mathsf{rf} = \bigcup_{\mathcal{T}_i \in \Gamma.\mathcal{T}/\mathsf{st}} (\mathsf{IRF}^1_i \cup \mathsf{IRF}^2_i)$
- $G.\mathsf{mo} = \Gamma.\mathsf{mo} \cup \left(\bigcup_{\mathcal{T}_i \in \Gamma.\mathcal{T}/\mathsf{st}} \mathsf{IMO}_i \right)^+$

Theorem 2 (Completeness). *For all PSI-consistent specification graphs Γ of a transactional program, there exists an RA-consistent execution graph G of the implementation in Fig. 1 that has the same program outcome.*

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

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

RTS. part 1
We proceed by contradiction. Let us assume that there exists $a, \mathcal{T}'_1$ such that $a \in \mathcal{T}'_1$ and $(a, a) \in G.\mathsf{hb}_{loc}$. From the auxiliary Lemma 6.2 in §1.4 we then have $(a, a) \in G.\mathsf{po}$, which is impossible given the construction of $G.\mathsf{po}$. This leads to a contradiction and we thus have $\mathsf{irreflexive}(G.\mathsf{hb}_{loc})$, as required.

RTS. part 2
We proceed by contradiction. Let us assume that there exist a, b such that $(a, b) \in G.\mathsf{mo}$ and $(b, a) \in G.\mathsf{hb}_{loc}$. From the auxiliaryLemma 6.3 in §1.4 we then have $(a, b) \in G.\mathsf{hb}$. As such, since $G.\mathsf{hb}$ is transitively closed, we have $(a, a) \in G.\mathsf{hb}_{loc}$. However, in the previous part (1) we demonstrate that $\forall a.\ (a, a) \notin G.\mathsf{hb}_{loc}$, resulting in a contradiction. We thus have $\mathsf{irreflexive}(G.\mathsf{mo}; G.\mathsf{hb}_{loc})$, as required.

RTS. part 3
We proceed by contradiction. Let us assume that there exist $\mathcal{T}'_1, \mathcal{T}'_2, a, b$ such that $a \in \mathcal{T}'_1$, $b \in \mathcal{T}'_2$, $(a, b) \in G.\mathsf{hb}_{loc}$ and $(b, a) \in G.\mathsf{rb}$. There are then two cases to consider: either 1) $\mathcal{T}'_1 = \mathcal{T}'_2$; or 2) $\mathcal{T}'_1 \neq \mathcal{T}'_2$.

In the former case (1), from the auxiliary Lemma 6.2 in §1.4 we know $(a, b) \in G.\mathsf{po}$. On the other hand, either a) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{x}$, for some shared location

$\mathtt{x}$; or b) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{vx}$, for some version lock $\mathtt{vx}$. In case (1.a), since b is a read event and a is a write event, given the structure of transaction we know $(b, a) \in G.\mathsf{po}$. In case (1.b) since a, b are both in the same transaction, given the construction of $G.\mathsf{rf}$, we know that $(b, a) \in G.\mathsf{po}$. As such, in both cases (1.a) and (1.b) we have $(a, b) \in G.\mathsf{po}$. Consequently, we have $a \stackrel{G.\mathsf{po}}{\to} b \stackrel{G.\mathsf{po}}{\to} a$. However, from the construction of G we know that $G.\mathsf{po}$ is acyclic. This leads to a contradiction and we thus have $\mathsf{irreflexive}(G.\mathsf{rb}; G.\mathsf{hb}_{loc})$, as required.

In the latter case (2), there are again two cases to consider: either a) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{x}$, for some shared location $\mathtt{vx}$; or b) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{vx}$, for a location lock $\mathtt{vx}$ associated with location $\mathtt{x}$. In the first case (2.a), from the construction of G we know that $a \in \mathcal{T}_1$ and there exist $c \in \mathcal{T}'_2$ such that $c \in \mathcal{T}_2$, $(b, c) \in G.\mathsf{po}$ and $(c, a) \in \Gamma.\mathsf{rb}$. On the other hand, since $(a, b) \in G.\mathsf{hb}$, $a \in \mathcal{T}'_1$, $b, c \in \mathcal{T}'_2$, $a \in \mathcal{T}_1$ and $c \in \mathcal{T}_2$, from the auxiliary Lemma 6.1 in §1.4 we have $(a, c) \in \Gamma.\mathsf{psi\text{-}hb}$. We thus have $a \stackrel{\Gamma.\mathsf{psi\text{-}hb}}{\to} c \stackrel{\Gamma.\mathsf{rb}}{\to} a$, contradicting our assumption that Γ is consistent.

In the second case (2.b), there are two final cases to consider: either i) $\mathtt{x} \notin \mathtt{WS}_{\mathcal{T}_2}$; or ii) $\mathtt{x} \in \mathtt{WS}_{\mathcal{T}_2}$. In case (2.b.i) from the construction of G we know there exist $w_{\mathtt{x}}, w'_{\mathtt{x}}, w_{\mathtt{vx}}, r$ such that $w'_{\mathtt{x}} \in \mathcal{T}_1$, $r \in \mathcal{T}_2$, $(w_{\mathtt{vx}}, b) \in G.\mathsf{rf}$, $(w_{\mathtt{x}}, r^{\mathtt{x}}_2) \in G.\mathsf{rf}$, $(w_{\mathtt{x}}, r) \in \Gamma.\mathsf{rf}$, $(w_{\mathtt{x}}, w'_{\mathtt{x}}) \in G.\mathsf{mo}$ and $(w_{\mathtt{x}}, w'_{\mathtt{x}}) \in \Gamma.\mathsf{mo}$. As such we have $(r^{\mathtt{x}}_2, w'_{\mathtt{x}}) \in G.\mathsf{rb}$ and $(r, w'_{\mathtt{x}}) \in \Gamma.\mathsf{rb}$. On the other hand, since $(a, b) \in G.\mathsf{hb}_{loc}$, from the auxiliary Lemma 6.1 in §1.4 we know that $(w'_{\mathtt{x}}, r) \in \Gamma.\mathsf{psi\text{-}hb}$. We thus have $w'_{\mathtt{x}} \stackrel{\Gamma.\mathsf{psi\text{-}hb}}{\to} r \stackrel{\Gamma.\mathsf{rb}}{\to} w'_{\mathtt{x}}$, contradicting our assumption that Γ is consistent.

Similarly, in case (2.b.ii) again from the construction of G we know there exist $w_{\mathtt{x}}, w'_{\mathtt{x}}, w_{\mathtt{vx}}, r$ such that $w'_{\mathtt{x}} \in \mathcal{T}_1$, $r \in \mathcal{T}_2$, $(w_{\mathtt{vx}}, L^{\mathtt{x}}_2) \in G.\mathsf{rf}$, $(w_{\mathtt{vx}}, L^{\mathtt{x}}_2) \in G.\mathsf{mo}|_{\text{imm}}$, $(w_{\mathtt{x}}, r^{\mathtt{x}}_2) \in G.\mathsf{rf}$, $(w_{\mathtt{x}}, r) \in \Gamma.\mathsf{rf}$, $(w_{\mathtt{x}}, w'_{\mathtt{x}}) \in G.\mathsf{mo}$ and $(w_{\mathtt{x}}, w'_{\mathtt{x}}) \in \Gamma.\mathsf{mo}$. As such we have $(r^{\mathtt{x}}_2, w'_{\mathtt{x}}) \in G.\mathsf{rb}$ and $(r, w'_{\mathtt{x}}) \in \Gamma.\mathsf{rb}$. On the other hand, since $(a, b) \in G.\mathsf{hb}_{loc}$, from the auxiliary Lemma 6.1 in §1.4 we know that $(w'_{\mathtt{x}}, r) \in \Gamma.\mathsf{psi\text{-}hb}$. We thus have $w'_{\mathtt{x}} \stackrel{\Gamma.\mathsf{psi\text{-}hb}}{\to} r \stackrel{\Gamma.\mathsf{rb}}{\to} w'_{\mathtt{x}}$, contradicting our assumption that Γ is consistent. □

1.3 Auxiliary Soundness Lemmata

Lemma 2. *For all specification graphs $\Gamma = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathcal{T})$ for all $a, b \in \Gamma.\mathcal{T}$:*

$$(a, b) \in \mathsf{psi\text{-}hb} \land [a]_{\mathsf{st}} \neq [b]_{\mathsf{st}} \Rightarrow [a]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq (\mathsf{po}_{\mathsf{T}} \cup \mathsf{rf}_{\mathsf{T}} \cup \mathsf{mo}_{\mathsf{T}})^+$$

Proof. Pick an arbitrary $\Gamma = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathcal{T})$. As $\mathsf{psi\text{-}hb}$ is a transitive closure, it is straightforward to demonstrate that $\mathsf{psi\text{-}hb} = \bigcup_{i \in \mathbb{N}} \mathsf{psi\text{-}hb}_i$, with $\mathsf{psi\text{-}hb}_0 = \mathsf{po} \cup \mathsf{rf} \cup \mathsf{rf}_{\mathsf{T}} \cup \mathsf{mo}_{\mathsf{T}}$ and for all $\mathsf{psi\text{-}hb}_{i+1} = \mathsf{psi\text{-}hb}_0; \mathsf{psi\text{-}hb}_i$. As such we demonstrate the following instead:

$$\begin{array}{l} \forall a, b \in \Gamma.\mathcal{T}.\ \forall i \in \mathbb{N}. \\ \quad (a, b) \in \mathsf{psi\text{-}hb}_i \land [a]_{\mathsf{st}} \neq [b]_{\mathsf{st}} \Rightarrow [a]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq (\mathsf{po}_{\mathsf{T}} \cup \mathsf{rf}_{\mathsf{T}} \cup \mathsf{mo}_{\mathsf{T}})^+ \end{array}$$

We proceed by induction on i.

Base case $i = 0$
Pick arbitrary $a, b \in \Gamma.\mathcal{T}$ such that $[a]_{\mathsf{st}} \neq [b]_{\mathsf{st}}$ and $(a, b) \in \mathsf{psi\text{-}hb}_0$. We then know that either i) $(a, b) \in \mathsf{po}$; or ii) $(a, b) \in \mathsf{rf}$; or iii) $(a, b) \in \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T}$. In (i) from the definition of po we have $[a]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq \mathsf{po}_\mathsf{T}$ and thus $[a]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+$, as required. In case (ii) from the definition of rf we have $[a]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq \mathsf{rf}_\mathsf{T}$ and thus $[a]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+$, as required. In case (iii) from the definitions of rf_T and mo_T we have $[a]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T}$ and thus $[a]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+$, as required.

Inductive case $i = n{+}1$

$$\begin{array}{l}\forall a, b \in \Gamma.\mathcal{T}.\ \forall j \leq n. \\ \quad (a, b) \in \mathsf{psi\text{-}hb}_i \land [a]_{\mathsf{st}} \neq [b]_{\mathsf{st}} \Rightarrow [a]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+\end{array} \tag{I.H.}$$

Pick arbitrary $a, b \in \Gamma.\mathcal{T}$ such that $[a]_{\mathsf{st}} \neq [b]_{\mathsf{st}}$ and $(a, b) \in \mathsf{psi\text{-}hb}_i$. From the definition of $\mathsf{psi\text{-}hb}_i$ We then know there exists c such that $(a, c) \in \mathsf{psi\text{-}hb}_0$ and $(c, b) \in \mathsf{psi\text{-}hb}_n$. There are three cases to consider: i) $[a]_{\mathsf{st}} = [c]_{\mathsf{st}}$; or ii) $[b]_{\mathsf{st}} = [c]_{\mathsf{st}}$; or iii) $[a]_{\mathsf{st}} \neq [c]_{\mathsf{st}}$ and $[b]_{\mathsf{st}} \neq [c]_{\mathsf{st}}$. In case (i) from (I.H.) we have $[a]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+$, as required. In case (ii) from the proof of the base case we have $[a]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+$, as required. In case (iii) from the proof of the base case we have $[a]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+$. Similarly, from (I.H.) we have $[b]_{\mathsf{st}} \times [c]_{\mathsf{st}} \subseteq (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+$. From the definition of transitive closures we thus have $[a]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+$, as required. □

Lemma 3. *For all specification graphs* $\Gamma{=}(E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathcal{T})$ *where* $\mathsf{irreflexive}((\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+; \mathsf{rb}_\mathsf{T}{}^?) \land \mathsf{rf}_\mathrm{I} \cup \mathsf{mo}_\mathrm{I} \cup \mathsf{rb}_\mathrm{I} \subseteq \mathsf{po}$ *holds, for all* $a, b \in \Gamma.\mathcal{T}$*:*

$$(a, b) \in \mathsf{psi\text{-}hb} \land [a]_{\mathsf{st}} = [b]_{\mathsf{st}} \Rightarrow (a, b) \in \mathsf{po}$$

Proof. Pick an arbitrary $\Gamma = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathcal{T})$ such that $\mathsf{irreflexive}((\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+; \mathsf{rb}_\mathsf{T}{}^?) \land \mathsf{rf}_\mathrm{I} \cup \mathsf{mo}_\mathrm{I} \cup \mathsf{rb}_\mathrm{I} \subseteq \mathsf{po}$ holds and $\mathsf{rf}_\mathrm{I} \subseteq \mathsf{po}$ holds. As $\mathsf{psi\text{-}hb}$ is a transitive closure, it is straightforward to demonstrate that $\mathsf{psi\text{-}hb} = \bigcup_{i \in \mathbb{N}} \mathsf{psi\text{-}hb}_i$, with $\mathsf{psi\text{-}hb}_0 = \mathsf{po} \cup \mathsf{rf} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T}$ and for all $\mathsf{psi\text{-}hb}_{i+1} = \mathsf{psi\text{-}hb}_0; \mathsf{psi\text{-}hb}_i$. As such we demonstrate the following instead:

$$\begin{array}{l}\forall a, b \in \Gamma.\mathcal{T}.\ \forall i \in \mathbb{N}. \\ \quad (a, b) \in \mathsf{psi\text{-}hb}_i \land [a]_{\mathsf{st}} = [b]_{\mathsf{st}} \Rightarrow (a, b) \in \mathsf{po}\end{array}$$

We proceed by induction on i.

Base case $i = 0$
Pick arbitrary $a, b \in \Gamma.\mathcal{T}$ such that $[a]_{\mathsf{st}} = [b]_{\mathsf{st}}$ and $(a, b) \in \mathsf{psi\text{-}hb}_0$. From the definition of $\mathsf{psi\text{-}hb}$ we know that either $(a, b) \in \mathsf{po}$ or $(a, b) \in \mathsf{rf}$. In the first case the desired result holds immediately. In the latter case we then have $(a, b) \in \mathsf{rf}_\mathrm{I}$ and thus from the assumption of the lemma we have $(a, b) \in \mathsf{po}$, as required.

Inductive case $i = n{+}1$

$$\begin{array}{l}\forall a, b \in \Gamma.\mathcal{T}.\ \forall j \leq n. \\ \quad (a, b) \in \mathsf{psi\text{-}hb}_i \land [a]_{\mathsf{st}} = [b]_{\mathsf{st}} \Rightarrow (a, b) \in \mathsf{po}\end{array} \tag{I.H.}$$

Pick arbitrary $a, b \in \Gamma.\mathcal{T}$ such that $[a]_{\mathsf{st}} = [b]_{\mathsf{st}}$ and $(a, b) \in \mathsf{psi\text{-}hb}_i$. From the definition of $\mathsf{psi\text{-}hb}_i$ We then know there exists c such that $(a, c) \in \mathsf{psi\text{-}hb}_0$ and $(c, b) \in \mathsf{psi\text{-}hb}_n$. There are two cases to consider: i) $[a]_{\mathsf{st}} = [c]_{\mathsf{st}}$; or ii) $[a]_{\mathsf{st}} \neq [c]_{\mathsf{st}}$. In case (i) from the proof of the base case we have $(a, c) \in \mathsf{po}$. Similarly, from (I.H.) we have $(c, b) \in \mathsf{po}$. As po is transitively closed, we have $(a, b) \in \mathsf{po}$ as required. In case (ii) from Lemma 2 we have $[a]_{\mathsf{st}} \times [c]_{\mathsf{st}} \subseteq (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+$ and $[c]_{\mathsf{st}} \times [b]_{\mathsf{st}} \subseteq (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+$, and thus from the definition of transitive closures we have $[a]_{\mathsf{st}} \times [a]_{\mathsf{st}} \subseteq (\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+$, contradicting the assumption that $\mathsf{irreflexive}((\mathsf{po}_\mathsf{T} \cup \mathsf{rf}_\mathsf{T} \cup \mathsf{mo}_\mathsf{T})^+)$ holds. □

Lemma 4. *For all consistent execution graphs of the implementation $G = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo})$ and its transaction set* Tx*, for all version lock locations* $\mathtt{vx}$*, and all transaction subsets* $\mathrm{Tx}_{\mathtt{vx}} \subseteq \mathrm{Tx}$ *with* $\mathtt{vx}$ *in their write sets* ($\forall \xi \in \mathrm{Tx}_{\mathtt{vx}}.\ \mathtt{x} \in \mathtt{WS}_\xi$)*:*

1. *there exists* $L = [\xi_1 \cdots \xi_m] = \mathtt{perm}(\mathrm{Tx}_{\mathtt{vx}})$*, such that:*

$$\xi_1.L_{\mathtt{vx}} \stackrel{\mathsf{mo}|_{imm}}{\rightarrow} \xi_1.U_{\mathtt{vx}} \stackrel{\mathsf{mo}|_{imm}}{\rightarrow} \cdots \stackrel{\mathsf{mo}|_{imm}}{\rightarrow} \xi_m.L_{\mathtt{vx}} \stackrel{\mathsf{mo}|_{imm}}{\rightarrow} \xi_m.U_{\mathtt{vx}}$$

   where $\xi_i.L_{\mathtt{vx}}$ denotes the event corresponding to the successful acquisition of the $\mathtt{vx}$ *lock in transaction ξ_i, and $\xi_i.U_{\mathtt{vx}}$ denotes the unlocking of* $\mathtt{vx}$ *in ξ_i (i.e. $\xi_i.L_{\mathtt{vx}} = \mathtt{U}(\mathtt{vx}, a, a{+}1)$ and $\xi_i.U_{\mathtt{vx}} = \mathtt{W}(\mathtt{vx}, a{+}2)$, for some a such that $a \bmod 2 = 0$).*
2. *for all $\xi_1, \xi_2 \in \mathrm{Tx}_{\mathtt{vx}}$, if $\xi_1 \neq \xi_2$, then either $\xi_1.U_{\mathtt{vx}} \stackrel{\mathsf{hb}}{\rightarrow} \xi_2.L_{\mathtt{vx}}$, or $\xi_2.U_{\mathtt{vx}} \stackrel{\mathsf{hb}}{\rightarrow} \xi_1.L_{\mathtt{vx}}$.*
3. *each write event to location* $\mathtt{vx}$ *in E, writes a unique value:*

$$\forall a, b \in G.\mathcal{W}_{\mathtt{vx}}.\ \mathtt{val}_{\mathtt{w}}(a) \neq \mathtt{val}_{\mathtt{w}}(b)$$

Proof (part 1). By induction on the length of $\mathrm{Tx}_{\mathtt{vx}}$.

Base case $\mathrm{Tx}_{\mathtt{vx}} = \{\}$.
This case holds vacuously.

Inductive case $|\mathrm{Tx}_{\mathtt{vx}}| = m$, where $m \geq 1$.
Given the trace of each transaction described above, we know that the set of write events on $\mathtt{vx}$ is given by $\mathcal{W}_{\mathtt{vx}} = \bigcup_{\xi_i \in \mathrm{Tx}_{\mathtt{vx}}} \{\xi_i.L_{\mathtt{vx}}, \xi_i.U_{\mathtt{vx}}\}$. Since the write events of $\mathtt{vx}$ are totally ordered by mo, we know there exists a minimal $e_0 \in \mathcal{W}_{\mathtt{vx}}$ such that $\forall e \in \mathcal{W}_{\mathtt{vx}} \setminus \{e_0\}.\ e_0 \stackrel{\mathsf{mo}}{\rightarrow} e$. That is, there exists $\xi_i \in \mathrm{Tx}_{\mathtt{vx}}$ such that either $e_0 = \xi_i.L_{\mathtt{vx}}$ or $e_0 = \xi_i.U_{\mathtt{vx}}$. Let us assume that $e_0 = \xi_i.U_{\mathtt{vx}}$; we then have $\xi_i.U_{\mathtt{vx}} \stackrel{\mathsf{mo}}{\rightarrow} \xi_i.L_{\mathtt{vx}}$. On the other hand, since we have $\xi_i.L_{\mathtt{vx}} \stackrel{\mathsf{po}}{\rightarrow} \xi_i.U_{\mathtt{vx}}$, we have $\xi_i.L_{\mathtt{vx}} \stackrel{\mathsf{po}}{\rightarrow} \xi_i.U_{\mathtt{vx}} \stackrel{\mathsf{mo}}{\rightarrow} \xi_i.L_{\mathtt{vx}}$, contradicting the assumption that G is consistent. We thus know that the minimal element is $e_0 = \xi_i.L_{\mathtt{vx}}$ for some $\xi_i \in \mathrm{Tx}_{\mathtt{vx}}$.

From the totality of mo on $\mathcal{W}_{\mathtt{vx}}$, we know that there exists $e_1 \in \mathcal{W}_{\mathtt{vx}} \setminus \{e_0\}$ such that $e_0 \stackrel{\mathsf{mo}|_{\mathrm{imm}}}{\rightarrow} e_1$. That is, either $e_1 = \xi_i.U_{\mathtt{vx}}$; or there exists $j \neq i$ such that $e_1 = \xi_j.L_{\mathtt{vx}}$ or $e_1 = \xi_i.U_{\mathtt{vx}}$. Let us pick an arbitrary $j \neq i$ and assume that $e_1 = \xi_j.L_{\mathtt{vx}}$. Since $e_0 \stackrel{\mathsf{mo}|_{\mathrm{imm}}}{\rightarrow} e_1$, the value read by $e_1 = \xi_j.L_{\mathtt{vx}}$, must be that written by $e_0 = \xi_i.L_{\texttt{vx}}$. However, the value written by e_0 is an odd number, whilst the value read by e_0 is an even number. We thus know that $e_1 \neq \xi_j.L_{\texttt{vx}}$ for all $j \neq i$. Similarly, let us pick an arbitrary $j \neq i$ and assume that $e_1 = \xi_j.U_{\texttt{vx}}$. We then have $\xi_j.U_{\texttt{vx}} \xrightarrow{\textsf{mo}} \xi_j.L_{\texttt{vx}}$. On the other hand, since we have $\xi_j.L_{\texttt{vx}} \xrightarrow{\textsf{po}} \xi_j.U_{\texttt{vx}}$, we have $\xi_j.L_{\texttt{vx}} \xrightarrow{\textsf{po}} \xi_j.U_{\texttt{vx}} \xrightarrow{\textsf{mo}} \xi_j.L_{\texttt{vx}}$, contradicting the assumption that G is consistent. We thus know that $e_1 \neq \xi_j.U_{\texttt{vx}}$ for all $j \neq i$. Consequently we have $e_1 = \xi_i.U_{\texttt{vx}}$.

Let $\text{Tx'}_{\texttt{vx}} = \text{Tx'}_{\texttt{vx}} \setminus \{\xi_i\}$. From the inductive hypothesis we then know there exist $L' = \texttt{perm}(\text{Tx'}_{\texttt{vx}})$ such that

$$L'|_1 .L_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} L'|_1 .U_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \cdots \xrightarrow{\textsf{mo}|_{\text{imm}}} L'|_{|L'|} .L_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} L'|_{|L'|} .U_{\texttt{vx}}$$

where $L'|_i$ denotes the i^{th} element of L'. On the other hand, since we have $e_0{=}\xi_i.L_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} e_1{=}\xi_i.U_{\texttt{vx}}$ and e_0 is the minimal element according to mo, we then have:

$$\begin{array}{l} \xi_i.L_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \xi_i.U_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \\ L'|_1 .L_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} L'|_1 .U_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \cdots \xrightarrow{\textsf{mo}|_{\text{imm}}} L'|_{|L'|} .L_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} L'|_{|L'|} .U_{\texttt{vx}} \end{array}$$

as required.

Proof (part 2). From part 1 we know there exists L, i, j such that $L|_1 .L_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} L|_1 .U_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \cdots \xrightarrow{\textsf{mo}|_{\text{imm}}} L|_{|L|} .L_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} L|_{|L|} .U_{\texttt{vx}}$ and $L|_i = \xi_1,\ L|_j = \xi_2$, ans either $i < j$ or $j < i$.

Let us assume the former case. Since each $U_{\texttt{vx}}$ event is a `rel` write event and each $L_{\texttt{vx}}$ event is an `acqrel` update event, we have $\cdots \xi_1.U_{\texttt{vx}} \xrightarrow{\textsf{rf}} L|_{i+1} .L_{\texttt{vx}} \xrightarrow{\textsf{po}} L|_{i+1} .U_{\texttt{vx}} \xrightarrow{\textsf{rf}} \cdots \xrightarrow{\textsf{rf}} \xi_2.L_{\texttt{vx}}$. On the other hand, since $\textsf{hb} = (\textsf{po} \cup \textsf{rf})^+$, we have $\xi_1.U_{\texttt{vx}} \xrightarrow{\textsf{hb}} \xi_2.L_{\texttt{vx}}$ as required. The proof of the latter case in analogous and is omitted here.

Proof (part 3). From part 1 we know that the write events in $G.\mathcal{W}_{\texttt{vx}}$ are ordered by mo as follows, where $L = [\xi_1 \cdots \xi_m] = \texttt{perm}(\text{Tx}_{\texttt{vx}})$:

$$\xi_1.L_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \xi_1.U_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \cdots \xrightarrow{\textsf{mo}|_{\text{imm}}} \xi_m.L_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \xi_m.U_{\texttt{vx}}$$

As such, the values written to `vx` by the write events ordered as above monotonically increase: each $\xi_i.L_{\texttt{vx}}$ event increments the value of `vx` by one (it updates `vx` from v to $v{+}1$); while each subsequent $\xi_i.U_{\texttt{vx}}$ event increments the value of `vx` by one (it updates `vx` from $v{+}1$ to $v{+}2$). Consequently, each value written by the write events ordered above is unique. □

Lemma 5. *For all consistent implementation execution graphs G and their counterpart specification graph Γ constructed as above,*

1. $\Gamma.\textsf{po} \subseteq G.\textsf{po}$

2. $\Gamma.\mathsf{rf}_{\mathsf{T}} \cup \Gamma.\mathsf{mo}_{\mathsf{T}} \subseteq G.\mathsf{hb}$
3. $\Gamma.\mathsf{mo} \subseteq G.\mathsf{hb}$
4. $\Gamma.\mathsf{rf} \subseteq G.\mathsf{hb}$
5. $\Gamma.\mathsf{psi\text{-}hb} \subseteq G.\mathsf{hb}$

Proof (Part 1). Immediate from the definitions of $\Gamma.\mathsf{po}$ and $G.\mathsf{po}$.

Proof (Part 2). In what follows we demonstrate that $\Gamma.\mathsf{mo}_{\mathsf{T}} \subseteq G.\mathsf{hb}$ and $\Gamma.\mathsf{rf}_{\mathsf{T}} \subseteq G.\mathsf{hb}$.

RTS. $\Gamma.\mathsf{rf}_{\mathsf{T}} \subseteq G.\mathsf{hb}$
Pick an arbitrary $(a, b) \in \Gamma.\mathsf{rf}_{\mathsf{T}}$; we are then required to show that $(a, b) \in G.\mathsf{hb}$.

From the definition of $\Gamma.\mathsf{rf}_{\mathsf{T}}$ and the construction of Γ we know there exist ξ_1, ξ_2, w, r such that $\xi_1 \neq \xi_2$, $(w, r) \in \Gamma.\mathsf{rf}$ $a, w \in \xi_1$ and $b, r \in \xi_2$. Let $\mathtt{loc}(w) = \mathtt{loc}(r) = \mathtt{x}$. We then know $\xi_1.L_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} w \stackrel{G.\mathsf{po}}{\to} \xi_1.U_{\mathtt{vx}}$, and $a \stackrel{G.\mathsf{po}^*}{\to} \xi_1.U_{\mathtt{vx}}$.

Let $w = \mathtt{W}(\mathtt{x}, v)$ and $r = \mathtt{R}(\mathtt{x}, v)$. From the construction of Γ we know there exists $rv_1, rx, rv_2 \in G.E$ and d such that $rv_1 = \mathtt{R}(\mathtt{vx}, d)$, $rx = \mathtt{R}(\mathtt{x}, v)$, $rv_2 = \mathtt{R}(\mathtt{vx}, d)$, $(w, rx) \in G.\mathsf{rf}$, $rv_2 \stackrel{\mathsf{po}}{\to} b$, and $rv_1 \stackrel{G.\mathsf{po}}{\to} rx \stackrel{G.\mathsf{po}}{\to} rv_2 \stackrel{G.\mathsf{po}}{\to} r$. There are two cases to consider: A) either $\mathtt{x} \in \mathtt{WS}_{\xi_2}$; or B) $\mathtt{x} \notin \mathtt{WS}_{\xi_2}$.

In the former case (A), from Lemma 4.2 we then know that either i) $\xi_2.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_1.L_{\mathtt{vx}}$; or ii) $\xi_1.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_2.L_{\mathtt{vx}}$. In case (A.i) we then have $\xi_2.U_{\mathtt{vx}} \stackrel{G.\mathsf{mo}}{\to} \xi_1.L_{\mathtt{vx}}$ (since otherwise we would have a cycle $\xi_2.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_1.L_{\mathtt{vx}} \stackrel{G.\mathsf{mo}}{\to} \xi_2.U_{\mathtt{vx}}$, contradicting our assumption that G is consistent). As such we have $(rv_2, \xi_1.L_{\mathtt{vx}}) \in G.\mathsf{rb}$.We then have $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.L_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} w \stackrel{G.\mathsf{rf}}{\to} rx \stackrel{G.\mathsf{po}}{\to} rv_2$. As $G.\mathsf{rf} \subseteq \mathsf{hb}$ and $G.\mathsf{po} \subseteq G.\mathsf{hb}$, we then have $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.L_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} rv_2$, contradicting the assumption that G is consistent.

In case (A.ii) we then have $a \stackrel{G.\mathsf{po}^*}{\to} \xi_1.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_2.L_{\mathtt{vx}} \stackrel{G.\mathsf{po}^*}{\to} b$. That is, since we have $G.\mathsf{po} \subseteq G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we have $a \stackrel{G.\mathsf{hb}}{\to} b$, as required.

In the latter case (B) we then know b (in $rv_1 = \mathtt{R}(\mathtt{vx}, b)$) is even. Additionally, since write events on $\mathtt{vx}$ have unique values, we know that either i) rv_1 reads from the initial write to $\mathtt{vx}$ and we thus have $rv_1 \stackrel{G.\mathsf{rb}}{\to} \xi_1.L_{\mathtt{vx}}$ and $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.L_{\mathtt{vx}}$; or ii) there exists ξ_3 such that $\mathtt{x} \in \mathtt{WS}_{\xi_3}$, $\xi_3.L_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} \xi_3.U_{\mathtt{vx}}$ and $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{rf}}{\to} rv_1$.

In case (B.i) we have $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.L_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} w \stackrel{G.\mathsf{rf}}{\to} rx \stackrel{G.\mathsf{po}}{\to} rv_2$. As $G.\mathsf{rf} \subseteq \mathsf{hb}$ and $G.\mathsf{po} \subseteq G.\mathsf{hb}$, we then have $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.L_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} rv_2$, contradicting the assumption that G is consistent.

In case (B.ii), since we have $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{rf}}{\to} rv_1$ and each write event on $\mathtt{vx}$ writes a unique value (Lemma 4.3), we also have $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{rf}}{\to} rv_2$. On the other hand, from Lemma 4.2 we know that either a) $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_1.L_{\mathtt{vx}}$; or b) $\xi_1.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_3.L_{\mathtt{vx}}$.

In case (B.ii.a), since $G.\mathsf{mo}$ on $\mathtt{vx}$ is totally ordered, from the consistency of Γ we know that $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{mo}}{\to} \xi_1.L_{\mathtt{vx}}$ (since otherwise we would have a cycle $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_1.L_{\mathtt{vx}} \stackrel{G.\mathsf{mo}}{\to} \xi_3.U_{\mathtt{vx}}$, contradicting $\mathsf{RA\text{-}consistent}(G)$). Consequently, since we have $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{rf}}{\to} rv_2$, and $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{mo}}{\to} \xi_1.L_{\mathtt{vx}}$, we have $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.L_{\mathtt{vx}}$. We thus have $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.L_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} w \stackrel{G.\mathsf{rf}}{\to} rx \stackrel{G.\mathsf{po}}{\to} rv_2$. As $G.\mathsf{rf} \subseteq \mathsf{hb}$ and $G.\mathsf{po} \subseteq G.\mathsf{hb}$, we have $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.L_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} rv_2$, contradicting the assumption that G is consistent.

In case (B.ii.b) we have $\xi_1.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_3.L_{\mathtt{vx}}$. Recall that we also have $a \stackrel{\mathsf{po}}{\to} \xi_1.U_{\mathtt{vx}}$, $\xi_3.L_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} \xi_3.U_{\mathtt{vx}}$, $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{rf}}{\to} rv_2$, and $rv_2 \stackrel{\mathsf{po}}{\to} b$. As $G.\mathsf{po}, G.\mathsf{rf} \in G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we thus have $a \stackrel{G.\mathsf{hb}}{\to} b$, as required.

RTS. $\Gamma.\mathsf{mo}_{\mathsf{T}} \subseteq G.\mathsf{hb}$
Pick an arbitrary $(a, b) \in \Gamma.\mathsf{mo}_{\mathsf{T}}$; we are then required to show that $(a, b) \in G.\mathsf{hb}$.

From the definition of $\Gamma.\mathsf{mo}_{\mathsf{T}}$ and the construction of Γ we know there exist ξ_1, ξ_2, c, d such that $\xi_1 \neq \xi_2$, $(c, d) \in \Gamma.\mathsf{mo}$, $a, c \in \xi_1$, $b, d \in \xi_2$. Let $\mathtt{loc}(c) = \mathtt{loc}(d) = \mathtt{x}$. We then know $a \stackrel{G.\mathsf{po}^*}{\to} \xi_1.U_{\mathtt{vx}}$, $\xi_1.L_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} c \stackrel{G.\mathsf{po}}{\to} \xi_1.U_{\mathtt{vx}}$, $\xi_2.L_{\mathtt{vx}} \stackrel{G.\mathsf{po}^*}{\to} b$ and $\xi_2.L_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} d \stackrel{G.\mathsf{po}}{\to} \xi_2.U_{\mathtt{vx}}$.

From Lemma 4.2 we then know that either $\xi_1.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_2.L_{\mathtt{vx}}$, or $\xi_2.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_1.L_{\mathtt{vx}}$. Let us assume that the latter holds. We then have $d \stackrel{G.\mathsf{po}}{\to} \xi_2.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_1.L_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} c \stackrel{G.\mathsf{mo}}{\to} d$. That is, since $G.\mathsf{po} \in G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we have $d \stackrel{G.\mathsf{hb}}{\to} c \stackrel{G.\mathsf{mo}}{\to} d$, contradicting the assumption that G is consistent. We thus know that $\xi_1.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_2.L_{\mathtt{vx}}$. As such, we have $a \stackrel{G.\mathsf{po}^*}{\to} \xi_1.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_2.L_{\mathtt{vx}} \stackrel{G.\mathsf{po}^*}{\to} b$. As $G.\mathsf{po} \in G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we have $a \stackrel{G.\mathsf{hb}}{\to} b$, as required.

Proof (Part 3). Pick an arbitrary $(a, b) \in \Gamma.\mathsf{mo}$ and let $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{x}$. There are then two cases to consider: either $[a]_{\mathsf{st}} = [b]_{\mathsf{st}}$, or $[a]_{\mathsf{st}} \neq [b]_{\mathsf{st}}$. In the latter case we then have $(a, b) \in \Gamma.\mathsf{mo}_{\mathsf{T}}$ and thus from part 2 we have $(a, b) \in G.\mathsf{hb}$, as required.

Now let us assume that $[a]_{\mathsf{st}} = [b]_{\mathsf{st}}$. From the construction of Γ we know there exists ξ such that $a, b \in \xi$, $(a, b) \in G.\mathsf{mo}$, and either $\xi.L_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} a \stackrel{G.\mathsf{po}}{\to} b \stackrel{G.\mathsf{po}}{\to} \xi.U_{\mathtt{vx}}$, or $\xi.L_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} b \stackrel{G.\mathsf{po}}{\to} a \stackrel{G.\mathsf{po}}{\to} \xi.U_{\mathtt{vx}}$. Let us assume that the latter case holds. We then have $a \stackrel{G.\mathsf{mo}}{\to} b \stackrel{G.\mathsf{po}}{\to} a$, contradicting the assumption that G is consistent. On the other hand, when the former case holds we have $a \stackrel{G.\mathsf{po}}{\to} b$ and thus $a \stackrel{G.\mathsf{hb}}{\to} b$, as required.

Proof (part 4). Pick an arbitrary $(w, r) \in \Gamma.\mathsf{rf}$ and let $\mathtt{loc}(w) = \mathtt{loc}(r) = \mathtt{x}$. There are then two cases to consider: either $[w]_{\mathsf{st}} = [r]_{\mathsf{st}}$, or $[w]_{\mathsf{st}} \neq [r]_{\mathsf{st}}$. In the latter case we then have $(w, r) \in \Gamma.\mathsf{rf}_{\mathsf{T}}$ and thus from part 2 we have $(w, r) \in G.\mathsf{hb}$, as required. Now let us assume that $[w]_{\mathsf{st}} = [r]_{\mathsf{st}}$ and let $r = \mathtt{R}(\mathtt{x}, v)$.

Proof (Part 5). Immediate from parts 1, 2 and 4. 'qed

1.4 Auxiliary Completeness Lemmata

In what follows, we write $\mathcal{T}'_i$ for the set of events in the *implementation* trace θ_i; that is, $\mathcal{T}'_i \triangleq \theta_i.E$. In other words, $\mathcal{T}'_i$ corresponds to the set of events in the implementation of the specification transaction class $\mathcal{T}_i$.

Lemma 6. *For all consistent specification execution graphs Γ and their counterpart implementation graphs G constructed as above,*

1. *for all $\mathcal{T}'_1, \mathcal{T}'_2$ and for all a, b:*

$$\mathcal{T}'_1 \neq \mathcal{T}'_2 \land a \in \mathcal{T}'_1 \land b \in \mathcal{T}'_2 \land (a, b) \in G.\mathsf{hb} \Rightarrow (\mathcal{T}_1 \times \mathcal{T}_2) \subseteq \Gamma.\mathsf{psi\text{-}hb}$$

2. *for all $\mathcal{T}'_i$ and for all a, b:*

$$a, b \in \mathcal{T}'_i \land (a, b) \in G.\mathsf{hb} \Rightarrow (a, b) \subseteq G.\mathsf{po}$$

3. $G.\mathsf{mo} \subseteq G.\mathsf{hb}$

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

$$\begin{array}{l} \forall i \in \mathbb{N}.\ \forall \mathcal{T}'_1, \mathcal{T}'_2.\ \forall a, b. \\ \quad \mathcal{T}'_1 \neq \mathcal{T}'_2 \land a \in \mathcal{T}'_1 \land b \in \mathcal{T}'_2 \land (a, b) \in G.\mathsf{hb}_i \Rightarrow (\mathcal{T}_1 \times \mathcal{T}_2) \subseteq \Gamma.\mathsf{psi\text{-}hb} \end{array}$$

Base case $i = 0$

Pick arbitrary $\mathcal{T}'_1, \mathcal{T}'_2$ and a, b, c, d such that $\mathcal{T}'_1 \neq \mathcal{T}'_2$, $a \in \mathcal{T}'_1$, $b \in \mathcal{T}'_2$, $(a, b) \in G.\mathsf{hb}_0$, $c \in \mathcal{T}_1$ and $d \in \mathcal{T}_2$. We are then required to show that $(c, d) \in \Gamma.\mathsf{hb}$. Observe from the construction above that $\mathcal{T}'_i \supseteq \mathcal{T}_i$ for all i. As such, we know that $c \in \mathcal{T}'_1$ and $d \in \mathcal{T}'_2$. As $(a, b) \in G.\mathsf{hb}_0$, there are two cases to consider: either A) $(a, b) \in G.\mathsf{po}$, or B) $(a, b) \in G.\mathsf{rf}$.

In case (A), since $\mathcal{T}'_1 \neq \mathcal{T}'_2$, $a, c \in \mathcal{T}'_1$, $b, d \in \mathcal{T}'_2$ and $(a, b) \in G.\mathsf{po}$, from the construction of G we thus know that $(c, d) \in G.\mathsf{po}$. As $G.\mathsf{po}$ does not introduce additional orderings between events of $\Gamma.E$ ($\forall e, f \in \Gamma.E.\ (e, f) \in \Gamma \Leftrightarrow (e, f) \in G.\mathsf{po}$), we thus know that $(c, d) \in \Gamma.\mathsf{po}$. As such, from the definition of $\Gamma.\mathsf{psi\text{-}hb}$ we have $(a, b) \in \Gamma.\mathsf{psi\text{-}hb}$, as required.

In case (B), there are two cases to consider: 1) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{x}$, for some shared location $\mathtt{x}$; or 2) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{vx}$, for some version lock $\mathtt{vx}$ associated with location $\mathtt{x}$. In case (1) from the construction of $G.\mathsf{rf}$ we know $a \in \mathcal{T}_1$ and that there exists $e \in \mathcal{T}_2$ such that $e \in \mathcal{T}'_2$, $(b, e) \in G.\mathsf{po}$ and that $(a, e) \in \Gamma.\mathsf{rf}$. That is, $(\mathcal{T}_1 \times \mathcal{T}_2) \subseteq \Gamma.\mathsf{rf}_\mathsf{T} \subseteq \Gamma.\mathsf{psi\text{-}hb}$. We thus have $(c, d) \in \Gamma.\mathsf{psi\text{-}hb}$, as required.

In case (2) from the construction of $G.\mathsf{rf}$ we know that there are two possible cases: i) either $\mathtt{x} \notin \mathtt{WS}_{\mathcal{T}_2}$ ($\mathcal{T}_2$ merely reads from $\mathtt{x}$); or ii) $\mathtt{x} \in \mathtt{WS}_{\mathcal{T}_2}$. In case (2.i) from the construction of $\Gamma.\mathsf{rf}$ we know there exist $e \in \mathcal{T}'_1$, $f \in \mathcal{T}'_2$ such that $(e, f) \in G.\mathsf{rf}$. We can then use the same steps as in case (A) to demonstrate that $(c, d) \in \Gamma.\mathsf{psi\text{-}hb}$, as required. In case (2.ii) from the construction of G we know there exist $e \in \mathcal{T}'_1$ $f \in \mathcal{T}'_2$ such that $(e, f) \in G.\mathsf{mo}$, $e \in \mathcal{T}_1$, $f \in \mathcal{T}_2$. From the construction of $G.\mathsf{mo}$ we then have $(e, f) \in \Gamma.\mathsf{mo}$. That is, $(\mathcal{T}_1 \times \mathcal{T}_2) \subseteq \Gamma.\mathsf{mo}_\mathsf{T} \subseteq \Gamma.\mathsf{psi\text{-}hb}$. We thus have $(c, d) \in \Gamma.\mathsf{psi\text{-}hb}$, as required.

Inductive case $i = n{+}1$

$$\begin{array}{l}\forall i \leq n.\ \forall \mathcal{T}_1', \mathcal{T}_2'.\ \forall a, b. \\ \ \mathcal{T}_1' \neq \mathcal{T}_2' \land a \in \mathcal{T}_1' \land b \in \mathcal{T}_2' \land (a,b) \in G.\mathsf{hb}_i \Rightarrow (\mathcal{T}_1 \times \mathcal{T}_2) \subseteq \Gamma.\mathsf{psi\text{-}hb}\end{array} \tag{I.H.}$$

Pick arbitrary $\mathcal{T}_1', \mathcal{T}_2'$ and a, b such that $\mathcal{T}_1' \neq \mathcal{T}_2'$, $a \in \mathcal{T}_1'$, and $b \in \mathcal{T}_2'$, $(a,b) \in G.\mathsf{hb}_{n+1}$. Since $(a,b) \in \mathsf{hb}_{n+1}$, from the definition of hb_{n+1} we know there exist $e, \mathcal{T}_3'$ such that $e \in \mathcal{T}_3'$, $(a,e) \in \mathsf{hb}_0$ and $(e,b) \in \mathsf{hb}_n$. There are three cases to consider.

Case 1. $\mathcal{T}_3' = \mathcal{T}_1'$
We then have $e \in \mathcal{T}_1' \land b \in \mathcal{T}_2' \land (e,b) \in G.\mathsf{hb}_i$. Consequently, from (I.H.) we have $(\mathcal{T}_1 \times \mathcal{T}_2) \subseteq \Gamma.\mathsf{psi\text{-}hb}$, as required.

Case 2. $\mathcal{T}_3' = \mathcal{T}_2'$
We then have $a \in \mathcal{T}_1' \land e \in \mathcal{T}_2' \land (a,e) \in G.\mathsf{hb}_0$. From the proof of the base case we then have $(\mathcal{T}_1 \times \mathcal{T}_2) \subseteq \Gamma.\mathsf{psi\text{-}hb}$, as required.

Case 3. $\mathcal{T}_3' \neq \mathcal{T}_1' \land \mathcal{T}_3' \neq \mathcal{T}_2'$
We then have $a \in \mathcal{T}_1' \land e \in \mathcal{T}_3' \land (a,e) \in G.\mathsf{hb}_0 \land \mathcal{T}_3' \neq \mathcal{T}_1'$. From the proof of the base case we then have $(\mathcal{T}_1 \times \mathcal{T}_3) \subseteq \Gamma.\mathsf{psi\text{-}hb}$. On the other hand, we have $e \in \mathcal{T}_3' \land b \in \mathcal{T}_2' \land (e,b) \in G.\mathsf{hb}_i \land \mathcal{T}_3' \neq \mathcal{T}_2'$. Consequently, from (I.H.) we have $(\mathcal{T}_3 \times \mathcal{T}_2) \subseteq \Gamma.\mathsf{psi\text{-}hb}$. Since we have $(\mathcal{T}_1 \times \mathcal{T}_3) \subseteq \Gamma.\mathsf{psi\text{-}hb}$ and $(\mathcal{T}_3 \times \mathcal{T}_2) \subseteq \Gamma.\mathsf{psi\text{-}hb}$, and $\Gamma.\mathsf{psi\text{-}hb}$ is transitively closed, we have $(\mathcal{T}_1 \times \mathcal{T}_2) \subseteq \Gamma.\mathsf{psi\text{-}hb}$, as required.

Proof (Part 2). As in part 1 we show instead that the desired result holds for all hb_i as defined above. That is,

$$\begin{array}{l}\forall i \in \mathbb{N}.\ \forall \mathcal{T}_i'.\ \forall a, b. \\ \ a, b \in \mathcal{T}_i' \land (a,b) \in G.\mathsf{hb}_i \Rightarrow (a,b) \in G.\mathsf{po}\end{array}$$

Base case $i = 0$
Pick arbitrary $\mathcal{T}_i'$ and a, b such that $a, b \in \mathcal{T}_1'$ and $(a,b) \in G.\mathsf{hb}_0$. There are two cases to consider: either $(a,b) \in G.\mathsf{po}$, or $(a,b) \in G.\mathsf{rf}$.

In the former case, the desired result holds immediately. In the latter case, since a, b are both in $\mathcal{T}_i'$, from the construction of $G.\mathsf{rf}$ we know that there exists a version lock $\mathtt{vx}$ and a value v such that $b = \mathtt{R}(\mathtt{x}, v)$, $a = \mathtt{U}(\mathtt{x}, v{-}1, v)$ and that $(a,b) \in G.\mathsf{po}$, as required.

Inductive case $i = n{+}1$

$$\begin{array}{l}\forall i \leq n.\ \forall \mathcal{T}_i'.\ \forall a, b. \\ \ a, b \in \mathcal{T}_1' \land (a,b) \in G.\mathsf{hb}_i \Rightarrow (a,b) \subseteq G.\mathsf{po}\end{array} \tag{I.H.}$$

Pick arbitrary $\mathcal{T}_i'$ and a, b such that $a, b \in \mathcal{T}_i'$ and $(a,b) \in G.\mathsf{hb}_{n+1}$. Since $(a,b) \in \mathsf{hb}_{n+1}$, from the definition of hb_{n+1} we know there exist $e, \mathcal{T}_j'$ such that $e \in \mathcal{T}_j'$, $(a,e) \in \mathsf{hb}_0$ and $(e,b) \in \mathsf{hb}_n$. There are two cases to consider.

Case 1. $\mathcal{T}_i' = \mathcal{T}_j'$
We then have $a, e \in \mathcal{T}_i' \wedge (a, e) \in G.\mathsf{hb}_0$. As such, from the proof of the base case we have $(a, e) \subseteq G.\mathsf{po}$. Similarly, we have $e, b \in \mathcal{T}_i' \wedge (e, b) \in G.\mathsf{hb}_i$. Consequently, from (I.H.) we have $(e, b) \subseteq G.\mathsf{po}$. As we have $(a, e) \subseteq G.\mathsf{po}$. and $(e, b) \subseteq G.\mathsf{po}$ and $G.\mathsf{po}$ is transitively closed, we have $(a, b) \subseteq G.\mathsf{po}$, as required.

Case 2. $\mathcal{T}_i' \neq \mathcal{T}_j'$
We then have $a \in \mathcal{T}_i' \wedge e \in \mathcal{T}_j' \wedge (a, e) \in G.\mathsf{hb}_0$. From the proof of part 1 we then have $(\mathcal{T}_i \times \mathcal{T}_j) \subseteq \Gamma.\mathsf{psi\text{-}hb}$. Similarly, we have $e \in \mathcal{T}_j' \wedge b \in \mathcal{T}_i' \wedge (e, b) \in G.\mathsf{hb}_i$ and thus from part 1 we have $(\mathcal{T}_j \times \mathcal{T}_i) \subseteq \Gamma.\mathsf{psi\text{-}hb}$. Pick arbitrary $c \in \mathcal{T}_i$ and $d \in \mathcal{T}_j$ (from the construction of G we know that the $\mathcal{T}_i$ and $\mathcal{T}_j$ sets are non-empty and thus such c and d exist). We then have $c \stackrel{\Gamma.\mathsf{psi\text{-}hb}}{\rightarrow} d \stackrel{\Gamma.\mathsf{psi\text{-}hb}}{\rightarrow} c$, contradicting the assumption that Γ is consistent.

Proof (Part 3). Pick arbitrary $\mathcal{T}_1', \mathcal{T}_2'$ and a, b such that $a \in \mathcal{T}_1'$, $b \in \mathcal{T}_2'$, $(a, b) \in G.\mathsf{mo}$. There are then two cases to consider: 1) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{vx}$, for some location lock $\mathtt{vx}$ associated with location $\mathtt{vx}$; or 2) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{x}$, for some shared location $\mathtt{x}$.

In case (1), from the construction of $G.\mathsf{mo}$, $G.\mathsf{rf}$ and $G.\mathsf{po}$ we know that $a \stackrel{(G.\mathsf{po}^*;G.\mathsf{rf}^*)^+}{\rightarrow} b$. Since $G.\mathsf{po}, G.\mathsf{rf} \subseteq G.\mathsf{hb}$, and $G.\mathsf{hb}$ is transitively closed, we have $a \stackrel{G.\mathsf{hb}}{\rightarrow} b$, as required.

In case (2), given the construction of G we know that $(a, b) \in \Gamma.\mathsf{mo}$. There are again two cases to consider: a) $\mathcal{T}_1' = \mathcal{T}_2'$; or b) $\mathcal{T}_1' \neq \mathcal{T}_2'$. In case (2.a), since $\mathcal{T}_1' = \mathcal{T}_2'$, we know that either $(a, b) \in G.\mathsf{po}$ or $(b, a) \in G.\mathsf{po}$. In the former case, since $G.\mathsf{po} \subseteq G.\mathsf{hb}$, we have $(a, b) \in G.\mathsf{hb}$, as required. In the latter case, from the construction of $G.\mathsf{po}$ we know that $(b, a) \in \Gamma.\mathsf{po}$. As such, we have $a \stackrel{\Gamma.\mathsf{mo}}{\rightarrow} b \stackrel{\Gamma.\mathsf{po}}{\rightarrow} a$; that is, $a \stackrel{\Gamma.\mathsf{mo}}{\rightarrow} b \stackrel{\Gamma.\mathsf{hb}}{\rightarrow} a$, contradicting the assumption that Γ is consistent.

In case (2.b) from the construction of G we know that $a \stackrel{G.\mathsf{po}}{\rightarrow} U_1^{\mathtt{x}}$ and $L_2^{\mathtt{x}} \stackrel{G.\mathsf{po}}{\rightarrow} b$. Moreover, from the construction of $G.\mathsf{mo}$ we know $U_1^{\mathtt{x}} \stackrel{G.\mathsf{mo}}{\rightarrow} L_2^{\mathtt{x}}$. From the proof of case (1) we then know $U_1^{\mathtt{x}} \stackrel{G.\mathsf{hb}}{\rightarrow} L_2^{\mathtt{x}}$. That is, we have $a \stackrel{G.\mathsf{po}}{\rightarrow} U_1^{\mathtt{x}} \stackrel{G.\mathsf{hb}}{\rightarrow} L_2^{\mathtt{x}} \stackrel{G.\mathsf{po}}{\rightarrow} b$. As $G.\mathsf{po} \subseteq G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we have $(a, b) \in G.\mathsf{hb}$, as required. $\square$

2 Robust Parallel Snapshot Isolation (RPSI)

2.1 Implementation Soundness

The RPSI implementation in Fig. 1 is *sound*: for each consistent implementation graph G, a corresponding specification graph Γ with the same program outcome can be constructed such that $\textsf{rpsi-consistent}(\Gamma)$ holds.

Constructing Consistent Specification Graphs Constructing an RPSI-consistent specification graph from the implementation graph is similar to the corresponding PSI construction described in §1.1. More concretely, the events associated with non-transactional events remain unchanged and are simply added to the specification graph. On the other hand, the events associated with transactional events are adapted in a similar way to those of PSI in §1.1. In particular, observe that given an execution of the RPSI implementation with t transactions, as with the PSI implementation, the trace of each transaction $i \in \{1 \cdots t\}$ is of the form $\theta_i = Ls_i \overset{\textsf{po}}{\to} FS_i \overset{\textsf{po}}{\to} S_i \overset{\textsf{po}}{\to} Ts_i \overset{\textsf{po}}{\to} Us_i$, with Ls_i, FS_i, S_i, Ts_i and Us_i denoting analogous sequences of events to those of PSI. The difference between an RPSI trace θ_i and a PSI one is in the FS_i and S_i sequences, obtaining the snapshot. In particular, the validation phases of FS_i and S_i in RPSI include an additional read for each location to rule out intermediate non-transactional writes. As in the PSI construction, for each transactional trace θ_i of our implementation, we construct a corresponding trace of the specification as $\theta'_i = B_i \overset{\textsf{po}}{\to} Ts'_i \overset{\textsf{po}}{\to} E_i$, with B_i, E_i and Ts'_i as defined in §1.1.

Given a consistent RPSI implementation graph $G = (E, \textsf{po}, \textsf{rf}, \textsf{mo})$, let $G.\mathcal{NT} \triangleq G.E \setminus \bigcup_{i\in\{1\cdots t\}} \theta.E$ denote the non-transactional events of G. We construct a consistent RPSI specification graph $\Gamma = (E, \textsf{po}, \textsf{rf}, \textsf{mo}, \mathcal{T})$ such that:

- $\Gamma.E \triangleq G.\mathcal{NT} \cup \bigcup_{i\in\{1\cdots t\}} \theta'_i.E$ – the $\Gamma.E$ events comprise the non-transactional events in G and the events in each transactional trace θ'_i of the specification;
- $\Gamma.\textsf{po} \triangleq G.\textsf{po}|_{\Gamma.E}$ – the $\Gamma.\textsf{po}$ is that of $G.\textsf{po}$ restricted to the events in $\Gamma.E$;
- $\Gamma.\textsf{rf} \triangleq \bigcup_{i\in\{1\cdots t\}} \textsf{RF}_i \cup G.\textsf{rf}; [G.\mathcal{NT}]$ – the $\Gamma.\textsf{rf}$ is the union of $\textsf{RF}_i$ relations for transactional reads as defined in §1.1, together with the $G.\textsf{rf}$ relation for non-transactional reads;
- $\Gamma.\textsf{mo} \triangleq G.\textsf{mo}|_{\Gamma.E}$ – the $\Gamma.\textsf{mo}$ is that of $G.\textsf{mo}$ restricted to the events in $\Gamma.E$;
- $\Gamma.\mathcal{T} \triangleq \bigcup_{i\in\{1\cdots t\}} \theta'_i.E$, where for each $e \in \theta'_i.E$, we define $\texttt{tx}(e) = i$.

Theorem 3 (Soundness). *Let P be a program that possibly mixes transactional and non-transactional code. If every RPSI-consistent execution graph of P satisfies the condition in ($*$) below, then for all RA-consistent implementation graphs G of the implementation in Fig. 1, there exists an RPSI-consistent specification graph Γ of the corresponding transactional program with the same program outcome.*

$$\begin{array}{l} \forall \texttt{x}.\ \forall r \in \mathcal{T} \cap \mathcal{R}_{\texttt{x}}.\ \forall w, w' \in \mathcal{NT} \cap \mathcal{W}_{\texttt{x}}. \\ \quad w \neq w' \land \texttt{val}_{\texttt{w}}(w) = \texttt{val}_{\texttt{w}}(w') \land (r, w) \notin \textsf{rpsi-hb} \land (r, w') \notin \textsf{rpsi-hb} \\ \qquad \Rightarrow (w, r) \in \textsf{rpsi-hb} \land (w', r) \in \textsf{rpsi-hb} \end{array} \qquad (*)$$

Proof. Pick an arbitrary G such that $\mathsf{RA\text{-}consistent}(G)$, and its associated Γ constructed as described above.

RTS. $\mathsf{rpsi\text{-}consistent}(\Gamma)$
It is sufficient to establish that $\mathsf{irreflexive}(\Gamma.\mathsf{rpsi\text{-}hb})$, $\mathsf{irreflexive}(\Gamma.\mathsf{rpsi\text{-}hb}; \Gamma.\mathsf{rb})$ and $\mathsf{irreflexive}(\Gamma.\mathsf{rpsi\text{-}hb}; \Gamma.\mathsf{mo})$ hold.

RTS. $\mathsf{irreflexive}(\Gamma.\mathsf{rpsi\text{-}hb})$
We proceed by contradiction. Let us assume that there exists $(a, a) \in \Gamma.\mathsf{rpsi\text{-}hb}$. From auxiliary Lemma 8.5 in 2.3 we then have $(a, a) \in G.\mathsf{hb}$, contradicting our assumption that G is consistent.

RTS. $\mathsf{irreflexive}(\Gamma.\mathsf{rpsi\text{-}hb}; \Gamma.\mathsf{mo})$
We proceed by contradiction. Assume that there exists $(w, w) \in \Gamma.\mathsf{rpsi\text{-}hb}; \Gamma.\mathsf{mo}$. That is, there exist w' such that $(w, w') \in \Gamma.\mathsf{rpsi\text{-}hb}$ and $(w', w) \in \Gamma.\mathsf{mo}$. From Lemma 8.5 we then have $(w, w') \in G.\mathsf{hb}$. On the other hand, from the construction of G we know that $(w', w) \in G.\mathsf{mo}$. As such, we have $w \stackrel{G.\mathsf{hb}}{\rightarrow} w' \stackrel{G.\mathsf{mo}}{\rightarrow}$, contradicting the assumption that G is consistent.

RTS. $\mathsf{irreflexive}(\Gamma.\mathsf{rpsi\text{-}hb}; \Gamma.\mathsf{rb})$
We proceed by contradiction. Let us assume that there exists $(r, r) \in \Gamma.\mathsf{rpsi\text{-}hb}; \Gamma.\mathsf{rb}$. That is, there exist w such that $(w, r) \in \Gamma.\mathsf{rpsi\text{-}hb}$ and $(r, w) \in \Gamma.\mathsf{rb}$. From Lemma 8.5 we then have $(w, r) \in G.\mathsf{hb}$. There are then three cases to consider: 1) $r \in \Gamma.\mathcal{NT}$; or 2) $r, w \in \Gamma.\mathcal{T} \wedge [w]_{\mathsf{st}} = [r]_{\mathsf{st}}$; or 3) $r \in \Gamma.\mathcal{T} \wedge (w \notin \Gamma.\mathcal{NT} \Rightarrow [w]_{\mathsf{st}} \neq [r]_{\mathsf{st}})$.

In case (1), from the construction of Γ we then know that $(r, w) \in G.\mathsf{rb}$. As such, we have $w \stackrel{G.\mathsf{hb}}{\rightarrow} r \stackrel{G.\mathsf{rb}}{\rightarrow} w$, contradicting the assumption that G is consistent.

In case (2), let $\mathtt{loc}(w) = \mathtt{loc}(r) = \mathtt{x}$. From the definition of $\Gamma.\mathsf{rb}$, $\Gamma.\mathsf{rf}$ and $\Gamma.\mathsf{mo}$ we know there exists w' such that $[w']_{\mathsf{st}} = [r]_{\mathsf{st}}$, $(w', r) \in \Gamma.\mathsf{rf}$, $(w', r) \in \Gamma.\mathsf{po}$, $(w', w) \in \Gamma.\mathsf{mo}$, $(w', w) \in G.\mathsf{mo}$ and for all e if $w' \stackrel{\Gamma.\mathsf{po}}{\rightarrow} e \stackrel{\Gamma.\mathsf{po}}{\rightarrow} r$ then $\mathtt{loc}(e) \neq \mathtt{x} \vee e \notin \mathcal{W}$. Now since $[w']_{\mathsf{st}} = [w]_{\mathsf{st}}$ we know that either a) $(w, w') \in \Gamma.\mathsf{po}$; or b) $(w', w) \in \Gamma.\mathsf{po}$.

In case (2.a) from the definition of $\Gamma.\mathsf{po}$ we have $(w, w') \in G.\mathsf{po}$. As such we have $w \stackrel{G.\mathsf{po}}{\rightarrow} w' \stackrel{G.\mathsf{mo}}{\rightarrow} w$, contradicting the assumption that $G.\mathsf{hb}$ is consistent.

In case (2.b) since $[w]_{\mathsf{st}} = [r]_{\mathsf{st}}$ we know that either i) $(w, r) \in \Gamma.\mathsf{po}$; or 2) $(r, w) \in \Gamma.\mathsf{po}$. Moreover, since we know for all e if $w' \stackrel{\Gamma.\mathsf{po}}{\rightarrow} e \stackrel{\Gamma.\mathsf{po}}{\rightarrow} r$ then $\mathtt{loc}(e) \neq \mathtt{x} \vee e \notin \mathcal{W}$, and we have $w \in \mathcal{W}$ and $\mathtt{loc}(w) = \mathtt{x}$, we thus know $(r, w) \in \Gamma.\mathsf{po}$. From the definition of $\Gamma.\mathsf{po}$ we then have $(r, w) \in G.\mathsf{po} \subseteq G.\mathsf{hb}$ We thus have $w \stackrel{G.\mathsf{hb}}{\rightarrow} r \stackrel{G.\mathsf{hb}}{\rightarrow} w$, contradicting the assumption that G is consistent.

In case (3) we then know there exist ξ such that $r \in \xi$ and $w \notin \xi$. Let $\mathtt{loc}(w) = \mathtt{loc}(r) = \mathtt{x}$ and $\mathtt{val_r}(r) = v$. From the construction of Γ we know there exist $rx_1, rx_2 \in G.E$ such that $rx_1 = \mathtt{R}(\mathtt{x}, v)$, $rx_2 = \mathtt{R}(\mathtt{x}, v)$, $(rx_1, w), (rx_2, w) \in G.\mathsf{rb}$ and $rx_1 \stackrel{G.\mathsf{po}}{\rightarrow} rx_2 \stackrel{G.\mathsf{po}}{\rightarrow} r = \mathtt{R}(\mathtt{s[x]}, v))$. It is straightforward to demonstrate that

$G.\mathsf{hb} = G.\mathsf{po}^+ \cup (G.\mathsf{po} \cup G.\mathsf{rf})^*; (G.\mathsf{rf} \setminus G.\mathsf{po}); G.\mathsf{po}^*$. There are then two cases to consider: a) $(w, r) \in G.\mathsf{po}^+$; or b) $(w, r) \in (G.\mathsf{po} \cup G.\mathsf{rf})^*; (G.\mathsf{rf} \setminus G.\mathsf{po}); G.\mathsf{po}^*$. In case (3.a), since we have $(w, r) \in G.\mathsf{po}^+$, $r \in \xi$ and $w \notin \xi$., we also have $(w, rx_1) \in G.\mathsf{po}^+$. As such, we have $w \stackrel{G.\mathsf{po}^+}{\rightarrow} rx_1 \stackrel{G.\mathsf{rb}}{\rightarrow} w$, contradicting the assumption that G is consistent.

In case (3.b), we know there exist d, e such that $(w, d) \in (G.\mathsf{po} \cup G.\mathsf{rf})^*$, $(d, e) \in (G.\mathsf{rf} \setminus G.\mathsf{po})$ and $(e, r) \in G.\mathsf{po}^*$. Either i) $e \notin \xi$ or ii) $e \in \xi$. In case (3.b.i) since we have $(e, r) \in G.\mathsf{po}^*$ and $r, rx_1 \in \xi$, we know that $(e, rx_1) \in G.\mathsf{po}$. As such, we have $w \stackrel{(G.\mathsf{po} \cup G.\mathsf{rf})^*}{\rightarrow} d \stackrel{G.\mathsf{rf}}{\rightarrow} e \stackrel{G.\mathsf{po}}{\rightarrow} rx_1 \stackrel{G.\mathsf{rb}}{\rightarrow} w$. That is, since $G.\mathsf{rf} \subseteq G.\mathsf{hb}$, $G.\mathsf{po} \subseteq G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we have $w \stackrel{G.\mathsf{hb}}{\rightarrow} rx_1 \stackrel{G.\mathsf{rb}}{\rightarrow} w$, contradicting the assumption that G is consistent.

In (3.b.ii) there are two additional cases to consider: either 1) $e \stackrel{G.\mathsf{po}}{\rightarrow} rx_2$, or 2) $rx_2 \stackrel{G.\mathsf{po}}{\rightarrow} e$. In the (3.b.ii.1) case we then have $rx_2 \stackrel{G.\mathsf{rb}}{\rightarrow} w \stackrel{(G.\mathsf{po} \cup G.\mathsf{rf})^*}{\rightarrow} d \stackrel{G.\mathsf{rf}}{\rightarrow} e \stackrel{G.\mathsf{po}}{\rightarrow} rx_2$. That is, since $G.\mathsf{rf} \subseteq G.\mathsf{hb}$, $G.\mathsf{po} \subseteq G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we have $rx_2 \stackrel{G.\mathsf{rb}}{\rightarrow} w \stackrel{G.\mathsf{hb}}{\rightarrow} rx_2$, contradicting the assumption that G is consistent.

In the (3.b.ii.2) case, given the structure of our implementation we know there exists e' such that $rx_1 \stackrel{G.\mathsf{po}}{\rightarrow} e' \stackrel{G.\mathsf{po}}{\rightarrow} rx_2$, $\mathtt{val_r}(e) = \mathtt{val_r}(e')$, $\mathtt{loc}(e) = \mathtt{loc}(e')$. In what follows we demonstrate that we also have $(w, e') \in G.\mathsf{rf}$. We thus have $rx_2 \stackrel{G.\mathsf{rb}}{\rightarrow} w \stackrel{(G.\mathsf{po} \cup G.\mathsf{rf})^*}{\rightarrow} d \stackrel{G.\mathsf{rf}}{\rightarrow} e' \stackrel{G.\mathsf{po}}{\rightarrow} rx_2$. That is, since $G.\mathsf{rf} \subseteq G.\mathsf{hb}$, $G.\mathsf{po} \subseteq G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we have $rx_2 \stackrel{G.\mathsf{rb}}{\rightarrow} w \stackrel{G.\mathsf{hb}}{\rightarrow} rx_2$, contradicting the assumption that G is consistent.

As our only remaining proof obligation let us show that above we also have $(w, e') \in G.\mathsf{rf}$. Either $\mathtt{loc}(e) = \mathtt{loc}(e') = \mathtt{vy}$ for some location lock, in which case from Lemma 7.3 in §2.3 we know that the writes to sequence locks write unique values and thus as we have $\mathtt{val_r}(e) = \mathtt{val_r}(e')$, we also have $(w, e') \in G.\mathsf{rf}$, as required. Or $\mathtt{loc}(e) = \mathtt{loc}(e') = \mathtt{vy}$ for some shared location $\mathtt{y}$. Now w is either a non-transactional write, in which case since we assume values written by non-transactional writes are unique and we have $\mathtt{val_r}(e) = \mathtt{val_r}(e')$ and $(w, e) \in G.\mathsf{rf}$, we also have $(w, e') \in G.\mathsf{rf}$, as required.

Now let us assume that w is a transactional event where $w \in \xi_w$ for some ξ_w, and let $(w', e') \in G.\mathsf{rf}$. We must then show that $w = w'$. As the values written by non-transactional writes are unique, we also know that $w' \in \xi_{w'}$ for some transaction $\xi_{w'}$. We also know that $(w', w) \in G.\mathsf{mo}$, since otherwise (when $(w, w') \in G.\mathsf{mo}$) we have $e \stackrel{G.\mathsf{rb}}{\rightarrow} w' \stackrel{G.\mathsf{rf}}{\rightarrow} e' \stackrel{G.\mathsf{po}}{\rightarrow} e$, contradicting the assumption that G is consistent. Furthermore, we know there exist yv_1, yv_2 such that $yv_1 \stackrel{G.\mathsf{po}}{\rightarrow} e' \stackrel{G.\mathsf{po}}{\rightarrow} e \stackrel{G.\mathsf{po}}{\rightarrow} yv_2$, $\mathtt{loc}(yv_1) = \mathtt{loc}(yv_2) = \mathtt{y}$, $\mathtt{val_r}(yv_1) = \mathtt{val_r}(yv_2)$. As a final proof obligation below we show that $(\xi_{w'}.U^{\mathtt{y}}, yv_1), (\xi_{w'}.U^{\mathtt{y}}, yv_2) \in G.\mathsf{rf}$. From Lemma 7.2 in §2.3 we then know that either $\xi_w.U^{\mathtt{y}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_{w'}.L^{\mathtt{y}}$ or $\xi_{w'}.U^{\mathtt{y}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_w.L^{\mathtt{y}}$. In the former case we then have $\xi_w.U^{\mathtt{y}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_{w'}.L^{\mathtt{y}} \stackrel{G.\mathsf{po}}{\rightarrow} w' \stackrel{G.\mathsf{mo}}{\rightarrow} w \stackrel{G.\mathsf{po}}{\rightarrow} \xi_w.U^{\mathtt{y}}$, i.e. $w' \stackrel{G.\mathsf{mo}}{\rightarrow} w \stackrel{G.\mathsf{hb}}{\rightarrow} w'$, contradicting the assumption that G is consistent. In the latter case, we also have $\xi_{w'}.U^{\mathtt{y}} \stackrel{G.\mathsf{mo}}{\rightarrow} \xi_w.L^{\mathtt{y}}$ (since otherwise we have $\xi_{w'}.U^{\mathtt{y}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_w.L^{\mathtt{y}} \stackrel{G.\mathsf{mo}}{\rightarrow} \xi_{w'}.U^{\mathtt{y}}$, contradicting the assumption that G is consistent). As such we have have $(yv_1, \xi_w.L^{\mathtt{y}}), (yv_2, \xi_w.L^{\mathtt{y}}) \in G.\mathsf{rb}$. We then have $w \stackrel{G.\mathsf{rf}}{\rightarrow} e \stackrel{G.\mathsf{po}}{\rightarrow} yv_2 \stackrel{G.\mathsf{rb}}{\rightarrow} \xi_w.L^{\mathtt{y}} \stackrel{G.\mathsf{po}}{\rightarrow} w$. That is, we have $yv_2 \stackrel{G.\mathsf{rb}}{\rightarrow}\stackrel{G.\mathsf{rb}}{\rightarrow} \xi_w.L^{\mathtt{y}} \stackrel{G.\mathsf{hb}}{\rightarrow} yv_2$, contradicting the assumption that G is consistent.

As our last proof obligation let us show that $(\xi_{w'}.U^{\mathtt{y}}, yv_1), (\xi_{w'}.U^{\mathtt{y}}, yv_2) \in G.\mathsf{rf}$. From Lemma 7.3 in §2.3 we know that the values written to `vy` are unique. As such, we know that there exists w'' such that $(w'', yv_1), (w'', yv_2) \in G.\mathsf{rf}$. Now either 1) $w'' \stackrel{G.\mathsf{mo}}{\rightarrow} \xi_{w'}.L^{\mathtt{y}}$; or 2) $\xi_{w'}.L^{\mathtt{y}} \stackrel{G.\mathsf{mo}}{\rightarrow} w''$.

In case (1) we then have $(yv_2, \xi_{w'}.L^{\mathtt{y}}) \in G.\mathsf{rb}$. As such we have $w' \stackrel{G.\mathsf{rf}}{\rightarrow} e \stackrel{G.\mathsf{po}}{\rightarrow} yv_2 \stackrel{G.\mathsf{rb}}{\rightarrow} \xi_{w'}.L^{\mathtt{y}} \stackrel{G.\mathsf{po}}{\rightarrow} w'$. That is, we have $\xi_{w'}.L^{\mathtt{y}} \stackrel{G.\mathsf{hb}}{\rightarrow} yv_2 \stackrel{G.\mathsf{rb}}{\rightarrow} \xi_{w'}.L^{\mathtt{y}}$, contradicting the assumption that G is consistent.

In case (2) we then know there exists $\xi_{w''}$ such that $w'' \in \xi_{w''}$ and $w'' = \xi_{w''}.U^{\mathtt{y}}$. We also know there exists a write event $w_y \in \xi_{w''}$ such that $\mathtt{loc}(w_y) = \mathtt{y}$ and $\xi_{w''}.L^{\mathtt{y}} \stackrel{G.\mathsf{po}}{\rightarrow} w_y \stackrel{G.\mathsf{po}}{\rightarrow} \xi_{w''}.U^{\mathtt{y}}$. From Lemma 7.2 in §2.3 we then have either $\xi_{w''}.U^{\mathtt{y}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_{w'}.L^{\mathtt{y}}$ or $\xi_{w'}.U^{\mathtt{y}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_{w''}.L^{\mathtt{y}}$. In the former case we then have $\xi_{w''}.U^{\mathtt{y}} \stackrel{G.\mathsf{mo}}{\rightarrow} \xi_{w'}.L^{\mathtt{y}}$ (since otherwise we have a cycle $\xi_{w''}.U^{\mathtt{y}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_{w'}.L^{\mathtt{y}} \stackrel{G.\mathsf{mo}}{\rightarrow} \xi_{w''}.L^{\mathtt{y}}$ contradicting the assumption that G is consistent.) As such we have, $(yv_2, \xi_{w'}.L^{\mathtt{y}}) \in G.\mathsf{rb}$. We then have $w' \stackrel{G.\mathsf{rf}}{\rightarrow} e \stackrel{G.\mathsf{po}}{\rightarrow} yv_2 \stackrel{G.\mathsf{rb}}{\rightarrow} \xi_{w'}.L^{\mathtt{y}} \stackrel{G.\mathsf{po}}{\rightarrow} w'$. That is, we have $\xi_{w'}.L^{\mathtt{y}} \stackrel{G.\mathsf{hb}}{\rightarrow} yv_2 \stackrel{G.\mathsf{rb}}{\rightarrow} \xi_{w'}.L^{\mathtt{y}}$, contradicting the assumption that G is consistent.

In the latter case we know that $(w', w_y) \in \mathsf{mo}$ (since otherwise we have a cycle $\xi_{w'}.U^{\mathtt{y}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_{w''}.L^{\mathtt{y}} \stackrel{G.\mathsf{po}}{\rightarrow} w_y \stackrel{G.\mathsf{mo}}{\rightarrow} w' \stackrel{G.\mathsf{po}}{\rightarrow} \xi_{w'}.U^{\mathtt{y}}$, contradicting the assumption that G is consistent.) As such, we have $(e, w_y) \in \mathsf{rb}$. We then have $w_y \stackrel{G.\mathsf{po}}{\rightarrow} \xi_{w''}.U^{\mathtt{y}} \stackrel{G.\mathsf{rf}}{\rightarrow} yv_1 \stackrel{G.\mathsf{po}}{\rightarrow} e \stackrel{G.\mathsf{rb}}{\rightarrow} w_y$, contradicting the assumption that G is consistent. ☐

2.2 Implementation Completeness

The RPSI implementation in Fig. 1 is *complete*: for each consistent specification graph Γ a corresponding implementation graph G can be constructed with the same program outcome such that $\mathsf{RA\text{-}consistent}(G)$ holds.

Constructing Consistent Implementation Graphs In order to construct an execution graph of the implementation G from the specification Γ, we follow similar steps as those in the corresponding PSI construction in §1.2. More concretely, the events associated with non-transactional events are unchanged and simply added to the implementation graph. For transactional events, given each trace θ'_i of a transaction in the specification, as before we construct an analogous trace of the implementation by inserting the appropriate events for acquiring and inspecting the version locks, as well as obtaining a snapshot. For each transaction class $\mathcal{T}_i \in \mathcal{T}/\mathsf{st}$, we first determine its read and write sets as before and subsequently decide the order in which the version locks are acquired and inspected. This then enables us to construct the 'reads-from' and 'modification-order' relations for the events associated with version locks.

Given a consistent execution graph of the specification $\Gamma = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathcal{T})$, and a transaction class $\mathcal{T}_i \in \Gamma.\mathcal{T}/\mathsf{st}$, we define $\mathtt{WS}_{\mathcal{T}_i}$ and $\mathtt{RS}_{\mathcal{T}_i}$ as described in §1.2. Determining the ordering of lock events hinges on a similar observation as that in the PSI construction. Given a consistent execution graph of the specification $\Gamma = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathcal{T})$, let for each location $\mathtt{x}$ the total order mo be given as: $w_1 \stackrel{\mathsf{mo}|_{\text{imm}}}{\to} \cdots \stackrel{\mathsf{mo}|_{\text{imm}}}{\to} w_{n_{\mathtt{x}}}$. This order can be broken into adjacent segments where the events of each segment are *either* non-transactional writes *or* belong to the *same* transaction. That is, given the transaction classes $\Gamma.\mathcal{T}/\mathsf{st}$, the order above is of the following form where $\mathcal{T}_1, \cdots, \mathcal{T}_m \in \Gamma.\mathcal{T}/\mathsf{st}$ and for each such $\mathcal{T}_i$ we have $\mathtt{x} \in \mathtt{WS}_{\mathcal{T}_i}$ and $w_{(i,1)} \cdots w_{(i,n_i)} \in \mathcal{T}_i$:

$$\underbrace{w_{(1,1)} \stackrel{\mathsf{mo}|_{\text{imm}}}{\to} \cdots \stackrel{\mathsf{mo}|_{\text{imm}}}{\to} w_{(1,n_1)}}_{\Gamma.\mathcal{NT} \cup \mathcal{T}_1} \stackrel{\mathsf{mo}|_{\text{imm}}}{\to} \cdots \stackrel{\mathsf{mo}|_{\text{imm}}}{\to} \underbrace{w_{(m,1)} \stackrel{\mathsf{mo}|_{\text{imm}}}{\to} \cdots \stackrel{\mathsf{mo}|_{\text{imm}}}{\to} w_{(m,n_m)}}_{\Gamma.\mathcal{NT} \cup \mathcal{T}_m}$$

Were this not the case and we had $w_1 \stackrel{\mathsf{mo}}{\to} w \stackrel{\mathsf{mo}}{\to} w_2$ such that $w_1, w_2 \in \mathcal{T}_i$ and $w \in \mathcal{T}_j \neq \mathcal{T}_i$, we would consequently have $w_1 \stackrel{\mathsf{mo}_\mathsf{T}}{\to} w \stackrel{\mathsf{mo}_\mathsf{T}}{\to} w_1$, contradicting the assumption that Γ is consistent. We thus define $\Gamma.\mathsf{MO}_{\mathtt{x}} = [\mathcal{T}_1 \cdots \mathcal{T}_m]$.

Note that each transactional execution trace of the specification is of the form $\theta'_i = B_i \stackrel{\mathsf{po}}{\to} Ts'_i \stackrel{\mathsf{po}}{\to} E_i$, with B_i, E_i and Ts'_i as described in §1.2. For each such θ'_i, we construct a corresponding trace of our implementation as $\theta_i = Ls_i \stackrel{\mathsf{po}}{\to} S_i \stackrel{\mathsf{po}}{\to} Ts_i \stackrel{\mathsf{po}}{\to} Us_i$, where Ls_i, Ts_i and Us_i are as defined in §1.2, and $S_i = tr_i^{\mathtt{x}_1} \stackrel{\mathsf{po}}{\to} \cdots \stackrel{\mathsf{po}}{\to} tr_i^{\mathtt{x}_p} \stackrel{\mathsf{po}}{\to} vr_i^{\mathtt{x}_1} \stackrel{\mathsf{po}}{\to} \cdots \stackrel{\mathsf{po}}{\to} vr_i^{\mathtt{x}_p}$ denotes the sequence of events obtaining a tentative snapshot ($tr_i^{\mathtt{x}_j}$) and subsequently validating it ($vr_i^{\mathtt{x}_j}$). Each $tr_i^{\mathtt{x}_j}$ sequence is of the form $ivr_i^{\mathtt{x}_j} \stackrel{\mathsf{po}}{\to} ir_i^{\mathtt{x}_j} \stackrel{\mathsf{po}}{\to} s_i^{\mathtt{x}_j}$, with $ivr_i^{\mathtt{x}_j}$, $ir_i^{\mathtt{x}_j}$ and $s_i^{\mathtt{x}_j}$ defined below (with fresh identifiers). Similarly, each $vr_i^{\mathtt{x}_j}$ sequence is of the form $fr_i^{\mathtt{x}_j} \stackrel{\mathsf{po}}{\to} fvr_i^{\mathtt{x}_j}$, with $fr_i^{\mathtt{x}_j}$ and $fvr_i^{\mathtt{x}_j}$ defined as follows (with fresh identifiers). We then define the rf relation for each of these read events in S_i in a similar way.

For each $(\mathtt{x}, r) \in \mathtt{RS}_{\mathcal{T}_i}$, when r (the event in the specification class $\mathcal{T}_i$ that reads the value of $\mathtt{x}$) reads from w in the specification graph ($(w, r) \in \Gamma.\mathsf{rf}$), we add $(w, ir_i^{\mathtt{x}})$ and $(w, fr_i^{\mathtt{x}})$ to the rf of G (the first line of IRF_i^2 below). For version locks, as before if transaction $\mathcal{T}_i$ also writes to $\mathtt{x}_j$, then $ivr_i^{\mathtt{x}_j}$ and $fvr_i^{\mathtt{x}_j}$ events (reading and validating $\mathtt{vx}_j$), read from the lock event in $\mathcal{T}_i$ that acquired $\mathtt{vx}_j$, namely $L_i^{\mathtt{x}_j}$. Similarly, if $\mathcal{T}_i$ does not write to $\mathtt{x}_j$ and it reads the value of $\mathtt{x}_j$ written by the initial write, $init_{\mathtt{x}}$, then $ivr_i^{\mathtt{x}_j}$ and $fvr_i^{\mathtt{x}_j}$ read the value written to $\mathtt{vx}_j$ by the initial write to $\mathtt{vx}$, $init_{\mathtt{vx}}$. Lastly, if transaction $\mathcal{T}_i$ does not write to $\mathtt{x}_j$ and it reads $\mathtt{x}_j$ from a write other than $init_{\mathtt{x}}$, then $ir_i^{\mathtt{x}_j}$ and $vr_i^{\mathtt{x}_j}$ read from the unlock event of a transaction $\mathcal{T}_j$ (i.e. $U_j^{\mathtt{x}}$), who has $\mathtt{x}$ in its write set and whose write to $\mathtt{x}$, $w_{\mathtt{x}}$, maximally 'RPSI-happens-before' r. That is, for all other such writes that 'RPSI-happen-before' r, then $w_{\mathtt{x}}$ 'RPSI-happens-after' them.

$$\mathsf{IRF}^2_i \triangleq \bigcup_{(\mathtt{x},r)\in \mathtt{RS}_{\mathcal{T}_i}} \left\{ \begin{array}{l} (w, ir^{\mathtt{x}}_i), \\ (w, fr^{\mathtt{x}}_i), \\ (w', ivr^{\mathtt{x}}_i), \\ (w', fvr^{\mathtt{x}}_i) \end{array} \middle| \begin{array}{l} (w,r) \in \Gamma.\mathsf{rf} \wedge (\mathtt{x} \in \mathtt{WS}_{\mathcal{T}_i} \Rightarrow w' {=} L^{\mathtt{x}}_i) \\ \wedge (\mathtt{x} \notin \mathtt{WS}_{\mathcal{T}_i} \wedge w {=} init_{\mathtt{x}} \Rightarrow w' {=} init_{\mathtt{vx}}) \\ \wedge (\mathtt{x} \notin \mathtt{WS}_{\mathcal{T}_i} \wedge w {\neq} init_{\mathtt{x}} \Rightarrow \\ \quad \exists w_{\mathtt{x}}, \mathcal{T}_j .\ w_{\mathtt{x}} \in \mathcal{T}_j \cap \mathcal{W}_{\mathtt{x}} \wedge w_{\mathtt{x}} \stackrel{\mathsf{rpsi\text{-}hb}}{\rightarrow} r \wedge w' {=} U^{\mathtt{x}}_j \\ \quad \wedge [\forall w'_{\mathtt{x}}, \mathcal{T}_k .\ w'_{\mathtt{x}} {\in} \mathcal{T}_k \cap \mathcal{W}_{\mathtt{x}} \wedge w'_{\mathtt{x}} \stackrel{\mathsf{rpsi\text{-}hb}}{\rightarrow} r \Rightarrow w'_{\mathtt{x}} \stackrel{\mathsf{rpsi\text{-}hb}}{\rightarrow} w_{\mathtt{x}}]) \end{array} \right\}$$

$$ir^{\mathtt{x}_j}_i {=} fr^{\mathtt{x}_j}_i {=} \mathtt{R}(\mathtt{x}_j, v) \quad s^{\mathtt{x}_j}_i {=} \mathtt{W}(\mathtt{s[x}_j\mathtt{]}, v) \quad \text{s.t. } \exists w.\ (w, ir^{\mathtt{x}_j}_i) \in \mathsf{IRF}^2_i \wedge \mathtt{val_w}(w) {=} v$$

$$ivr^{\mathtt{x}_j}_i {=} fvr^{\mathtt{x}_j}_i {=} \mathtt{R}(\mathtt{vx}_j, v) \text{ s.t. } \exists w.\ (w, ivr^{\mathtt{x}_j}_i) \in \mathsf{IRF}^2_i \wedge \mathtt{val_w}(w) {=} v$$

We are now in a position to construct our implementation graph. Given a consistent execution graph Γ of the specification, we construct an execution graph of the implementation, $G = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo})$, such that:

- $G.E = \bigcup_{\mathcal{T}_i \in \Gamma.\mathcal{T}/\mathsf{st}} \theta_i.E \cup \Gamma.\mathcal{NT}$;
- $G.\mathsf{po}$ is defined as $\Gamma.\mathsf{po}$ extended by the po for the additional events of G, given by the θ_i traces defined above;
- $G.\mathsf{rf} = \bigcup_{\mathcal{T}_i \in \Gamma.\mathcal{T}/\mathsf{st}} (\mathsf{IRF}^1_i \cup \mathsf{IRF}^2_i)$, with IRF^1_i as in §1.2 and IRF^2_i defined above;
- $G.\mathsf{mo} = \Gamma.\mathsf{mo} \cup \left(\bigcup_{\mathcal{T}_i \in \Gamma.\mathcal{T}/\mathsf{st}} \mathsf{IMO}_i \right)^+$, with IMO_i as defined in §1.2.

Theorem 4 (Completeness). *For all RPSI-consistent specification graphs Γ of a program, there exists an RA-consistent execution graph G of the implementation in Fig. 1 that has the same program outcome.*

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

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

RTS. part 1

We proceed by contradiction. Let us assume that there exists a such that $(a, a) \in G.\mathsf{hb}_{loc}$. There are then two cases to consider: 1) $a \in G.\mathcal{NT}$; or 2) $a \in \mathcal{T}'_a$ for some $\mathcal{T}'_a \in \bigcup_{\mathcal{T}_i \in \Gamma.\mathcal{T}/\mathsf{st}} \mathcal{T}'_i$. In case (1), from auxiliary Lemma 9 in §2.4 we have $(a, a) \in \Gamma.\mathsf{rpsi\text{-}hb}$, contradicting the assumption that Γ is consistent. In case (2) from auxiliary Lemma 9 in §2.4 we have $(a, a) \in G.\mathsf{po}$, which is impossible given the construction of $G.\mathsf{po}$.

RTS. part 2

We proceed by contradiction. Let us assume that there exist a, b such that $(a, b) \in G.\mathsf{mo}$ and $(b, a) \in G.\mathsf{hb}_{loc}$. We now need to consider five cases: 1) $a, b \in G.\mathcal{NT}$; or 2) $a \in G.\mathcal{NT} \wedge b \in \mathcal{T}'_b$; or 3) $a \in \mathcal{T}'_a \wedge b \in G.\mathcal{NT}$; or 4)

$a{\in}\mathcal{T}'_a \wedge b{\in}\mathcal{T}'_a \wedge \mathcal{T}'_a{\neq}\mathcal{T}'_b$; or 5) $a, b \in \mathcal{T}'_a$.

Case 1. $a, b \in G.\mathcal{NT}$
From the construction of $G.\mathsf{mo}$ we then have $(a, b) \in \Gamma.\mathsf{mo}$. On the other hand, from auxiliary Lemma 9 in §2.4 we have $(b, a) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As such, we have $a \stackrel{\Gamma.\mathsf{mo}}{\rightarrow} b \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\rightarrow} a$, contradicting the assumption that Γ is consistent.

Case 2. $a \in G.\mathcal{NT} \wedge b \in \mathcal{T}'_b$
From the construction of $G.\mathsf{mo}$ we then have $(a, b) \in \Gamma.\mathsf{mo}$ and that $b \in \mathcal{T}_b$ (since b is a write event). On the other hand, from auxiliary Lemma 9 in §2.4 we know there exists $(b, a) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As such, we have $a \stackrel{\Gamma.\mathsf{mo}}{\rightarrow} b \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\rightarrow} a$, contradicting the assumption that Γ is consistent.

Case 3. $a \in \mathcal{T}'_a \wedge b \in G.\mathcal{NT}$
From the construction of $G.\mathsf{mo}$ we then have $(a, b) \in \Gamma.\mathsf{mo}$. On the other hand, from auxiliary Lemma 9 in §2.4 we have $(\{b\} \times \mathcal{T}_a) \in \Gamma.\mathsf{rpsi\text{-}hb}$. In particular, we have $(b, a) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As such, we have $a \stackrel{\Gamma.\mathsf{mo}}{\rightarrow} b \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\rightarrow} a$, contradicting the assumption that Γ is consistent.

Case 4. $a{\in}\mathcal{T}'_a \wedge b{\in}\mathcal{T}'_a \wedge \mathcal{T}'_a{\neq}\mathcal{T}'_b$
There are then two cases to consider: 1) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{x}$, for some shared location $\mathtt{x}$; or 2) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{vx}$, for some location lock $\mathtt{vx}$ associated with location $\mathtt{x}$.

In case (1), since a, b are both write events, from the construction of $G.\mathsf{mo}$ we have $(a, b) \in \Gamma.\mathsf{mo}$. On the other hand, from auxiliary Lemma 9 in §2.4 we know $(b, a) \in \Gamma.\mathsf{rpsi\text{-}hb}$. We then have $a \stackrel{\Gamma.\mathsf{mo}}{\rightarrow} b \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\rightarrow} a$, contradicting the assumption that Γ is consistent.

In case (2), from the construction of $G.\mathsf{mo}$ we know there exists $d \in \mathcal{T}_a$ and $e \in \mathcal{T}_b$ such that $\mathtt{loc}(d) = \mathtt{loc}(e) = \mathtt{x}$ and that $(d, e) \in \Gamma.\mathsf{mo}$. On the other hand, from auxiliary Lemma 9 in §2.4 we know there exists $c \in \mathcal{T}_b$ such that $(\{c\} \times \mathcal{T}_a) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. In particular, we have $(c, d) \in \Gamma.\mathsf{rpsi\text{-}hb}$. Moreover, since $(d, e) \in \Gamma.\mathsf{mo}$ and $e, c \in \mathcal{T}_b \neq \mathcal{T}_a \ni a$, we have $(d, c) \in \Gamma.\mathsf{mo}_{\mathsf{T}}$, and thus $(d, c) \in \Gamma.\mathsf{rpsi\text{-}hb}$. We then have $d \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\rightarrow} c \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\rightarrow} d$, contradicting the assumption that Γ is consistent.

Case 5. $a, b \in \mathcal{T}'_a$
From auxiliary Lemma 9 in §2.4 we have $(b, a) \in G.\mathsf{po}$. There are now two cases to consider: a) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{x}$, for some shared location $\mathtt{x}$; or b) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{vx}$, for some location lock $\mathtt{vx}$ associated with location $\mathtt{x}$.

In case (a), since a, b are both write events, from the construction of G we know that $a, b \in \Gamma.E$. As $G.\mathsf{po}$ does not alter the orderings between events of $\Gamma.E$ we also have $(b, a) \in \Gamma.\mathsf{po}$. On the other hand, from the construction of $G.\mathsf{mo}$ we have $(a, b) \in \Gamma.\mathsf{mo}$. We then have $a \stackrel{\Gamma.\mathsf{mo}}{\rightarrow} b \stackrel{\Gamma.\mathsf{po}}{\rightarrow} a$; that is, $a \stackrel{\Gamma.\mathsf{mo}}{\rightarrow} b \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\rightarrow} a$, contradicting the assumption that Γ is consistent.

In case (b), from the construction of G.mo for location locks we have $(a, b) \in G$.hb. We thus have $a \stackrel{G.\mathsf{po}}{\rightarrow} b \stackrel{G.\mathsf{po}}{\rightarrow} a$, which is impossible given our construction of G.po.

RTS. part 3

Let us assume that there exist a, b such that $(a, b) \in G$.rb and $(b, a) \in G$.hb. We now need to consider five cases: 1) $a, b \in G.\mathcal{NT}$; or 2) $a \in G.\mathcal{NT} \land b \in \mathcal{T}'_b$; or 3) $a \in \mathcal{T}'_a \land b \in G.\mathcal{NT}$; or 4) $a{\in}\mathcal{T}'_a \land b{\in}\mathcal{T}'_a \land \mathcal{T}'_a{\neq}\mathcal{T}'_b$; or 5) $a, b \in \mathcal{T}'_a$.

Case 1. $a, b \in G.\mathcal{NT}$

From the construction of G.rb we then have $(a, b) \in \Gamma$.rb. On the other hand, from auxiliary Lemma 9 in §2.4 we have $(b, a) \in \Gamma$.rpsi-hb. As such, we have $a \stackrel{\Gamma.\mathsf{rb}}{\rightarrow} b \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\rightarrow} a$, contradicting the assumption that Γ is consistent.

Case 2. $a \in G.\mathcal{NT} \land b \in \mathcal{T}'_b$

From the construction of G.rb we then have $(a, b) \in \Gamma$.rb. On the other hand, from auxiliary Lemma 9 in §2.4 we know $(b, a) \in \Gamma$.rpsi-hb. As such, we have $a \stackrel{\Gamma.\mathsf{rb}}{\rightarrow} b \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\rightarrow} a$, contradicting the assumption that Γ is consistent.

Case 3. $a \in \mathcal{T}'_a \land b \in G.\mathcal{NT}$

From the construction of G.rb we then know there exists $c \in \mathcal{T}_a$ such that $(c, b) \in \Gamma$.rb. On the other hand, from auxiliary Lemma 9 in §2.4 we have $(\{b\} \times \mathcal{T}_a) \in \Gamma$.rpsi-hb. In particular, we have $(b, c) \in \Gamma$.rpsi-hb. As such, we have $c \stackrel{\Gamma.\mathsf{rb}}{\rightarrow} b \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\rightarrow} c$, contradicting the assumption that Γ is consistent.

Case 4. $a{\in}\mathcal{T}'_a \land b{\in}\mathcal{T}'_a \land \mathcal{T}'_a{\neq}\mathcal{T}'_b$

There are then two cases to consider: a) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{x}$, for some shared location $\mathtt{x}$; or b) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{vx}$, for some location lock $\mathtt{vx}$ associated with location $\mathtt{x}$.

In case (a), from the construction of G.rb we then know there exists $c \in \mathcal{T}_a$ such that $(c, b) \in \Gamma$.rb. On the other hand, from auxiliary Lemma 9 in §2.4 we know $(\{b\} \times \mathcal{T}_a) \in \Gamma$.rpsi-hb. In particular, we have $(b, c) \in \Gamma$.rpsi-hb. As such, we have $c \stackrel{\Gamma.\mathsf{rb}}{\rightarrow} b \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\rightarrow} c$, contradicting the assumption that Γ is consistent.

In case (b), there are again two cases to consider: 1) $\mathtt{x} \in \mathtt{WS}_{\xi_a}$; or 2) $\mathtt{x} \notin \mathtt{WS}_{\xi_a}$.

In (b.1) from the construction of G.rf, G.mo and G.rb we know that there exist $wx_a \in \mathcal{T}_a$ and $wx_b \in \mathcal{T}_b$ such that $(wx_a, wx_b) \in \Gamma$.mo. On the other hand, from auxiliary Lemma 9 in §2.4 we know that there exist $c \in \mathcal{T}_b$ such that $(\{c\} \times \mathcal{T}_a) \subseteq \Gamma$.rpsi-hb. In particular, we have $(c, wx_a) \in \Gamma$.rpsi-hb. Moreover, since $(wx_a, wx_b) \in \Gamma$.mo, we have $(wx_a, c) \in \Gamma.\mathsf{mo}_\mathsf{T} \subseteq \Gamma$.rpsi-hb. As such, we have $c \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\rightarrow} wx_a \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\rightarrow} c$. That is, $c \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\rightarrow} c$, contradicting the assumption that Γ is consistent.

In (b.2), from the construction of G.rf for location locks we know there exist $\mathcal{T}'_c$, $wx_c \in \mathcal{T}_c$, $wxv_c \in \mathcal{T}'_c$, and $rx \in \mathcal{T}_a$ such that $\mathcal{T}'_c \neq \mathcal{T}'_a$, $\mathcal{T}'_c \neq \mathcal{T}'_b$, $rx{=}\mathtt{R}(\mathtt{x}, -)$, $wx_c = \mathtt{W}(\mathtt{x}, -)$, $(wx_c, rx) \in \Gamma$.rpsi-hb, $(wxv_c, a) \in G$.rf, $(wxv_c, b) \in$

G.mo, $\mathcal{T}_c \in HB^{\mathtt{x}}_a$ and that $\mathcal{T}_c$ is the maximal element of $HB^{\mathtt{x}}_a$: $\forall \mathcal{T}_d, j, k.\ \mathcal{T}_d \in HB^{\mathtt{x}}_a \wedge \mathsf{MO}_{\mathtt{x}}|_j = \mathcal{T}_c \wedge \mathsf{MO}_{\mathtt{x}}|_k = \mathcal{T}_d \Rightarrow j \geq k$). Moreover, since $(wxv_c, b) \in G.\mathsf{mo}$, from the construction of G.mo we know that there exists $wx_b \in \mathcal{T}_b$ such that $(wx_c, wx_b) \in \Gamma.\mathsf{mo}$. On the other hand, from auxiliary Lemma 9 in §2.4 and since $(b, a) \in G.\mathsf{hb}$, we know $\exists e \in \mathcal{T}_b.\ (\{e\} \times \mathcal{T}_a) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. Moreover, since $wx_b \in \mathcal{T}_b$ and $rx \in \mathcal{T}_a$, we have $(\mathcal{T}_b \times \mathcal{T}_a) \subseteq \Gamma.(\mathsf{st}; ([\mathcal{W}]; \mathsf{st}; \mathsf{hb}; \mathsf{st}; [\mathcal{R}])_{loc}; \mathsf{st}) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. In particular, we have $(wx_b, rx) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As such, we have $\mathcal{T}_b \in HB^{\mathtt{x}}_a$. Pick m, n such that $\mathsf{MO}_{\mathtt{x}}|_m = \mathcal{T}_c \wedge \mathsf{MO}_{\mathtt{x}}|_n = \mathcal{T}_b$. We then know that $m < n$ (since $(wx_c, wx_b) \in \Gamma.\mathsf{mo}$). This however contradicts the assumption that $\mathcal{T}_b$ is the maximal element of $HB^{\mathtt{x}}_a$.

Case 5. $a, b \in \mathcal{T}'_a$
From auxiliary Lemma 9 in §2.4 we have $(b, a) \in G.\mathsf{po}$. On the other hand, from the construction of G and since $(a, b) \in G.\mathsf{rb}$, we have $(a, b) \in G.\mathsf{po}$. We thus have $a \stackrel{G.\mathsf{po}}{\rightarrow} b \stackrel{G.\mathsf{po}}{\rightarrow} a$, which is impossible given our construction of G.po. □

2.3 Auxiliary Soundness Lemmata

Lemma 7. *For all consistent execution graphs of the implementation* $G = (E, \mathsf{po}, \mathsf{rf}, \mathsf{mo})$ *and its transaction set* Tx*, for all version lock locations* $\mathtt{vx}$*, and all transaction subsets* $\mathrm{Tx}_{\mathtt{vx}} \subseteq \mathrm{Tx}$ *with* $\mathtt{vx}$ *in their write sets* ($\forall \xi \in \mathrm{Tx}_{\mathtt{vx}}.\ \mathtt{x} \in \mathtt{WS}_\xi$)*:*

1. *there exists* $L = [\xi_1 \cdots \xi_m] = \mathtt{perm}(\mathrm{Tx}_{\mathtt{vx}})$*, such that:*

$$\xi_1.SL_{\mathtt{vx}} \stackrel{\mathsf{mo}|_{imm}}{\rightarrow} \xi_1.U_{\mathtt{vx}} \stackrel{\mathsf{mo}|_{imm}}{\rightarrow} \cdots \stackrel{\mathsf{mo}|_{imm}}{\rightarrow} \xi_m.SL_{\mathtt{vx}} \stackrel{\mathsf{mo}|_{imm}}{\rightarrow} \xi_m.U_{\mathtt{vx}}$$

   where $\xi_i.SL_{\mathtt{vx}}$ *denotes the event corresponding to the successful acquisition of the* $\mathtt{vx}$ *lock in transaction* ξ_i*, and* $\xi_i.U_{\mathtt{vx}}$ *denotes the unlocking of* $\mathtt{vx}$ *in* ξ_i *(i.e.* $\xi_i.SL_{\mathtt{vx}} = \mathtt{U}(\mathtt{vx}, a, a{+}1)$ *and* $\xi_i.U_{\mathtt{vx}} = \mathtt{W}(\mathtt{vx}, a{+}2)$*, for some* a *such that* $a \bmod 2 = 0$*).*
2. *for all* $\xi_1, \xi_2 \in \mathrm{Tx}_{\mathtt{vx}}$*, if* $\xi_1 \neq \xi_2$*, then either* $\xi_1.U_{\mathtt{vx}} \stackrel{\mathsf{hb}}{\rightarrow} \xi_2.SL_{\mathtt{vx}}$*, or* $\xi_2.U_{\mathtt{vx}} \stackrel{\mathsf{hb}}{\rightarrow} \xi_1.SL_{\mathtt{vx}}$*.*
3. *each write event to location* $\mathtt{vx}$ *in* E*, writes a unique value:*

$$\forall a, b \in G.\mathcal{W}_{\mathtt{vx}}.\ \mathtt{val_w}(a) \neq \mathtt{val_w}(b)$$

Proof (part 1). By induction on the length of $\mathrm{Tx}_{\mathtt{vx}}$.

Base case $\mathrm{Tx}_{\mathtt{vx}} = \{\}$.
This case holds vacuously.

Inductive case $|\mathrm{Tx}_{\mathtt{vx}}| = m$, where $m \geq 1$.
Given the trace of each transaction described above, we know that the set of write events on $\mathtt{vx}$ is given by $\mathcal{W}_{\mathtt{vx}} = \bigcup_{\xi_i \in \mathrm{Tx}_{\mathtt{vx}}} \{\xi_i.SL_{\mathtt{vx}}, \xi_i.U_{\mathtt{vx}}\}$. Since the write events of $\mathtt{vx}$ are totally ordered by mo, we know there exists a minimal $e_0 \in \mathcal{W}_{\mathtt{vx}}$ such that $\forall e \in \mathcal{W}_{\mathtt{vx}} \setminus \{e_0\}.\ e_0 \stackrel{\mathsf{mo}}{\rightarrow} e$. That is, there exists $\xi_i \in \mathrm{Tx}_{\mathtt{vx}}$ such that

either $e_0 = \xi_i.SL_{\texttt{vx}}$ or $e_0 = \xi_i.U_{\texttt{vx}}$. Let us assume that $e_0 = \xi_i.U_{\texttt{vx}}$; we then have $\xi_i.U_{\texttt{vx}} \xrightarrow{\textsf{mo}} \xi_i.SL_{\texttt{vx}}$. On the other hand, since we have $\xi_i.SL_{\texttt{vx}} \xrightarrow{\textsf{po}} \xi_i.U_{\texttt{vx}}$, we have $\xi_i.SL_{\texttt{vx}} \xrightarrow{\textsf{po}} \xi_i.U_{\texttt{vx}} \xrightarrow{\textsf{mo}} \xi_i.SL_{\texttt{vx}}$, contradicting the assumption that G is consistent. We thus know that the minimal element is $e_0 = \xi_i.SL_{\texttt{vx}}$ for some $\xi_i \in \text{Tx}_{\texttt{vx}}$.

From the totality of mo on $\mathcal{W}_{\texttt{vx}}$, we know that there exists $e_1 \in \mathcal{W}_{\texttt{vx}} \setminus \{e_0\}$ such that $e_0 \xrightarrow{\textsf{mo}|_{\text{imm}}} e_1$. That is, either $e_1 = \xi_i.U_{\texttt{vx}}$; or there exists $j \neq i$ such that $e_1 = \xi_j.SL_{\texttt{vx}}$ or $e_1 = \xi_i.U_{\texttt{vx}}$. Let us pick an arbitrary $j \neq i$ and assume that $e_1 = \xi_j.SL_{\texttt{vx}}$. Since $e_0 \xrightarrow{\textsf{mo}|_{\text{imm}}} e_1$, the value read by $e_1 = \xi_j.SL_{\texttt{vx}}$, must be that written by $e_0 = \xi_i.SL_{\texttt{vx}}$. However, the value written by e_0 is an odd number, whilst the value read by e_0 is an even number. We thus know that $e_1 \neq \xi_j.SL_{\texttt{vx}}$ for all $j \neq i$. Similarly, let us pick an arbitrary $j \neq i$ and assume that $e_1 = \xi_j.U_{\texttt{vx}}$. We then have $\xi_j.U_{\texttt{vx}} \xrightarrow{\textsf{mo}} \xi_j.SL_{\texttt{vx}}$. On the other hand, since we have $\xi_j.SL_{\texttt{vx}} \xrightarrow{\textsf{po}} \xi_j.U_{\texttt{vx}}$, we have $\xi_j.SL_{\texttt{vx}} \xrightarrow{\textsf{po}} \xi_j.U_{\texttt{vx}} \xrightarrow{\textsf{mo}} \xi_j.SL_{\texttt{vx}}$, contradicting the assumption that G is consistent. We thus know that $e_1 \neq \xi_j.U_{\texttt{vx}}$ for all $j \neq i$. Consequently we have $e_1 = \xi_i.U_{\texttt{vx}}$.

Let $\text{Tx'}_{\texttt{vx}} = \text{Tx'}_{\texttt{vx}} \setminus \{\xi_i\}$. From the inductive hypothesis we then know there exist $L' = \texttt{perm}(\text{Tx'}_{\texttt{vx}})$ such that

$$L'|_1.SL_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} L'|_1.U_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \cdots \xrightarrow{\textsf{mo}|_{\text{imm}}} L'|_{|L'|}.SL_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} L'|_{|L'|}.U_{\texttt{vx}}$$

where $L'|_i$ denotes the i^{th} element of L'. On the other hand, since we have $e_0{=}\xi_i.SL_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} e_1{=}\xi_i.U_{\texttt{vx}}$ and e_0 is the minimal element according to mo, we then have:

$$\begin{array}{l} \xi_i.SL_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \xi_i.U_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \\ L'|_1.SL_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} L'|_1.U_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \cdots \xrightarrow{\textsf{mo}|_{\text{imm}}} L'|_{|L'|}.SL_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} L'|_{|L'|}.U_{\texttt{vx}} \end{array}$$

as required.

Proof (part 2). From part 1 we know there exists L, i, j such that $L|_1.SL_{\texttt{x1}} \xrightarrow{\textsf{mo}|_{\text{imm}}} L|_1.U_{\texttt{x1}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \cdots \xrightarrow{\textsf{mo}|_{\text{imm}}} L|_{|L|}.SL_{\texttt{x1}} \xrightarrow{\textsf{mo}|_{\text{imm}}} L|_{|L|}.U_{\texttt{x1}}$ and $L|_i = \xi_1$, $L|_j = \xi_2$, ans either $i < j$ or $j < i$.

Let us assume the former case. Since each $U_{\texttt{x1}}$ event is a `rel` write event and each $SL_{\texttt{x1}}$ event is an `acqrel` update event, we have $\cdots \xi_1.U_{\texttt{x1}} \xrightarrow{\textsf{rf}} L|_{i+1}.SL_{\texttt{x1}} \xrightarrow{\textsf{po}} L|_{i+1}.U_{\texttt{x1}} \xrightarrow{\textsf{rf}} \cdots \xrightarrow{\textsf{rf}} \xi_2.SL_{\texttt{x1}}$. On the other hand, since $\textsf{hb} = (\textsf{po} \cup \textsf{rf})^+$, we have $\xi_1.U_{\texttt{x1}} \xrightarrow{\textsf{hb}} \xi_2.SL_{\texttt{x1}}$ as required. The proof of the latter case in analogous and is omitted here.

Proof (part 3). From part 1 we know that the write events in $G.\mathcal{W}_{\texttt{vx}}$ are ordered by mo as follows, where $L = [\xi_1 \cdots \xi_m] = \texttt{perm}(\text{Tx}_{\texttt{vx}})$:

$$\xi_1.SL_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \xi_1.U_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \cdots \xrightarrow{\textsf{mo}|_{\text{imm}}} \xi_m.SL_{\texttt{vx}} \xrightarrow{\textsf{mo}|_{\text{imm}}} \xi_m.U_{\texttt{vx}}$$

As such, the values written to $\mathtt{vx}$ by the write events ordered as above monotonically increase: each $\xi_i.SL_{\mathtt{vx}}$ event increments the value of $\mathtt{vx}$ by one (it updates $\mathtt{vx}$ from v to $v{+}1$); while each subsequent $\xi_i.U_{\mathtt{vx}}$ event increments the value of $\mathtt{vx}$ by one (it updates $\mathtt{vx}$ from $v{+}1$ to $v{+}2$). Consequently, each value written by the write events ordered above is unique. ☐

Lemma 8. *For all consistent implementation execution graphs G and their counterpart specification graph Γ constructed as above,*

1. $\Gamma.\mathsf{po} \subseteq G.\mathsf{po}$
2. $\Gamma.\mathsf{rf} \subseteq G.\mathsf{hb}$
3. $\Gamma.(\mathsf{mo_T} \cup [\mathcal{NT}];\mathsf{rf};\mathsf{st}) \subseteq G.\mathsf{hb}$
4. $\forall i \in \mathbb{N}.\ \Gamma.\mathsf{rpsi\text{-}hb}_i \subseteq G.\mathsf{hb}$, *where* $\Gamma.\mathsf{rpsi\text{-}hb}_0 = \Gamma.(\mathsf{po} \cup \mathsf{rf} \cup \mathsf{mo_T} \cup [\mathcal{NT}];\mathsf{rf};\mathsf{st})$, *and for all* $i > 0$,
   $\Gamma.\mathsf{rpsi\text{-}hb}_{i+1} = \mathsf{st};([|\mathcal{W}];\mathsf{st};\mathsf{rpsi\text{-}hb}_i;\mathsf{st};[\mathcal{R}])_{loc};\mathsf{st}$.
5. $\Gamma.\mathsf{rpsi\text{-}hb} \subseteq G.\mathsf{hb}$

Proof (Part 1). Immediate from the definitions of $\Gamma.\mathsf{po}$ and $G.\mathsf{po}$.

Proof (part 2). Pick an arbitrary $(w, r) \in \Gamma.\mathsf{rf}$ and let $\mathtt{loc}(w) = \mathtt{loc}(r) = \mathtt{x}$. There are then four cases to consider: 1) $r \in \Gamma.\mathcal{NT}$; 2) $w \in \Gamma.\mathcal{NT}, r \in \Gamma.\mathcal{T}$; or 3) $w, r \in \Gamma.\mathcal{T} \wedge [w]_{\mathsf{st}} = [r]_{\mathsf{st}}$, or 4) $w, r \in \Gamma.\mathcal{T} \wedge [w]_{\mathsf{st}} \neq [r]_{\mathsf{st}}$.

Case 1. $r \in \Gamma.\mathcal{NT}$
Since $r \in \Gamma$, from the construction of $\Gamma \mathsf{rf}$ we then have $(w, r) \in G.\mathsf{rf}$. Consequently, since $G.\mathsf{rf} \subseteq G.\mathsf{hb}$, we have $(w, r) \in G.\mathsf{hb}$, as required.

Case 2. $w \in \Gamma.\mathcal{NT}, r \in \Gamma.\mathcal{T}$
Pick arbitrary ξ_i such that $r \in \mathcal{T}_{\xi_i}$. Let $\mathtt{val_w}(w) = \mathtt{val_r}(r) = v$. From the construction of Γ we know there exist rx_1 such that $rx_1 = \mathtt{R}(\mathtt{x}, v)$ and $rx_1 \stackrel{G.\mathsf{po}^*}{\to} r$, and $(w, rx_1) \in G.\mathsf{rf}$. We thus have $w \stackrel{G.\mathsf{rf}}{\to} rx_1 \stackrel{G.\mathsf{po}^*}{\to} r$. As $G.\mathsf{rf} \subseteq G.hb$ and $G.\mathsf{hb}$ is transitively closed, we have $w \stackrel{G.\mathsf{hb}}{\to} a$, as required.

Case 3. $w, r \in \Gamma.\mathcal{T} \wedge [w]_{\mathsf{st}} = [r]_{\mathsf{st}}$
From the construction of Γ we know there exists ξ $w, r \in \xi$, and $\xi_1.SL_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} w \stackrel{G.\mathsf{po}}{\to} \xi_1.U_{\mathtt{vx}}$. Let $r = \mathtt{R}(\mathtt{x}, v)$. From the construction of Γ we then know there exists $rx = \mathtt{R}(\mathtt{x}, v)$ such that $(w, rx) \in G.\mathsf{rf}$ and $(rx, r) \in G.\mathsf{po}$. On the other hand, given the shape of the traces of our implementation we know that $(rx, w) \in G.\mathsf{po}$. As such, we have $w \stackrel{G\mathsf{rf}}{\to} rx \stackrel{G.\mathsf{po}}{\to} w$. Since $G.\mathsf{po} \subseteq G.\mathsf{hb}$, we have $w \stackrel{G.\mathsf{rf}}{\to} rx \stackrel{G.\mathsf{hb}}{\to} w$, contradicting the assumption that G is consistent.

Case 4. $w, r \in \Gamma.\mathcal{T} \wedge [w]_{\mathsf{st}} \neq [r]_{\mathsf{st}}$
From the construction of Γ we know there exist ξ_1 and ξ_2 such that $w \in \xi_1$, $r \in \xi_2$, $\xi_1.SL_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} w \stackrel{G.\mathsf{po}}{\to} \xi_1.U_{\mathtt{vx}}$. Let $w = \mathtt{W}(\mathtt{x}, v)$ and $r = \mathtt{R}(\mathtt{x}, v)$. From the construction

of Γ we know there exists $rv_1, rx, rv_2 \in G.E$ and b such that $rv_1 = \mathtt{R}(\mathtt{vx}, b)$, $rx = \mathtt{R}(\mathtt{x}, v)$, $rv_2 = \mathtt{R}(\mathtt{vx}, b)$, $(w, rx) \in G.\mathsf{rf}$, and $rv_1 \stackrel{G.\mathsf{po}}{\rightarrow} rx \stackrel{G.\mathsf{po}}{\rightarrow} rv_2 \stackrel{G.\mathsf{po}}{\rightarrow} r$.

There are two cases to consider: A) either $\mathtt{x} \in \mathtt{WS}_{\xi_2}$; or B) $\mathtt{x} \notin \mathtt{WS}_{\xi_2}$. In the former case (A), from Lemma 7.2 we then know that either i) $\xi_2.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_1.SL_{\mathtt{vx}}$; or ii) $\xi_1.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_1.SL_{\mathtt{vx}}$.

In case (A.i) we then have $\xi_2.U_{\mathtt{vx}} \stackrel{G.\mathsf{mo}}{\rightarrow} \xi_1.SL_{\mathtt{vx}}$ (since otherwise we would have a cycle $\xi_2.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_1.SL_{\mathtt{vx}} \stackrel{G.\mathsf{mo}}{\rightarrow} \xi_2.U_{\mathtt{vx}}$, contradicting our assumption that G is consistent). As such we have $(rv_2, \xi_1.SL_{\mathtt{vx}}) \in G.\mathsf{rb}$. We then have $rv_2 \stackrel{G.\mathsf{rb}}{\rightarrow} \xi_1.SL_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\rightarrow} w \stackrel{G.\mathsf{rf}}{\rightarrow} rx \stackrel{G.\mathsf{po}}{\rightarrow} rv_2$. As $G.\mathsf{rf} \subseteq \mathsf{hb}$ and $G.\mathsf{po} \subseteq G.\mathsf{hb}$, we then have $rv_2 \stackrel{G.\mathsf{rb}}{\rightarrow} \xi_1.SL_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\rightarrow} rv_2$, contradicting the assumption that G is consistent.

In case (A.ii) we then have $w \stackrel{G.\mathsf{po}^*}{\rightarrow} \xi_1.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_2.SL_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\rightarrow} r$. That is, since we have $G.\mathsf{po} \subseteq G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we have $w \stackrel{G.\mathsf{hb}}{\rightarrow} r$, as required.

In the latter case (B) we then know b (in $rv_1 = \mathtt{R}(\mathtt{vx}, b)$) is even. Additionally, since write events on $\mathtt{vx}$ have unique values, we know that either i) rv_1 reads from the initial write to $\mathtt{vx}$ and we thus have $rv_1 \stackrel{G.\mathsf{rb}}{\rightarrow} \xi_1.SL_{\mathtt{vx}}$ and $rv_2 \stackrel{G.\mathsf{rb}}{\rightarrow} \xi_1.SL_{\mathtt{vx}}$; or ii) there exists ξ_3 such that $\mathtt{x} \in \mathtt{WS}_{\xi_3}$, $\xi_3.SL_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\rightarrow} \xi_3.U_{\mathtt{vx}}$ and $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{rf}}{\rightarrow} rv_1$.

In case (B.i) we have $rv_2 \stackrel{G.\mathsf{rb}}{\rightarrow} \xi_1.SL_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\rightarrow} w \stackrel{G.\mathsf{rf}}{\rightarrow} rx \stackrel{G.\mathsf{po}}{\rightarrow} rv_2$. As $G.\mathsf{rf} \subseteq \mathsf{hb}$ and $G.\mathsf{po} \subseteq G.\mathsf{hb}$, we then have $rv_2 \stackrel{G.\mathsf{rb}}{\rightarrow} \xi_1.SL_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\rightarrow} rv_2$, contradicting the assumption that G is consistent.

In case (B.ii), since we have $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{rf}}{\rightarrow} rv_1$ and each write event on $\mathtt{vx}$ writes a unique value (Lemma 7.3), we also have $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{rf}}{\rightarrow} rv_2$. That is, $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\rightarrow} rv_1$, $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\rightarrow} rv_2$. On the other hand, from Lemma 7.2 we know that either a) $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_1.SL_{\mathtt{vx}}$; or b) $\xi_1.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_3.SL_{\mathtt{vx}}$.

In case (B.ii.a), since $G.\mathsf{mo}$ on $\mathtt{vx}$ is totally ordered, from the consistency of Γ we know that $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{mo}}{\rightarrow} \xi_1.SL_{\mathtt{vx}}$ (since otherwise we would have a cycle $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_1.SL_{\mathtt{vx}} \stackrel{G.\mathsf{mo}}{\rightarrow} \xi_3.U_{\mathtt{vx}}$, contradicting $\mathsf{RA\text{-}consistent}(G)$). Consequently, since we have $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{rf}}{\rightarrow} rv_2$, and $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{mo}}{\rightarrow} \xi_1.SL_{\mathtt{vx}}$, we have $rv_2 \stackrel{G.\mathsf{rb}}{\rightarrow} \xi_1.SL_{\mathtt{vx}}$. We thus have $rv_2 \stackrel{G.\mathsf{rb}}{\rightarrow} \xi_1.SL_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\rightarrow} w \stackrel{G.\mathsf{rf}}{\rightarrow} rx \stackrel{G.\mathsf{po}}{\rightarrow} rv_2$. As $G.\mathsf{rf} \subseteq \mathsf{hb}$ and $G.\mathsf{po} \subseteq G.\mathsf{hb}$, we have $rv_2 \stackrel{G.\mathsf{rb}}{\rightarrow} \xi_1.SL_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\rightarrow} rv_2$, contradicting the assumption that G is consistent.

In case (B.ii.b) we have $\xi_1.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\rightarrow} \xi_3.SL_{\mathtt{vx}}$. Recall that we also have $w \stackrel{\mathsf{po}}{\rightarrow} \xi_1.U_{\mathtt{vx}}$, $\xi_3.SL_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\rightarrow} \xi_3.U_{\mathtt{vx}}$, $\xi_3.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\rightarrow} rv_2$, and $rv_2 \stackrel{\mathsf{po}}{\rightarrow} r$. As $G.\mathsf{po} \in G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we thus have $w \stackrel{G.\mathsf{hb}}{\rightarrow} r$, as required.

Proof (part 3). We show that $\Gamma.\mathsf{mo}_{\mathsf{T}} \subseteq G.\mathsf{hb}$, and $\Gamma.([\mathcal{NT}]; \mathsf{rf}; \mathsf{st}) \subseteq G.\mathsf{hb}$.

RTS. $\Gamma.\mathsf{mo}_{\mathsf{T}} \subseteq G.\mathsf{hb}$
Pick an arbitrary $(a, b) \in \Gamma.\mathsf{mo}_{\mathsf{T}}$; we then need to show that $(a, b) \in G.\mathsf{hb}$.

From the definition of $\Gamma.\mathsf{mo_T}$ and the construction of Γ we know there exist ξ_1, ξ_2, c, d such that $\xi_1 \neq \xi_2$, $(c, d) \in \Gamma.\mathsf{mo}$, $a, c \in \xi_1$, $b, d \in \xi_2$. Let $\mathtt{loc}(c) = \mathtt{loc}(d) = \mathtt{x}$. We then know $a \stackrel{G.\mathsf{po}^*}{\to} \xi_1.U_{\mathtt{vx}}$, $\xi_1.SL_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} c \stackrel{G.\mathsf{po}}{\to} \xi_1.U_{\mathtt{vx}}$, $\xi_2.SL_{\mathtt{vx}} \stackrel{G.\mathsf{po}^*}{\to} b$ and $\xi_2.SL_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} d \stackrel{G.\mathsf{po}}{\to} \xi_2.U_{\mathtt{vx}}$.

From Lemma 7.2 we then know that either $\xi_1.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_2.SL_{\mathtt{vx}}$, or $\xi_2.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_1.SL_{\mathtt{vx}}$. Let us assume that the latter holds. We then have $d \stackrel{G.\mathsf{po}}{\to} \xi_2.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_1.SL_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} c \stackrel{G.\mathsf{mo}}{\to} d$. That is, since $G.\mathsf{po} \in G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we have $d \stackrel{G.\mathsf{hb}}{\to} c \stackrel{G.\mathsf{mo}}{\to} d$, contradicting the assumption that G is consistent. We thus know that $\xi_1.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_2.SL_{\mathtt{vx}}$. As such, we have $a \stackrel{G.\mathsf{po}^*}{\to} \xi_1.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_2.SL_{\mathtt{vx}} \stackrel{G.\mathsf{po}^*}{\to} b$. As $G.\mathsf{po} \in G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we have $a \stackrel{G.\mathsf{hb}}{\to} b$, as required.

RTS. $\Gamma.([\mathcal{NT}]; \mathsf{rf}; \mathsf{st}) \subseteq G.\mathsf{hb}$
Pick arbitrary ξ_i, w, r, a such that $r, a \in \mathcal{T}_{\xi_i}$ and $(w, r) \in \Gamma.\mathsf{rf}$. We are then required to show $(w, a) \in G.\mathsf{hb}$. Let $\mathtt{loc}(w) = \mathtt{loc}(r) = \mathtt{x}$ and $\mathtt{val_w}(w) = \mathtt{val_r}(r) = v$. From the construction of Γ we know there exist rx_1 such that $rx_1 = \mathtt{R}(\mathtt{x}, v)$ and $rx_1 \stackrel{G.\mathsf{po}^*}{\to} a$, and $(w, rx_1) \in G.\mathsf{rf}$. We thus have $w \stackrel{G.\mathsf{rf}}{\to} rx_1 \stackrel{G.\mathsf{po}^*}{\to} a$. As $G.\mathsf{rf} \subseteq G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we have $w \stackrel{G.\mathsf{hb}}{\to} a$, as required.

Proof (part 4). We proceed by induction on i.

Base case: $i = 0$
The proof of this case is immediate from the definition of $\Gamma.\mathsf{rpsi\text{-}hb}_0$ and parts 1-3.

Inductive case: $i = n+1$

$$\forall j \leq n.\ \Gamma.(\mathsf{st}; ([\mathcal{W}]; \mathsf{st}; \mathsf{rpsi\text{-}hb}_j; \mathsf{st}; [\mathcal{R}])_{loc}; \mathsf{st}) \subseteq G.\mathsf{hb} \qquad \text{(I.H.)}$$

Pick arbitrary $(a, b) \in \Gamma.\big(\mathsf{st}; ([\mathcal{W}]; \mathsf{st}; \mathsf{rpsi\text{-}hb}_i; \mathsf{st}; [\mathcal{R}])_{loc}; \mathsf{st}\big)$. We are then required to show $(a, b) \in G.\mathsf{hb}$.

From the definition of $\Gamma.\big(\mathsf{st}; ([\mathcal{W}]; \mathsf{st}; \mathsf{rpsi\text{-}hb}_i; \mathsf{st}; [\mathcal{R}])_{loc}; \mathsf{st}\big)$ we know there exist c, d, w, r, ξ_1, ξ_2, such that $\mathcal{T}_{\xi_1}, \mathcal{T}_{\xi_2} \in \Gamma.\mathcal{T}/\mathsf{st}$, $\xi_1 \neq \xi_2$, $a, c, w \in \mathcal{T}_{\xi_1}$, $w \in \Gamma.\mathcal{W}$, $b, d, r \in \mathcal{T}_{\xi_2}$, $r \in \Gamma.\mathcal{R}$, $\mathtt{loc}(w) = \mathtt{loc}(r)$ and $(c, d) \in \Gamma.\mathsf{rpsi\text{-}hb}_n$.

Let $\mathtt{loc}(w) = \mathtt{loc}(r) = \mathtt{x}$ and $\mathtt{val_r}(r) = v$. Since $\mathtt{loc}(w) = \mathtt{x}$, $w \in \Gamma.\mathcal{W}$ and $w \in \mathcal{T}_{\xi_1}$, from the construction of Γ we have $\mathtt{x} \in \mathtt{WS}_{\mathcal{T}'_{\xi_1}}$ and that $\xi_1.SL_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} c \stackrel{G.\mathsf{po}}{\to} \xi_1.U_{\mathtt{vx}}$ and $\xi_1.SL_{\mathtt{vx}} \stackrel{G.\mathsf{po}}{\to} a \stackrel{G.\mathsf{po}}{\to} \xi_1.U_{\mathtt{vx}}$. Similarly, since $\mathtt{loc}(r) = \mathtt{x}$, $r \in \Gamma.\mathcal{R}$ and $r \in \mathcal{T}_{\xi_1}$, from the construction of Γ we have $\mathtt{x} \in \mathtt{RS}_{\mathcal{T}'_{\xi_2}}$ and that there exist $rv_1, rv_2, rx \in \mathcal{T}'_{\xi_2}$ and e, such that $rv_1 = \mathtt{R}(\mathtt{vx}, e)$, $rv_2 = \mathtt{R}(\mathtt{vx}, e)$, $rx = \mathtt{R}(\mathtt{x}, v)$, $rv_1 \stackrel{G.\mathsf{po}}{\to} rx \stackrel{G.\mathsf{po}}{\to} rv_2 \stackrel{G.\mathsf{po}}{\to} d$, and $rv_1 \stackrel{G.\mathsf{po}}{\to} rx \stackrel{G.\mathsf{po}}{\to} rv_2 \stackrel{G.\mathsf{po}}{\to} b$.

There are now two cases to consider: 1) either $\mathtt{x} \in \mathtt{WS}_{\xi_2}$; or $\mathtt{x} \notin \mathtt{WS}_{\xi_2}$. In the former case (1), from Lemma 7.2 we then know that either i) $\xi_1.U_{\mathtt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_2.SL_{\mathtt{vx}}$;

or ii) $\xi_2.U_{\texttt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_1.SL_{\texttt{vx}}$. In case (1.i), we then have $a \stackrel{G.\mathsf{po}}{\to} \xi_1.U_{\texttt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_2.SL_{\texttt{vx}} \stackrel{G.\mathsf{hb}}{\to} b$. That is, as $G.\mathsf{po} \subseteq G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we have $a \stackrel{G.\mathsf{hb}}{\to} b$, as required. In case (1.ii), since we have $(c, d) \in \Gamma.\mathsf{rpsi\text{-}hb}_n$, from (I.H.), we have $(c, d) \in G.\mathsf{hb}$. As such, we have $d \stackrel{G.\mathsf{po}}{\to} \xi_2.U_{\texttt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_1.SL_{\texttt{vx}} \stackrel{G.\mathsf{po}}{\to} c \stackrel{G.\mathsf{hb}}{\to} d$. That is, since we have $G.\mathsf{po} \subseteq G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we have $c \stackrel{G.\mathsf{hb}}{\to} c$, contradicting the assumption that G is consistent.

In the latter case (2) we then know f (in rv_1 and rv_2) is even. As such, from our implementation we know there exists ξ_3 such that $\texttt{x} \in \texttt{WS}_{\xi_3}$, $\xi_3.SL_{\texttt{vx}} \stackrel{\mathsf{po}}{\to} \xi_3.U_{\texttt{vx}}$, and that $(\xi_3.U_{\texttt{vx}}, rv_1) \in G.\mathsf{rf}$. Since the values written to $\texttt{vx}$ are unique (Lemma 7.3) and $\texttt{val}_{\texttt{r}}(rv_1) = \texttt{val}_{\texttt{r}}(rv_2) = f$, we also have $(\xi_3.U_{\texttt{vx}}, rv_2) \in G.\mathsf{rf}$. On the other hand, from Lemma 7.2 we have either i) $\xi_1.U_{\texttt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_3.SL_{\texttt{vx}}$, or ii) $\xi_3.U_{\texttt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_1.SL_{\texttt{vx}}$. In case (2.i), we then have $a \stackrel{G.\mathsf{po}}{\to} \xi_1.U_{\texttt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_3.SL_{\texttt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_3.U_{\texttt{vx}} \stackrel{G.\mathsf{rf}}{\to} rv_1 \stackrel{G.\mathsf{po}}{\to} b$. That is, as $G.\mathsf{po}, G.\mathsf{rf} \subseteq G.\mathsf{hb}$ and $G.\mathsf{hb}$ is transitively closed, we have $a \stackrel{G.\mathsf{hb}}{\to} b$, as required.

In case (2.ii) we then have $\xi_3.U_{\texttt{vx}} \stackrel{G.\mathsf{mo}}{\to} \xi_1.SL_{\texttt{vx}}$ (since otherwise we would have a cycle $\xi_3.U_{\texttt{vx}} \stackrel{G.\mathsf{hb}}{\to} \xi_1.SL_{\texttt{vx}} \stackrel{G.\mathsf{mo}}{\to} \xi_3.U_{\texttt{vx}}$, contradicting our assumption that G is consistent). As such, we have $(rv_1, \xi_1.SL_{\texttt{vx}}) \in G.\mathsf{rb}$ and $(rv_2, \xi_1.SL_{\texttt{vx}}) \in G.\mathsf{rb}$. On the other hand, since we have $(c, d) \in \Gamma.\mathsf{rpsi\text{-}hb}_n$, from (I.H.) we have $(c, d) \in G.\mathsf{hb}$. It is straightforward to demonstrate that $G.\mathsf{hb} = G.(\mathsf{po} \cup \mathsf{rf})^+ = \Gamma.(\mathsf{po}^+ \cup (\mathsf{po} \cup \mathsf{rf})^*; \mathsf{rf} \setminus \mathsf{po}; \mathsf{po}^*)$. There are thus two cases to consider: a) $(c, d) \in G.\mathsf{po}^+$; or b) $(c, d) \in G.((\mathsf{po} \cup \mathsf{rf})^*; \mathsf{rf} \setminus \mathsf{po}; \mathsf{po}^*)$.

In case (2.ii.a), since $c \in \mathcal{T}'_{\xi_1}$, $d, rv_2 \in \mathcal{T}'_{\xi_2}$ and $\mathcal{T}'_{\xi_1} \neq \mathcal{T}'_{\xi_2}$, we know that $(c, rv_2) \in G.\mathsf{po}$. As such, we have $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.SL_{\texttt{vx}} \stackrel{G.\mathsf{po}}{\to} c \stackrel{G.\mathsf{po}}{\to} rv_2$. That is, we have $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.SL_{\texttt{vx}} \stackrel{G.\mathsf{hb}}{\to} rv_2$, contradicting the assumption that G is consistent.

In case (2.ii.b), we then know there exist m, n such that $(c, m) \in G.(\mathsf{po} \cup \mathsf{rf})^*$, $(m, n) \in G.(\mathsf{rf} \setminus \mathsf{po})$, $(n, d) \in G.\mathsf{po}^*$. There are now two additional cases to consider: either 1) $n \notin \mathcal{T}'_{\xi_2}$; or 2) $n \in \mathcal{T}'_{\xi_2}$.

In (2.ii.b.1), since $rv_2, d \in \mathcal{T}'_{\xi_2}$, $n \notin \mathcal{T}'_{\xi_2}$ and $(n, d) \in G.\mathsf{po}$, we also have $(n, rv_2) \in G.\mathsf{po}$. As such, we have $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.SL_{\texttt{vx}} \stackrel{G.\mathsf{po}}{\to} c \stackrel{G.(\mathsf{po} \cup \mathsf{rf})^*}{\to} m \stackrel{G.\mathsf{rf}}{\to} n \stackrel{G.\mathsf{po}^*}{\to} rv_2$. That is, we have $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.SL_{\texttt{vx}} \stackrel{G.\mathsf{hb}}{\to} rv_2$, contradicting the assumption that G is consistent.

In (2.ii.b.2), from our implementation we know that either i) $n \stackrel{G.\mathsf{po}}{\to} rv_2$; or ii) $rv_2 \stackrel{G.\mathsf{po}}{\to} n$. In (2.ii.b.2.i) we then have $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.SL_{\texttt{vx}} \stackrel{G.\mathsf{po}}{\to} c \stackrel{G.\mathsf{hb}}{\to} m \stackrel{G.\mathsf{rf}}{\to} n \stackrel{G.\mathsf{po}}{\to} rv_2$. That is, we have $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.SL_{\texttt{vx}} \stackrel{G.\mathsf{hb}}{\to} rv_2$, contradicting the assumption that G is consistent. In case (2.ii.b.2.ii) we then know that there exists n' such that $(m, n') \in G.\mathsf{rf}$ and $(n', rv_2) \in G.\mathsf{po}$. As such, we $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.SL_{\texttt{vx}} \stackrel{G.\mathsf{po}}{\to} c \stackrel{G.\mathsf{hb}}{\to} m \stackrel{G.\mathsf{rf}}{\to} n' \stackrel{G.\mathsf{po}}{\to} rv_2$. That is, we have $rv_2 \stackrel{G.\mathsf{rb}}{\to} \xi_1.SL_{\texttt{vx}} \stackrel{G.\mathsf{hb}}{\to} rv_2$, contradicting the assumption that G is consistent.

Proof (part 5). Immediate from parts 1-4 and the fact that $G.\mathsf{hb}$ is transitively closed. ◻

2.4 Auxiliary Completeness Lemmata

Given an execution graph of the implementation G, we write Tx for the set of transactions executed by the program.In what follows, we write $\mathcal{T}'_i$ for the set of events in the *implementation* trace θ_i; that is, $\mathcal{T}'_i \triangleq \theta_i.E$. In other words, $\mathcal{T}'_i$ corresponds to the set of events in the implementation of the specification transaction class $\mathcal{T}_i$.

Lemma 9. *For all consistent abstraction execution graphs Γ and their counterpart implementation graphs G constructed as above with transaction classes $\bigcup_{\mathcal{T}_i \in \Gamma.\mathcal{T}/\mathsf{st}} \mathcal{T}'_i$, for all $\mathcal{T}'_a, \mathcal{T}'_b$ and for all a, b, if $(a,b) \in G.\mathsf{hb}$, then*

$$\begin{aligned}
a, b \in G.\mathcal{NT} &\Rightarrow (a,b) \in \Gamma.\mathsf{rpsi\text{-}hb} \\
a \in G.\mathcal{NT} \land b \in \mathcal{T}'_b &\Rightarrow (\{a\} \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{rpsi\text{-}hb} \\
a \in \mathcal{T}'_a \land b \in G.\mathcal{NT} &\Rightarrow \exists c \in \mathcal{T}_a.\, (c,b) \in \Gamma.\mathsf{rpsi\text{-}hb} \\
&\quad\ \land (a \in \mathcal{T}_a \Rightarrow c{=}a) \\
a \in \mathcal{T}'_a \land b \in \mathcal{T}'_a \land \mathcal{T}'_a {\neq} \mathcal{T}'_b &\Rightarrow \exists c \in \mathcal{T}_a.\, (\{c\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb} \\
&\quad\ \land (a \in \mathcal{T}_a \Rightarrow c{=}a) \\
a, b \in \mathcal{T}'_a &\Rightarrow (a,b) \in G.\mathsf{po}
\end{aligned}$$

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

$$\begin{aligned}
&\forall i \in \mathbb{N}.\ \forall \mathcal{T}'_a, \mathcal{T}'_b.\ \forall a, b.\ (a,b) \in G.\mathsf{hb} \Rightarrow \\
&\qquad a, b \in G.\mathcal{NT} \Rightarrow (a,b) \in \Gamma.\mathsf{rpsi\text{-}hb} \\
&\qquad a \in G.\mathcal{NT} \land b \in \mathcal{T}'_b \Rightarrow (\{a\} \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{rpsi\text{-}hb} \\
&\qquad a \in \mathcal{T}'_a \land b \in G.\mathcal{NT} \Rightarrow \exists c \in \mathcal{T}_a.\, (c,b) \in \Gamma.\mathsf{rpsi\text{-}hb} \\
&\qquad\qquad \land (a \in \mathcal{T}_a \Rightarrow c{=}a) \\
&\qquad a \in \mathcal{T}'_a \land b \in \mathcal{T}'_a \land \mathcal{T}'_a {\neq} \mathcal{T}'_b \Rightarrow \exists c \in \mathcal{T}_a.\, (\{c\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb} \\
&\qquad\qquad \land (a \in \mathcal{T}_a \Rightarrow c{=}a) \\
&\qquad a, b \in \mathcal{T}'_a \Rightarrow (a,b) \in G.\mathsf{po}
\end{aligned}$$

Base case $i = 0$
Pick arbitrary $\mathcal{T}'_a, \mathcal{T}'_b$ and a, b such that $a \in \mathcal{T}'_a$, $b \in \mathcal{T}'_b$ and $(a,b) \in G.\mathsf{hb}_0$. There are then two cases to consider: 1) $(a,b) \in G.\mathsf{po}$; or 2) $(a,b) \in G.\mathsf{rf}$.

In case (1), we need to consider 5 cases: a) $a, b \in G.\mathcal{NT}$; or b) $a \in G.\mathcal{NT} \land b \in \mathcal{T}'_b$; or c) $a \in \mathcal{T}'_a \land b \in G.\mathcal{NT}$; or d) $a {\in} \mathcal{T}'_a \land b {\in} \mathcal{T}'_a \land \mathcal{T}'_a {\neq} \mathcal{T}'_b$; or e) $a, b \in \mathcal{T}'_a$. In case (1.a), from the definition $G.\mathsf{po}$ we simply have $(a,b) \in \Gamma.\mathsf{po}$ and thus $(a,b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In case (1.b), we then know that $(\{a\} \times \mathcal{T}'_b) \subseteq G.\mathsf{po}$. Consequently, from the definition $G.\mathsf{po}$ we have $(\{a\} \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{po}$, and thus $(\{a\} \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In case (1.c) we then know that $(\mathcal{T}'_a \times \{b\}) \subseteq G.\mathsf{po}$. Consequently, from the definition $G.\mathsf{po}$ we have $(\mathcal{T}_a \times \{b\}) \subseteq \Gamma.\mathsf{po}$. That is, $\exists c \in \mathcal{T}_a.\ (c,b) \in \Gamma.\mathsf{po}$ and $a \in \mathcal{T}_a \Rightarrow c{=}a$. As such, since $\Gamma.\mathsf{po} \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$, we have $\exists c \in \mathcal{T}_a.\ (c,b) \in \Gamma.\mathsf{rpsi\text{-}hb}$ and $a \in \mathcal{T}_a \Rightarrow c{=}a$, as required.

In case (1.d) we then know that $(\mathcal{T}'_a \times \mathcal{T}'_b) \subseteq G.\mathsf{po}$. Consequently, from the definition $G.\mathsf{po}$ we have $(\mathcal{T}_a \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{po}$. That is, $\exists c \in \mathcal{T}_a.\ (\{c\} \times \mathcal{T}_b) \in \Gamma.\mathsf{po}$ and $a \in \mathcal{T}_a \Rightarrow c{=}a$. As such, since $\Gamma.\mathsf{po} \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$, we have $\exists c \in \mathcal{T}_a.\ (\{c\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$ and $a \in \mathcal{T}_a \Rightarrow c{=}a$, as required.

In case (1.e) the desired result holds immediately.

In case (2) we again need to consider 5 cases: a) $a, b \in G.\mathcal{NT}$; or b) $a \in G.\mathcal{NT} \wedge b \in \mathcal{T}'_b$; or c) $a \in \mathcal{T}'_a \wedge b \in G.\mathcal{NT}$; or d) $a {\in} \mathcal{T}'_a \wedge b {\in} \mathcal{T}'_a \wedge \mathcal{T}'_a {\neq} \mathcal{T}'_b$; or e) $a, b \in \mathcal{T}'_a$.

In case (2.a), from the definition $G.\mathsf{rf}$ we simply have $(a, b) \in \Gamma.\mathsf{rf}$ and thus $(a, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In case (2.b), we then know that there exists $c \in \mathcal{T}_b$ such that $(a, c) \in \Gamma.\mathsf{rf}$. As such, we have $(\{a\} \times \mathcal{T}_b) \subseteq \Gamma.([\mathcal{NT}]; \mathsf{rf}; \mathsf{st})$. Since $\Gamma.([\mathcal{NT}]; \mathsf{rf}; \mathsf{st}) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$, we have $(\{a\} \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In case (2.c), from the definition of $G.\mathsf{rf}$ we then know that $(a, b) \in \Gamma.\mathsf{rf}$. Since $\Gamma.\mathsf{rf} \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$, we thus have $(a, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In case (2.d), since $\mathcal{T}'_a \neq \mathcal{T}'_b$, we have $\mathcal{T}_a \neq \mathcal{T}_b$. There are now two cases to consider: i) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{x}$ for some shared location $\mathtt{x}$; or ii) $\mathtt{loc}(a) = \mathtt{loc}(b) = \mathtt{vx}$, for a version lock $\mathtt{vx}$ associated with some location $\mathtt{x}$. In (2.d.i), from the definition of $G.\mathsf{rf}$ we know that there exists $c \in \mathcal{T}_b$ such that $(a, c) \in \Gamma.\mathsf{rf}$. As such, since $\Gamma.\mathsf{rf} \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$, we have $(\mathcal{T}_a \times \mathcal{T}_b) \subseteq \Gamma.\Big(\mathsf{st}; \big([\mathcal{W}]; \mathsf{st}; \mathsf{rf} \setminus \mathsf{st}; \mathsf{st}; [\mathcal{R}]\big)_{loc}; \mathsf{st}\Big) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. That is, $(\{a\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required. In (2.d.ii), from the construction of $G.\mathsf{rf}$ we then know that there exists $c = \mathtt{W}(\mathtt{x}, -) \in \mathcal{T}_a$ and $d = \mathtt{R}(\mathtt{x}, -) \in \mathcal{T}_b$, such that $(c, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As such, we have $(\mathcal{T}_a \times \mathcal{T}_b) \subseteq \Gamma.\Big(\mathsf{st}; \big([\mathcal{W}]; \mathsf{st}; \mathsf{rf} \setminus \mathsf{st}; \mathsf{st}; [\mathcal{R}]\big)_{loc}; \mathsf{st}\Big) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. In particular we have $\exists c \in \mathcal{T}_a.\ (\{c\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In case (2.e), since a, b are both in $\mathcal{T}'_j$, from the construction of $G.\mathsf{rf}$ we know that there exists a location lock $\mathtt{vx}$ and a value v such that $b = \mathtt{R}(\mathtt{x}, v)$, $a = \mathtt{U}(\mathtt{x}, v{-}1, v)$ and that $(a, b) \in G.\mathsf{po}$, as required.

Inductive case $i = n{+}1$

$$\begin{array}{rl}
\forall i \leq n.\ \forall \mathcal{T}'_a, \mathcal{T}'_b.\ \forall a, b.\ (a, b) \in G.\mathsf{hb} \Rightarrow & \\
a, b \in G.\mathcal{NT} \Rightarrow & (a, b) \in \Gamma.\mathsf{rpsi\text{-}hb} \\
a \in G.\mathcal{NT} \wedge b \in \mathcal{T}'_b \Rightarrow & (\{a\} \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{rpsi\text{-}hb} \\
a \in \mathcal{T}'_a \wedge b \in G.\mathcal{NT} \Rightarrow & \exists c \in \mathcal{T}_a.\ (c, b) \in \Gamma.\mathsf{rpsi\text{-}hb} \\
& \wedge\, (a \in \mathcal{T}_a \Rightarrow c{=}a) \\
a \in \mathcal{T}'_a \wedge b \in \mathcal{T}'_a \wedge \mathcal{T}'_a {\neq} \mathcal{T}'_b \Rightarrow & \exists c \in \mathcal{T}_a.\ (\{c\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb} \\
& \wedge\, (a \in \mathcal{T}_a \Rightarrow c{=}a) \\
a, b \in \mathcal{T}'_a \Rightarrow & (a, b) \in G.\mathsf{po}
\end{array} \tag{I.H.}$$

Pick arbitrary $\mathcal{T}'_a, \mathcal{T}'_b$ and a, b such that $a \in \mathcal{T}'_a$, $b \in \mathcal{T}'_b$ and $(a, b) \in G.\mathsf{hb}_0$. From the definition of hb_{n+1} we then know there exists c such that $(a, c) \in \mathsf{hb}_0$ and $(c, b) \in \mathsf{hb}_n$. We now need to consider five cases: 1) $a, b \in G.\mathcal{NT}$; or 2) $a \in G.\mathcal{NT} \wedge b \in \mathcal{T}'_b$; or 3) $a \in \mathcal{T}'_a \wedge b \in G.\mathcal{NT}$; or 4) $a {\in} \mathcal{T}'_a \wedge b {\in} \mathcal{T}'_a \wedge \mathcal{T}'_a {\neq} \mathcal{T}'_b$; or

5) $a, b \in \mathcal{T}'_a$.

Case 1. $a, b \in G.\mathcal{NT}$

There are two cases to consider: a) $c \in G.\mathcal{NT}$; or b) $c \in \mathcal{T}'_c$ for some $\mathcal{T}'_c$.

In case (1.a), from the proof of the base case we have $(a, c) \in \Gamma.\mathsf{rpsi\text{-}hb}$. On the other hand, from (I.H.) we have $(c, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$. Since $\Gamma.\mathsf{rpsi\text{-}hb}$ is transitively closed, we have $(a, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In case (1.b), from the proof of the base case we have $(\{a\} \times \mathcal{T}_c) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. On the other hand, from (I.H.) we have $\exists d \in \mathcal{T}_c.\ (d, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$. Since $\Gamma.\mathsf{rpsi\text{-}hb}$ is transitively closed, we have $(a, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

Case 2. $a \in G.\mathcal{NT} \wedge b \in \mathcal{T}'_b$

There are two cases to consider: a) $c \in G.\mathcal{NT}$; or b) $c \in \mathcal{T}'_c$ for some $\mathcal{T}'_c \in \bigcup_{\mathcal{T}_i \in \Gamma.\mathcal{T}/\mathsf{st}} \mathcal{T}'_i$.

In case (2.a), from the proof of the base case we have $(a, c) \in \Gamma.\mathsf{rpsi\text{-}hb}$. On the other hand, from (I.H.) we have $(\{c\} \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. Since $\Gamma.\mathsf{rpsi\text{-}hb}$ is transitively closed, we have $(\{a\} \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In case (2.b), from the proof of the base case we have $(\{a\} \times \mathcal{T}_c) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. There are now two cases to consider: i) $\mathcal{T}'_c \neq \mathcal{T}'_b$; or ii) $\mathcal{T}'_c = \mathcal{T}'_b$.

In case (2.b.i), from (I.H.) we have $\exists d \in \mathcal{T}_c.\ (\{d\} \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. Consequently, as we have $(\{a\} \times \mathcal{T}_c) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$ and $\Gamma.\mathsf{rpsi\text{-}hb}$ is transitively closed, we have $(\{a\} \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In case (2.b.ii), since we have $(\{a\} \times \mathcal{T}_c) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$ and $\mathcal{T}_c = \mathcal{T}_b$, we have $(\{a\} \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

Case 3. $a \in \mathcal{T}'_a \wedge b \in G.\mathcal{NT}$

There are two cases to consider: a) $c \in G.\mathcal{NT}$; or b) $c \in \mathcal{T}'_c$ for some $\mathcal{T}'_c$.

In case (3.a), from the proof of the base case we $\exists d \in \mathcal{T}_a.\ (d, c) \in \Gamma.\mathsf{rpsi\text{-}hb} \wedge (a \in \mathcal{T}_a \Rightarrow d = a)$. On the other hand, from (I.H.) we have $(c, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$. Since $\Gamma.\mathsf{rpsi\text{-}hb}$ is transitively closed, we have $\exists d \in \mathcal{T}_a.\ (d, b) \in \Gamma.\mathsf{rpsi\text{-}hb} \wedge (a \in \mathcal{T}_a \Rightarrow d = a)$, as required.

In case (3.b), from (I.H.) we have $\exists e \in \mathcal{T}_c.\ (e, b) \in \Gamma.\mathsf{rpsi\text{-}hb} \wedge (c \in \mathcal{T}_c \Rightarrow e = c)$. There are now two cases to consider: i) $\mathcal{T}'_a \neq \mathcal{T}'_c$; or ii) $\mathcal{T}'_a = \mathcal{T}'_c$.

In case (3.b.i), from the base case we have $\exists d \in \mathcal{T}_a.\ (\{d\} \times \mathcal{T}_c) \in \Gamma.\mathsf{rpsi\text{-}hb} \wedge (a \in \mathcal{T}_a \Rightarrow d = a)$. On the other hand, from (I.H.) we have $\exists e \in \mathcal{T}_c.\ (e, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As $\Gamma \mathsf{hb}$ is transitively closed, we have $\exists d \in \mathcal{T}_a.\ (d, b) \in \Gamma.\mathsf{rpsi\text{-}hb} \wedge (a \in \mathcal{T}_a \Rightarrow d = a)$, as required.

In case (3.b.ii), from the proof of the base case we then have $(a, c) \in G.\mathsf{po}$. Recall that we have $\exists e \in \mathcal{T}_c.\ (e, b) \in \Gamma.\mathsf{rpsi\text{-}hb} \wedge (c \in \mathcal{T}_c \Rightarrow e = c)$. That is, since $\mathcal{T}_c = \mathcal{T}_a$, we have $\exists e \in \mathcal{T}_a.\ (e, b) \in \Gamma.\mathsf{rpsi\text{-}hb} \wedge (c \in \mathcal{T}_a \Rightarrow e = c)$. There are now three cases to consider: 1) $a \notin \mathcal{T}_a$; 2) $a, c \in \mathcal{T}_a$; 3) $a \in \mathcal{T}_a$ and $c \notin \mathcal{T}_a$. In case (3.b.ii.1) we have $\exists e \in \mathcal{T}_a.\ (e, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required. In case (3.b.ii.2), we then have $(c, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$. On the other hand, since we have $(a, c) \in G.\mathsf{po}$ and $G.\mathsf{po}$ does not change the orderings between events of Γ, we also have $(a, c) \in \Gamma.\mathsf{po}$. As

$\Gamma.\mathsf{po} \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$ and $\Gamma.\mathsf{rpsi\text{-}hb}$ is transitively closed, we have $(a, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In case (3.b.ii.3), we are required to show that $(a, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$. it is easy to demonstrate that $\mathsf{hb}_i = G.(\mathsf{po}^+ \cup \mathsf{po}^*; \mathsf{rf} \setminus \mathsf{po}; (\mathsf{po} \cup \mathsf{rf})^*)$. That is, as we have $(c, b) \in \mathsf{hb}_i$, we either have a) $(c, b) \in G.\mathsf{po}^+$; or b) $(c, b) \in G.(\mathsf{po}^*; \mathsf{rf} \setminus \mathsf{po}; (\mathsf{po} \cup \mathsf{rf})^*)$. In (3.b.ii.3.a), we then have $(a, b) \in G.\mathsf{po}$ and from the proof of the base case we have $(a, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In (3.b.ii.3.b), we then know there exist f, g, j such that $(c, f) \in G.\mathsf{po}^*$, $(f, g) \in G.(\mathsf{rf} \setminus \mathsf{po})$, $(g, b) \in G.(\mathsf{po} \cup \mathsf{rf})^*$ and that $j < i$, $(f, g) \in \mathsf{hb}_0$ and $(g, b) \in \mathsf{hb}_j$. We thus have $(a, f) \in G.\mathsf{po}$. There are again three cases to consider: i) $f \in G.\mathcal{NT}$, ii) $f \notin \mathcal{T}'_a \wedge f \in \mathcal{T}'_f$ for some $\mathcal{T}'_f \neq \mathcal{T}'_a$, or iii) $f \in \mathcal{T}'_a$.

In (3.b.ii.3.b.i), from the proof of the base case we have $(a, f) \in \Gamma.\mathsf{rpsi\text{-}hb}$. Similarly, from (I.H.) we have $(f, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As such, since $\Gamma.\mathsf{rpsi\text{-}hb}$ is transitively closed, we have $(a, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In (3.b.ii.3.b.ii), from the proof of the base case we have $(\{a\} \times \mathcal{T}_f) \in \Gamma.\mathsf{rpsi\text{-}hb}$. Similarly, from (I.H.) we have $\exists h \in \mathcal{T}_f.\ (h, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As such, since $\Gamma.\mathsf{rpsi\text{-}hb}$ is transitively closed, we have $(a, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In (3.b.ii.3.b.iii), since $a \overset{G.\mathsf{po}}{\to} c \overset{G.\mathsf{po}}{\to} f$, $a \in \mathcal{T}_a$, $c \notin \mathcal{T}_a$, and f is a write event $((f, g) \in G.(\mathsf{rf} \setminus \mathsf{po}))$, we know that f is an unlock event associated with some version lock $\mathtt{vx}$. As such we know that g is a transactional event in some $\mathcal{T}'_g$. Moreover, we know that there exist $p = \mathtt{W}(\mathtt{x}, -) \in \mathcal{T}_a$ and $q = \mathtt{R}(\mathtt{x}, -) \in \mathcal{T}_b$, such that $(p, q) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As such, we have $(\mathcal{T}_a \times \mathcal{T}_g) \subseteq \Gamma.\Big(\mathsf{st}; \big([\mathcal{W}]; \mathsf{st}; \mathsf{rf} \setminus \mathsf{st}; \mathsf{st}; [\mathcal{R}]\big)_{loc}; \mathsf{st}\Big) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. In particular we have $(\{a\} \times \mathcal{T}_g) \in \Gamma.\mathsf{rpsi\text{-}hb}$. On the other hand, since we have $(g, b) \in \mathsf{hb}_j$ and $j < i$, from (I.H.) we have $\exists o \in \mathcal{T}_g.\ (o, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As such, since $\Gamma.\mathsf{rpsi\text{-}hb}$ is transitively closed, we have $(a, b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

Case 4. $a \in \mathcal{T}'_a \wedge b \in \mathcal{T}'_a \wedge \mathcal{T}'_a \neq \mathcal{T}'_b$
There are two cases to consider: a) $c \in G.\mathcal{NT}$; or b) $c \in \mathcal{T}'_c$ for some $\mathcal{T}'_c$.

In case (4.a), from (I.H. we have $(\{c\} \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. On the other hand, from the proof of the base case we have $\exists d \in \mathcal{T}_a.\ (d, c) \in \Gamma.\mathsf{rpsi\text{-}hb} \wedge (a \in \mathcal{T}_a \Rightarrow d = a)$. As such, since $\Gamma.\mathsf{rpsi\text{-}hb}$ is transitively closed, we have $\exists d \in \mathcal{T}_a.\ (\{d\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb} \wedge (a \in \mathcal{T}_a \Rightarrow d = a)$, as required.

In case (4.b), there are three cases to consider: i) $\mathcal{T}'_c \neq \mathcal{T}'_a \wedge \mathcal{T}'_c \neq \mathcal{T}'_a$; or ii) $\mathcal{T}'_c = \mathcal{T}'_b$; or iii) $\mathcal{T}'_c = \mathcal{T}'_a$.

In case (4.b.i), since $(a, c) \in \mathsf{hb}_0$ and $\mathcal{T}'_c \neq \mathcal{T}'_a$, from the proof of the base case we know $\exists d \in \mathcal{T}_a.\ (\{d\} \times \mathcal{T}_c) \subseteq \Gamma.\mathsf{rpsi\text{-}hb} \wedge (a \in \mathcal{T}_a \Rightarrow d = a)$. Similarly, since $(c, b) \in \mathsf{hb}_n$ and $\mathcal{T}'_c \neq \mathcal{T}'_b$, from (I.H.) we know $\exists e \in \mathcal{T}_c.\ (\{e\} \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. As such, since $\Gamma.\mathsf{rpsi\text{-}hb}$ is transitively closed, we have $\exists d \in \mathcal{T}_a.\ (\{d\} \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{rpsi\text{-}hb} \wedge (a \in \mathcal{T}_a \Rightarrow d = a)$, as required.

In case (4.b.ii), since $(a, c) \in \mathsf{hb}_0$ and $\mathcal{T}'_c = \mathcal{T}'_b$ (and thus $\mathcal{T}_c = \mathcal{T}_b$) and $\mathcal{T}'_a \neq \mathcal{T}'_b$, from the proof of the base case we know $\exists d \in \mathcal{T}_a.\ (\{d\} \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{rpsi\text{-}hb} \wedge (a \in \mathcal{T}_a \Rightarrow d = a)$, as required.

In case (4.b.iii), Since $\mathcal{T}'_c = \mathcal{T}'_a$, $\mathcal{T}'_a \neq \mathcal{T}'_b$, and $(c, b) \in \mathsf{hb}_n \subseteq G.\mathsf{hb}$, from (I.H.) we have $(\mathcal{T}_a \times \mathcal{T}_b) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In case (4.b.iii), since $\mathcal{T}'_c = \mathcal{T}'_a$, $\mathcal{T}'_a \neq \mathcal{T}'_b$, and $(c, b) \in \mathsf{hb}_n \subseteq G.\mathsf{hb}$, from (I.H.) we have $\exists e \in \mathcal{T}_c.\ (\{e\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb} \wedge (c \in \mathcal{T}_a \Rightarrow e = c)$. There are now three cases to consider: 1) $a \notin \mathcal{T}_a$; 2) $a, c \in \mathcal{T}_a$; 3) $a \in \mathcal{T}_a$ and $c \notin \mathcal{T}_a$. In case (4.b.iii.1) we have $\exists e \in \mathcal{T}_a.\ (\{e\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required. In case (4.b.iii.2), we then have $(\{c\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$. On the other hand, since we have $(a, c) \in G.\mathsf{po}$ and $G.\mathsf{po}$ does not change the orderings between events of Γ, we also have $(a, c) \in \Gamma.\mathsf{po}$. As $\Gamma.\mathsf{po} \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$ and $\Gamma.\mathsf{rpsi\text{-}hb}$ is transitively closed, we have $(\{a\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In case (4.b.iii.3), we are required to show that $(\{a\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$. It is easy to demonstrate that $\mathsf{hb}_i = G.(\mathsf{po}^+ \cup \mathsf{po}^*; \mathsf{rf} \setminus \mathsf{po}; (\mathsf{po} \cup \mathsf{rf})^*)$. That is, as we have $(c, b) \in \mathsf{hb}_i$, we either have a) $(c, b) \in G.\mathsf{po}^+$; or b) $(c, b) \in G.(\mathsf{po}^*; \mathsf{rf} \setminus \mathsf{po}; (\mathsf{po} \cup \mathsf{rf})^*)$. In (4.b.iii.3.a), we then have $(a, b) \in G.\mathsf{po}$ and from the proof of the base case we have $(\{a\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In (4.b.iii.3.b), we then know there exist f, g, j such that $(c, f) \in G.\mathsf{po}^*$, $(f, g) \in G.(\mathsf{rf} \setminus \mathsf{po})$, $(g, b) \in G.(\mathsf{po} \cup \mathsf{rf})^*$ and that $j < i$, $(f, g) \in \mathsf{hb}_0$ and $(g, b) \in \mathsf{hb}_j$. We thus have $(a, f) \in G.\mathsf{po}$. There are again three cases to consider: i) $f \in G.\mathcal{NT}$; ii) $f \in \mathcal{T}'_b$; iii) $f \in \mathcal{T}'_f \wedge f \notin \mathcal{T}'_a \wedge f \notin \mathcal{T}'_b$ for some $\mathcal{T}'_f \neq \mathcal{T}'_a$, or iv) $f \in \mathcal{T}'_a$.

In (4.b.iii.3.b.i), from the proof of the base case we have $(a, f) \in \Gamma.\mathsf{rpsi\text{-}hb}$. Similarly, from (I.H.) we have $(\{f\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As such, since $\Gamma.\mathsf{rpsi\text{-}hb}$ is transitively closed, we have $(\{a\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In (4.b.iii.3.b.ii), from the proof of the base case we have $(\{a\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In (4.b.iii.3.b.iii), from the proof of the base case we have $(\{a\} \times \mathcal{T}_f) \in \Gamma.\mathsf{rpsi\text{-}hb}$. Similarly, from (I.H.) we have $\exists h \in \mathcal{T}_f.\ (\{h\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As such, since $\Gamma.\mathsf{rpsi\text{-}hb}$ is transitively closed, we have $(\{a\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

In (4.b.iii.3.b.iv), since $a \overset{G.\mathsf{po}}{\to} c \overset{G.\mathsf{po}}{\to} f$, $a \in \mathcal{T}_a$, $c \notin \mathcal{T}_a$, and f is a write event ($(f, g) \in G.(\mathsf{rf} \setminus \mathsf{po})$), we know that f is an unlock event associated with some version lock $\mathtt{vx}$. As such we know that g is a transactional event in some $\mathcal{T}'_g$. Moreover, we know that there exist $p = \mathtt{W}(\mathtt{x}, -) \in \mathcal{T}_a$ and $q = \mathtt{R}(\mathtt{x}, -) \in \mathcal{T}_b$, such that $(p, q) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As such, we have $(\mathcal{T}_a \times \mathcal{T}_g) \subseteq \Gamma.\Big(\mathsf{st}; \big([\mathcal{W}]; \mathsf{st}; \mathsf{rf} \setminus \mathsf{st}; \mathsf{st}; [\mathcal{R}]\big)_{loc}; \mathsf{st}\Big) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. In particular we have $(\{a\} \times \mathcal{T}_g) \in \Gamma.\mathsf{rpsi\text{-}hb}$. Now either $\mathcal{T}'_g = \mathcal{T}'_b$ (and thus $\mathcal{T}_g = \mathcal{T}_b$) and we have $(\{a\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required. Or, $\mathcal{T}'_g \neq \mathcal{T}'_b$ and since $(g, b) \in \mathsf{hb}_j$ and $j < i$, from (I.H.) we have $\exists o \in \mathcal{T}_g.\ (\{o\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As such, since $\Gamma.\mathsf{rpsi\text{-}hb}$ is transitively closed, we have $(\{a\} \times \mathcal{T}_b) \in \Gamma.\mathsf{rpsi\text{-}hb}$, as required.

Case 5. $a, b \in \mathcal{T}'_a$
There are three cases to consider: a) $c \in G.\mathcal{NT}$; or b) $c \in \mathcal{T}'_c$ for some $\mathcal{T}'_c$ and $\mathcal{T}'_c \neq \mathcal{T}'_a$; or c) $c \in \mathcal{T}'_a$.

In case (5.a), from the proof of base case we know $\exists d \in \mathcal{T}_a.\ (d, c) \in \Gamma.\mathsf{rpsi\text{-}hb}$. On the other hand, from I.H. we know $(\{c\} \times \mathcal{T}_a) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. In particular, since $d \in \mathcal{T}_a$, we have $(c, d) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As such, we have $\exists d \in \mathcal{T}_a.\ d \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\to} c \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\to} d$, contradicting the assumption that Γ is consistent.

In case (5.b), from the proof of base case we know $\exists d \in \mathcal{T}_a.\ (\{d\} \times \mathcal{T}_c) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. On the other hand, from I.H. we know $\exists e \in \mathcal{T}_c.\ (\{e\} \times \mathcal{T}_a) \subseteq \Gamma.\mathsf{rpsi\text{-}hb}$. In particular, since $d \in \mathcal{T}_a$ and $e \in \mathcal{T}_c$, we have $(d, e) \in \Gamma.\mathsf{rpsi\text{-}hb}$ and $(e, d) \in \Gamma.\mathsf{rpsi\text{-}hb}$. As such, we have $\exists d \in \mathcal{T}_a, e \in \mathcal{T}_c.\ d \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\to} e \stackrel{\Gamma.\mathsf{rpsi\text{-}hb}}{\to} d$, contradicting the assumption that Γ is consistent.

In case (5.c), from the proof of base case we know $(a, c) \in G.\mathsf{po}$. Similarly, from (I.H.) we know $(c, b) \in G.\mathsf{po}$. As $G.\mathsf{po}$ is transitively closed, we have $(a, b) \in G.\mathsf{po}$, as required. $\square$