

RustBelt Relaxed - Technical Appendix

This work is accompanied by a machine-checked formalization in Coq, which includes all definitions, theorems, lemmas and proofs in this appendix, with the exception of the correspondence proof (§2).

Contents

1	Language	2
1.1	Grammar	2
1.2	Operational Semantics	2
2	Correspondence of ORC11 to RC11	13
2.1	Executions	13
2.1.1	Consistent Executions	14
2.2	Declarative Semantics	14
2.3	Operational Graph Semantics (OGS)	16
2.4	OGS to ORC11	18
3	Lifetime Logic for Views	23
3.1	Proof Rules	23
3.2	Derived Forms of Borrowing	26
4	GPFSL	29
5	Case Study: Arc	39
5.1	The Full APIs of Arc	39
5.2	Insufficient Synchronization in <code>get_mut</code>	42

1 Language

1.1 Grammar

Our language is an extension of the original RustBelt’s λ_{Rust} with the relaxed memory semantics of ORC11 (§1.2). λ_{Rust} is a lambda calculus with integers, locations with explicit allocation and deallocation, and a notion of poison value ⊥ . Instead of **sc** for atomic accesses, we use release **rel**, acquire **acq**, and relaxed **rlx** accesses together with fences.

The grammar is given in Fig. 1. Several syntactic sugars are taken as-is from the original RustBelt, given in Fig. 2. We refer the reader to the original RustBelt appendix (Jung et al. [2017]) for more explanation of the grammar and syntactic sugars.

1.2 Operational Semantics

Following iGPS (Kaiser et al. [2017]) we use an operational semantics for relaxed memory so that it can be instantiated in Iris. For this work, we extend iGPS’s operational semantics for RA+NA to include relaxed accesses and fences.

The semantics, called ORC11, is defined by three sub semantics: the expressions semantics (Fig. 5), the machine semantics (Fig. 9), and the race-detecting semantics (Fig. 6 and Fig. 7). The combined thread pool semantics is given in Fig. 10 and Fig. 11. In §2, we sketch a proof of correspondence that relates ORC11 to the axiomatic semantics from Lahav et al. [2017].

$$\begin{aligned}
& z \in \mathbb{Z} \\
\text{Expr} \ni e & ::= v \mid x \\
& \mid e.e \mid e + e \mid e - e \mid e \leq e \mid e == e \\
& \mid e(\bar{e}) \\
& \mid *^o e \\
& \mid e_1 :=_o e_2 \\
& \mid \mathbf{CAS}(e_0, e_1, e_2, o_f, o_r, o_w) \\
& \mid \mathbf{alloc}(e) \\
& \mid \mathbf{free}(e_1, e_2) \\
& \mid \mathbf{case } e \text{ of } \bar{e} \\
& \mid \mathbf{fork} \{ e \} \\
& \mid \mathbf{fence}_o \\
\text{Val} \ni v & ::= \ast \mid \ell \mid z \mid \mathbf{rec } f(\bar{x}) := e \\
\text{Loc} \ni \ell & ::= (i, n) \\
\text{Order} \ni o & ::= \mathbf{acq} \mid \mathbf{rel} \mid \mathbf{rlx} \mid \mathbf{na} \\
\text{Ctx} \ni K & ::= \bullet \\
& \mid K.e \mid v.K \mid K + e \mid v + K \mid K - e \mid v - K \\
& \mid K \leq e \mid v \leq K \mid K == e \mid v == K \\
& \mid K(\bar{e}) \mid v(\bar{v} \uparrow [K] \uparrow \bar{e}) \\
& \mid *^o K \mid K :=_o e \mid v :=_o K \\
& \mid \mathbf{CAS}(K, e_1, e_2, o_f, o_r, o_w) \\
& \mid \mathbf{CAS}(v_0, K, e_2, o_f, o_r, o_w) \\
& \mid \mathbf{CAS}(v_0, v_1, K, o_f, o_r, o_w) \\
& \mid \mathbf{alloc}(K) \\
& \mid \mathbf{free}(K, e_2) \\
& \mid \mathbf{free}(e_1, K) \\
& \mid \mathbf{case } K \text{ of } \bar{e}
\end{aligned}$$

Figure 1: Language syntax.

```

funrec  $f(\bar{x})$  ret  $k := e := \mathbf{rec}$   $f([k] \uparrow \bar{x}) := e$ 
    let  $x = e$  in  $e' := (\mathbf{rec\_}([x]) := e')(e)$ 
     $e'; e := \mathbf{let\_} = e' \mathbf{in}$   $e$ 
letcont  $k(\bar{x}) := e$  in  $e' := \mathbf{let}$   $k = (\mathbf{rec}$   $k(\bar{x}) := e) \mathbf{in}$   $e'$ 
    jump  $k(\bar{e}) := k(\bar{e})$ 
    call  $f(\bar{e})$  ret  $k := f([k] \uparrow \bar{e})$ 

    false := 0
    true := 1
    if  $e_0$  then  $e_1$  else  $e_2 := \mathbf{case}$   $e_0$  of  $[e_1, e_2]$ 

     $*e := *_{\mathbf{na}} e$ 
 $e_1 := e_2 := e_1 :=_{\mathbf{na}} e_2$ 
    new := rec  $\mathbf{new}(size) :=$ 
        if  $size == 0$  then (42, 1337) else alloc( $size$ )
    delete := rec  $\mathbf{delete}(size, ptr) :=$ 
        if  $size == 0$  then  $\emptyset$  else free( $size, ptr$ )
    memcpy := rec  $\mathbf{memcpy}(dst, len, src) :=$ 
        if  $len \leq 0$  then  $\emptyset$  else
             $dst.0 := src.0;$ 
             $\mathbf{memcpy}(dst.1, len - 1, src.1)$ 
 $e_1 :=_n *e_2 := \mathbf{memcpy}(e_1, n, e_2)$ 
 $e :=_{\mathbf{inj}^i} () := e.0 := i$ 
 $e_1 :=_{\mathbf{inj}^i} e_2 := e_1.0 := i; e_1.1 := e_2$ 
 $e_1 :=_{\mathbf{inj}^i} *e_2 := e_1.0 := i; e_1.1 :=_n *e_2$ 

    skip := let  $x = \del{\emptyset} in  $\emptyset$ 
    newlft :=  $\emptyset$ 
    endlft := skip$ 
```

Figure 2: Syntactic sugars.

$$\begin{aligned}
\pi \in \textit{Thread} &::= \mathbb{N} \\
t \in \textit{Time} &::= \mathbb{N}^+ \\
\omega \in \textit{MsgVal} &::= \dagger \mid \spadesuit \mid v \in \textit{Val} \\
\textit{ActionIds} &::= 2^{\mathbb{N}^+} \\
V \in \textit{View} &::= \textit{Loc} \xrightarrow{\text{fin}} \{w : \textit{Time}, \text{aw} : \textit{ActionIds}, \text{nr} : \textit{ActionIds}, \text{ar} : \textit{ActionIds}\} \\
\mathcal{V} \in \textit{ThreadView} &::= \left\{ \text{rel} : \textit{Loc} \xrightarrow{\text{fin}} \textit{View}, \text{frel} : \textit{View}, \text{cur} : \textit{View}, \text{acq} : \textit{View} \right\} \\
m \in \textit{ExtMsg} &::= \left\{ \text{ts} : \textit{Time}, \text{val} : \textit{MsgVal}, \text{view} : \textit{View}^? \right\} \\
\mathcal{M} \in \textit{MsgPool} &::= \textit{Loc} \xrightarrow{\text{fin}} \textit{Time} \xrightarrow{\text{fin}} \left\{ \text{val} : \textit{MsgVal}, \text{view} : \textit{View}^? \right\} \\
\mathcal{N} \in \textit{NARace} &::= \textit{View} \\
\varsigma \in \textit{GlobalState} &::= \textit{MsgPool} \times \textit{NARace} \\
\textit{MemEvent} \ni \varepsilon &::= \langle \text{Alloc}, \ell, n \in \mathbb{N}^+ \rangle \\
& \quad \mid \langle \text{Dealloc}, \ell, n \in \mathbb{N}^+ \rangle \\
& \quad \mid \langle \text{Read}, \ell, v, o \rangle \\
& \quad \mid \langle \text{Write}, \ell, v, o \rangle \\
& \quad \mid \langle \text{Update}, \ell, v_r, v_w, o_r, o_w \rangle \\
& \quad \mid \langle \text{Fence}, o \rangle
\end{aligned}$$

Figure 3: Machine state definitions.

$$\omega \in \text{Readable}(\ell, \mathcal{M}, \mathcal{V}) := \exists t. \mathcal{M}(\ell)(t) = (\omega, _) \wedge t \leq \mathcal{V}.\text{cur}(\ell)$$

$$\boxed{\omega \equiv v}$$

$$v \equiv v$$

$$\dagger \equiv \text{⊗}$$

$$\boxed{\ell \in \text{unalloc}(\mathcal{M})}$$

$$\frac{\ell \notin \text{dom}(\mathcal{M})}{\ell \in \text{unalloc}(\mathcal{M})}$$

$$\frac{\exists t. \mathcal{M}(\ell)(t) = (\dagger, _)}{\ell \in \text{unalloc}(\mathcal{M})}$$

$$\boxed{\mathcal{M} \vdash v_1 = v_2}$$

$$\mathcal{M} \vdash z = z$$

$$\mathcal{M} \vdash \ell = \ell$$

$$\frac{\ell_1 \in \text{unalloc}(\mathcal{M}) \vee \ell_2 \in \text{unalloc}(\mathcal{M})}{\mathcal{M} \vdash \ell_1 = \ell_2}$$

$$\boxed{\vdash v_1 \neq v_2}$$

$$\frac{z_1 \neq z_2}{\vdash z_1 \neq z_2}$$

$$\frac{\ell_1 \neq \ell_2}{\vdash \ell_1 \neq \ell_2}$$

$$\vdash \ell \neq 0$$

$$\vdash 0 \neq \ell$$

$$\boxed{\vdash v_1 =^? v_2}$$

$$\vdash z_1 =^? z_2$$

$$\vdash \ell_1 =^? \ell_2$$

$$\vdash \ell =^? 0$$

$$\vdash 0 =^? \ell$$

$$\boxed{o_1 \sqsubseteq o_2}$$

$$\mathbf{na} \sqsubseteq \mathbf{rlx}$$

$$\mathbf{na} \sqsubseteq \mathbf{acq}$$

$$\mathbf{na} \sqsubseteq \mathbf{rel}$$

$$\mathbf{rlx} \sqsubseteq \mathbf{acq}$$

$$\mathbf{rlx} \sqsubseteq \mathbf{rel}$$

Figure 4: Auxilliary relations.

$$\boxed{\mathcal{M}, \mathcal{V} \vdash e \xrightarrow{\varepsilon^?} e'_1, e'_2}$$

$$\begin{array}{c}
\text{OE-ECTX} \\
\frac{e \rightarrow e'_1, e'_2}{\mathcal{M}, \mathcal{V} \vdash K[e] \rightarrow K[e'_1], e'_2}
\end{array}
\quad
\begin{array}{c}
\text{OE-PROJ} \\
\mathcal{M}, \mathcal{V} \vdash \ell.n \rightarrow \ell + n
\end{array}
\quad
\begin{array}{c}
\text{OE-ADD} \\
\frac{z_1 + z_2 = z'}{\mathcal{M}, \mathcal{V} \vdash z_1 + z_2 \rightarrow z'}
\end{array}$$

$$\begin{array}{c}
\text{OE-SUB} \\
\frac{z_1 - z_2 = z'}{\mathcal{M}, \mathcal{V} \vdash z_1 - z_2 \rightarrow z'}
\end{array}
\quad
\begin{array}{c}
\text{OE-LE-TRUE} \\
\frac{z_1 \leq z_2}{\mathcal{M}, \mathcal{V} \vdash z_1 \leq z_2 \rightarrow 1}
\end{array}
\quad
\begin{array}{c}
\text{OE-LE-FALSE} \\
\frac{z_1 > z_2}{\mathcal{M}, \mathcal{V} \vdash z_1 \leq z_2 \rightarrow 0}
\end{array}$$

$$\begin{array}{c}
\text{OE-EQ-TRUE} \\
\frac{\mathcal{M} \vdash v_1 = v_2}{\mathcal{M}, \mathcal{V} \vdash v_1 == v_2 \rightarrow 1}
\end{array}
\quad
\begin{array}{c}
\text{OE-EQ-FALSE} \\
\frac{\vdash v_1 \neq v_2}{\mathcal{M}, \mathcal{V} \vdash v_1 == v_2 \rightarrow 0}
\end{array}$$

$$\begin{array}{c}
\text{OE-ALLOC} \\
\frac{n > 0}{\mathcal{M}, \mathcal{V} \vdash \mathbf{alloc}(n) \xrightarrow{\langle \text{Alloc}, \ell, n \rangle} \ell}
\end{array}
\quad
\begin{array}{c}
\text{OE-FREE} \\
\frac{n > 0}{\mathcal{M}, \mathcal{V} \vdash \mathbf{free}(n, \ell) \xrightarrow{\langle \text{Dealloc}, \ell, n \rangle} \clubsuit}
\end{array}$$

$$\begin{array}{c}
\text{OE-DEREF} \\
\frac{\omega \equiv v}{\mathcal{M}, \mathcal{V} \vdash *o \ell \xrightarrow{\langle \text{Read}, \ell, \omega, o \rangle} v}
\end{array}
\quad
\begin{array}{c}
\text{OE-ASSIGN} \\
\mathcal{M}, \mathcal{V} \vdash \ell :=_o v \xrightarrow{\langle \text{Write}, \ell, v, o \rangle} \clubsuit
\end{array}$$

$$\begin{array}{c}
\text{OE-CAS-FAIL} \\
\mathbf{rlx} \sqsubseteq o_f \\
\mathbf{rlx} \sqsubseteq o_r \quad (\forall \omega \in \text{Readable}(\ell, \mathcal{M}, \mathcal{V}). \exists v'. \omega \equiv v' \wedge \vdash v_1 =? v') \quad \vdash v_1 \neq v_r \\
\mathbf{rlx} \sqsubseteq o_w \\
\hline
\mathcal{M}, \mathcal{V} \vdash \mathbf{CAS}(\ell, v_1, v_2, o_f, o_r, o_w) \xrightarrow{\langle \text{Read}, \ell, v_r, o_f \rangle} 0
\end{array}$$

$$\begin{array}{c}
\text{OE-CAS-SUC} \\
\mathbf{rlx} \sqsubseteq o_f \\
\mathbf{rlx} \sqsubseteq o_r \quad (\forall \omega \in \text{Readable}(\ell, \mathcal{M}, \mathcal{V}). \exists v'. \omega \equiv v' \wedge \vdash v_1 =? v') \quad \mathcal{M} \vdash v_1 = v_r \\
\mathbf{rlx} \sqsubseteq o_w \\
\hline
\mathcal{M}, \mathcal{V} \vdash \mathbf{CAS}(\ell, v_1, v_2, o_f, o_r, o_w) \xrightarrow{\langle \text{Update}, \ell, v_r, v_2, o_r, o_w \rangle} 1
\end{array}$$

$$\begin{array}{c}
\text{OE-FENCE} \\
\mathcal{M}, \mathcal{V} \vdash \mathbf{fence}_o \xrightarrow{\langle \text{Fence}, o \rangle} \clubsuit
\end{array}
\quad
\begin{array}{c}
\text{OE-CASE} \\
\mathcal{M}, \mathcal{V} \vdash \mathbf{case } i \text{ of } (\bar{e}) \rightarrow \bar{e}_i
\end{array}$$

$$\begin{array}{c}
\text{OE-APP} \\
\mathcal{M}, \mathcal{V} \vdash (\mathbf{rec } f(\bar{x}) := e)(\bar{v}) \rightarrow e[\mathbf{rec } f(\bar{x}) := e/f, \bar{v}/\bar{x}]
\end{array}
\quad
\begin{array}{c}
\text{OE-FORK} \\
\mathcal{M}, \mathcal{V} \vdash \mathbf{fork} \{ e \} \rightarrow \clubsuit, e
\end{array}$$

Figure 5: Expression semantics.

$$\boxed{\mathcal{M}, \mathcal{N}, V \vdash \text{RaceFree}(\varepsilon)}$$

$$\frac{\text{DRF-READ-NA} \quad \forall t \in \text{dom}(\mathcal{M}(\ell)). t \leq \text{cur}(\ell).\text{w} \quad \mathcal{N}(\ell).\text{aw} \sqsubseteq \text{cur}(\ell).\text{aw}}{\mathcal{M}, \mathcal{N}, (\text{rel}, \text{frel}, \text{cur}, \text{acq}) \vdash \text{RaceFree}(\langle \text{Read}, \ell, v, \mathbf{na} \rangle)}$$

$$\frac{\text{DRF-WRITE-NA} \quad \mathcal{N}(\ell).\text{aw} \sqsubseteq \text{cur}(\ell).\text{aw} \quad \mathcal{N}(\ell).\text{nr} \sqsubseteq \text{cur}(\ell).\text{nr} \quad \mathcal{N}(\ell).\text{ar} \sqsubseteq \text{cur}(\ell).\text{ar} \quad \forall t \in \text{dom}(\mathcal{M}(\ell)). t \leq \text{cur}(\ell).\text{w} < t_w}{\mathcal{M}, \mathcal{N}, (\text{rel}, \text{frel}, \text{cur}, \text{acq}) \vdash \text{RaceFree}(\langle \text{Write}, \ell, v, \mathbf{na} \rangle)}$$

$$\frac{\text{DRF-READ-AT} \quad \mathbf{rlx} \sqsubseteq o \quad \mathcal{N}(\ell).\text{w} \leq \text{cur}(\ell).\text{w}}{\mathcal{M}, \mathcal{N}, (\text{rel}, \text{frel}, \text{cur}, \text{acq}) \vdash \text{RaceFree}(\langle \text{Read}, \ell, v, o \rangle)}$$

$$\frac{\text{DRF-WRITE-AT} \quad \mathbf{rlx} \sqsubseteq o \quad \mathcal{N}(\ell).\text{w} \leq \text{cur}(\ell).\text{w} \quad \mathcal{N}(\ell).\text{nr} \sqsubseteq \text{cur}(\ell).\text{nr}}{\mathcal{M}, \mathcal{N}, (\text{rel}, \text{frel}, \text{cur}, \text{acq}) \vdash \text{RaceFree}(\langle \text{Write}, \ell, v, o \rangle)}$$

$$\frac{\text{DRF-UPDATE} \quad \mathcal{M}, \mathcal{N}, V \vdash \text{RaceFree}(\langle \text{Read}, \ell, v_r, o_r \rangle) \quad \mathcal{M}, \mathcal{N}, V \vdash \text{RaceFree}(\langle \text{Write}, \ell, v_w, o_w \rangle)}{\mathcal{M}, \mathcal{N}, V \vdash \text{RaceFree}(\langle \text{Update}, \ell, v_r, v_w, o_r, o_w \rangle)}$$

$$\frac{\text{DRF-ALLOC}}{\mathcal{M}, \mathcal{N}, V \vdash \text{RaceFree}(\langle \text{Alloc}, \ell, n \rangle)}$$

$$\frac{\text{DRF-DEALLOC} \quad \forall i \in [< n], t' \in \text{dom}(\mathcal{M}(\ell + i)). t' \leq \text{cur}(\ell).\text{w} \quad \forall i \in [< n]. \mathcal{N}(\ell + i).\text{aw} \sqsubseteq \text{cur}(\ell + i).\text{aw} \quad \forall i \in [< n]. \mathcal{N}(\ell + i).\text{nr} \sqsubseteq \text{cur}(\ell + i).\text{nr} \quad \forall i \in [< n]. \mathcal{N}(\ell + i).\text{ar} \sqsubseteq \text{cur}(\ell + i).\text{ar}}{\mathcal{M}, \mathcal{N}, (\text{rel}, \text{frel}, \text{cur}, \text{acq}) \vdash \text{RaceFree}(\langle \text{Dealloc}, \ell, n \rangle)}$$

Figure 6: Data-race-free (DRF) pre condition, detailing the exact requirements on the local and global race detector state for any particular memory event.

$$\boxed{\mathcal{N} \xrightarrow{\varepsilon, t, m^*} \mathcal{N}'}$$

$$\frac{\text{DRF-POST-READ-NA} \quad r \notin \mathcal{N}(\ell).\text{nr} \quad \mathcal{N}' = \mathcal{N}[\ell \leftarrow \{\mathcal{N}(\ell) \text{ with nr} := \mathcal{N}(\ell).\text{nr} \cup \{r\}\}]}{\mathcal{N} \xrightarrow{\langle \text{Read}, \ell, v, \text{na} \rangle, r, []} \mathcal{N}'}$$

$$\frac{\text{DRF-POST-WRITE-NA} \quad \mathcal{N}' = \mathcal{N}[\ell \leftarrow \{\mathcal{N}(\ell) \text{ with w} := m.\text{ts}\}]}{\mathcal{N} \xrightarrow{\langle \text{Write}, \ell, v, \text{na} \rangle, \perp, [m]} \mathcal{N}'}$$

$$\frac{\text{DRF-POST-READ-AT} \quad \text{rlx} \sqsubseteq o \quad r \notin \mathcal{N}(\ell).\text{ar} \quad \mathcal{N}' = \mathcal{N}[\ell \leftarrow \{\mathcal{N}(\ell) \text{ with ar} := \mathcal{N}(\ell).\text{ar} \cup \{r\}\}]}{\mathcal{N} \xrightarrow{\langle \text{Read}, \ell, v, o \rangle, r, []} \mathcal{N}'}$$

$$\frac{\text{DRF-POST-WRITE-AT} \quad \text{rlx} \sqsubseteq o \quad \mathcal{N}' = \mathcal{N}[\ell \leftarrow \{\mathcal{N}(\ell) \text{ with aw} := \mathcal{N}(\ell).\text{aw} \cup \{m.\text{ts}\}\}]}{\mathcal{N} \xrightarrow{\langle \text{Write}, \ell, v, o \rangle, \perp, [m]} \mathcal{N}'}$$

$$\frac{\text{DRF-POST-UPDATE} \quad r \notin \mathcal{N}(\ell).\text{ar} \quad \mathcal{N}' = \mathcal{N}[\ell \leftarrow \{\mathcal{N}(\ell) \text{ with ar} := \mathcal{N}(\ell).\text{ar} \cup \{r\}\}] \quad \mathcal{N}' = \mathcal{N}[\ell \leftarrow \{\mathcal{N}(\ell) \text{ with aw} := \mathcal{N}(\ell).\text{aw} \cup \{m.\text{ts}\}\}]}{\mathcal{N} \xrightarrow{\langle \text{Update}, \ell, v_r, v_w, o_r, o_w \rangle, r, [m]} \mathcal{N}'}$$

$$\frac{\text{DRF-POST-ALLOC} \quad \mathcal{N}' = \mathcal{N}[\ell + i \leftarrow \{w := m_i.\text{ts}, \text{aw} := \emptyset, \text{nr} := \emptyset, \text{ar} := \emptyset\} \mid i \in [< n]]}{\mathcal{N} \xrightarrow{\langle \text{Alloc}, \ell, n \rangle, \perp, [m_0 \dots m_{n-1}]} \mathcal{N}'}$$

$$\frac{\text{DRF-POST-DEALLOC} \quad \mathcal{N}' = \mathcal{N}[\ell + i \leftarrow \{\mathcal{N}(\ell + i) \text{ with w} := m_i.\text{ts}\} \mid i \in [< n]]}{\mathcal{N} \xrightarrow{\langle \text{Dealloc}, \ell, n \rangle, \perp, [m_0 \dots m_{n-1}]} \mathcal{N}'}$$

Figure 7: Data-race-free (DRF) post condition, detailing the change to the global race detector state on a per-event basis.

OM-READ-HELPER

$$\begin{array}{c}
\text{cur}(\ell).w \leq t \quad R(\ell) \leq t \\
V = [\ell \leftarrow \{\mathbf{w} := t, \mathbf{aw} := \emptyset, \mathbf{nr} := \text{if } o = \mathbf{na} \text{ then } \{r\} \text{ else } \emptyset, \mathbf{ar} := \text{if } o \sqsubseteq \mathbf{rlx} \text{ then } \{r\} \text{ else } \emptyset\}] \\
\text{cur}' = \text{if } \mathbf{acq} \sqsubseteq o \text{ then } \text{cur} \sqcup V \sqcup R \text{ else } \text{cur} \sqcup V \\
\text{acq}' = \text{if } \mathbf{rlx} \sqsubseteq o \text{ then } \text{acq} \sqcup V \sqcup R \text{ else } \text{acq} \sqcup V \\
\hline
(\text{rel}, \text{frel}, \text{cur}, \text{acq}) \xrightarrow{\langle \mathbf{R}:o,\ell,t,R \rangle, r} (\text{rel}, \text{frel}, \text{cur}', \text{acq}')
\end{array}$$

OM-WRITE-HELPER

$$\begin{array}{c}
\text{cur}(\ell).w < t \\
V = [\ell \leftarrow \{\mathbf{w} := t, \mathbf{aw} := \text{if } \mathbf{rlx} \sqsubseteq o \text{ then } \{t\} \text{ else } \emptyset, \mathbf{nr} := \emptyset, \mathbf{ar} := \emptyset\}] \\
\text{cur}' = \text{cur} \sqcup V \quad \text{acq}' = \text{acq} \sqcup V \\
V' = \text{rel}(\ell) \sqcup \text{if } \mathbf{rel} \sqsubseteq o \text{ then } \text{cur}' \text{ else } V \quad \text{rel}' = \text{rel}[\ell \leftarrow V'] \\
R_w = \text{if } \mathbf{rlx} \sqsubseteq o \text{ then } V' \sqcup \text{frel} \sqcup R_r \text{ else } \perp \\
\hline
(\text{rel}, \text{frel}, \text{cur}, \text{acq}) \xrightarrow{\langle \mathbf{W}:o,\ell,t,R_r,R_w \rangle, \perp} (\text{rel}', \text{frel}, \text{cur}', \text{acq}')
\end{array}$$

Figure 8: View-helper relations.

$$\boxed{\varsigma \mid \mathcal{V} \xrightarrow{\varepsilon} \varsigma' \mid \mathcal{V}'}$$

OM-ALLOC

$$\frac{\begin{array}{l} \ell = (i, n') \quad \{i\} \times \mathbb{N} \# \text{dom}(\mathcal{M}) \\ \mathcal{M}' = \mathcal{M}[\ell + m \leftarrow [t_m \leftarrow (\dagger, \perp)] \mid m \in [< n]] \\ \mathcal{V} \xrightarrow{\langle \text{W:na}, \ell + 0, t_0, \perp, \perp \rangle} \dots \xrightarrow{\langle \text{W:na}, \ell + m, t_m, \perp, \perp \rangle} \dots \xrightarrow{\langle \text{W:na}, \ell + (n-1), t_{(n-1)}, \perp, \perp \rangle} \mathcal{V}' \\ ms = [(t_m, \dagger, \perp) \mid m \in [< n]] \end{array}}{\mathcal{M} \mid \mathcal{V} \xrightarrow{\langle \text{Alloc}, \ell, n \rangle, \perp, ms} \mathcal{M}' \mid \mathcal{V}'}$$

OM-FREE

$$\frac{\begin{array}{l} \ell = (i, n') \quad \text{dom}(\mathcal{M}) \cap \{i\} \times \mathbb{N} = \{i\} \times ([\geq n', < n' + n]) \\ \forall m \in [< n], t \in \text{dom}(\mathcal{M}(\ell + m)). t \leq \mathcal{V}.\text{cur}(\ell + m).\text{w} < t_m \wedge \mathcal{M}(\ell + m)(t).\text{val} \neq \Phi \\ \forall m \in [< n]. \text{dom}(\mathcal{M}(\ell + m)) \neq \emptyset \\ \mathcal{M}' = \mathcal{M}[\ell + m \leftarrow [t_m \leftarrow (\Phi, \perp)] \mid m \in [< n]] \\ \mathcal{V} \xrightarrow{\langle \text{W:na}, \ell + 0, t_0, \perp, \perp \rangle} \dots \xrightarrow{\langle \text{W:na}, \ell + m, t_m, \perp, \perp \rangle} \dots \xrightarrow{\langle \text{W:na}, \ell + (n-1), t_{(n-1)}, \perp, \perp \rangle} \mathcal{V}' \\ ms = [(t_m, \dagger, \perp) \mid m \in [< n]] \end{array}}{\mathcal{M} \mid \mathcal{V} \xrightarrow{\langle \text{Dealloc}, \ell, n \rangle, \perp, ms} \mathcal{M}' \mid \mathcal{V}'}$$

OM-READ

$$\frac{\begin{array}{l} \ell \notin \text{unalloc}(\mathcal{M}) \quad \mathcal{M}(\ell)(t) = (v, R) \\ \mathcal{V}'' \xrightarrow{\langle \text{R:o}, \ell, t, R \rangle, r} \mathcal{V}' \end{array}}{\mathcal{M} \mid \mathcal{V} \xrightarrow{\langle \text{Read}, \ell, v, o \rangle, r, []} \mathcal{M} \mid \mathcal{V}'}$$

OM-WRITE

$$\frac{\begin{array}{l} \ell \notin \text{unalloc}(\mathcal{M}) \quad t \notin \mathcal{M}(\ell) \\ \mathcal{M}' = \mathcal{M}[\ell \leftarrow \mathcal{M}(\ell)[t \leftarrow (v, R)]] \\ \mathcal{V} \xrightarrow{\langle \text{W:o}, \ell, t, \perp, R \rangle} \mathcal{V}' \end{array}}{\mathcal{M} \mid \mathcal{V} \xrightarrow{\langle \text{Write}, \ell, v, o \rangle, [(t, v, R)]} \mathcal{M}' \mid \mathcal{V}'}$$

OM-UPDATE

$$\frac{\begin{array}{l} \ell \notin \text{unalloc}(\mathcal{M}) \quad \mathcal{M}(\ell)(t_r) = (v_r, R_r) \quad t_w = t_r + 1 \quad t_w \notin \mathcal{M}(\ell) \\ \mathcal{M}' = \mathcal{M}[\ell \leftarrow \mathcal{M}(\ell)[t_w \leftarrow (v_w, R_w)]] \\ \mathcal{V} \xrightarrow{\langle \text{R:o}_r, \ell, t_r, R_r \rangle, r} \xrightarrow{\langle \text{W:o}_w, \ell, t_w, R_r, R_w \rangle} \mathcal{V}' \end{array}}{\mathcal{M} \mid \mathcal{V} \xrightarrow{\langle \text{Update}, \ell, v_r, v_w, o_r, o_w \rangle, r, [(t_w, v_w, R_w)]} \mathcal{M}' \mid \mathcal{V}'}$$

OM-ACQ-FENCE

$$\mathcal{M} \mid \mathcal{V} \xrightarrow{\langle \text{Fence}, \text{acq} \rangle} \mathcal{M} \mid (\mathcal{V}.\text{rel}, \mathcal{V}.\text{frel}, \mathcal{V}.\text{acq}, \mathcal{V}.\text{acq})$$

OM-REL-FENCE

$$\mathcal{M} \mid \mathcal{V} \xrightarrow{\langle \text{Fence}, \text{rel} \rangle} \mathcal{M} \mid ([\ell \leftarrow \mathcal{V}.\text{cur} \mid \ell \in \text{dom}(\mathcal{V}.\text{rel})], \mathcal{V}.\text{cur}, \mathcal{V}.\text{cur}, \mathcal{V}.\text{acq})$$

Figure 9: Machine semantics.

$$\boxed{\varsigma \mid \mathcal{V} \mid e \xrightarrow{\varepsilon} \pi \varsigma' \mid \mathcal{V}' \mid e'}$$

$$\text{COMBRED-PURE} \quad \frac{\mathcal{M}, \mathcal{V} \vdash e \rightarrow e', es}{(\mathcal{M}, \mathcal{N}, V) \mid e \xrightarrow{\perp, es} (\mathcal{M}, \mathcal{N}, V) \mid e'}$$

$$\text{COMBRED-EVENT} \quad \frac{\forall \varepsilon', \mathcal{M}'', \mathcal{V}'', e'', r', ms'. \mathcal{M}, \mathcal{V} \vdash e \xrightarrow{\varepsilon'} e', [] \wedge \mathcal{M} \mid \mathcal{V} \xrightarrow{\varepsilon, r', ms'} \mathcal{M}'' \mid \mathcal{V}'' \implies \mathcal{M}, \mathcal{N}, \mathcal{V} \vdash \text{RaceFree}(\varepsilon')}{\mathcal{M}, \mathcal{V} \vdash e^\varepsilon \rightarrow e', [] \quad \mathcal{M} \mid V \xrightarrow{\varepsilon, r, ms} \mathcal{M}' \mid V' \quad \mathcal{N} \xrightarrow{\varepsilon, r, ms} \mathcal{N}'}$$

$$\frac{}{(\mathcal{M}, \mathcal{N}, V) \mid e \xrightarrow{\varepsilon, []} (\mathcal{M}, \mathcal{N}, V) \mid e'}$$

Figure 10: Combined machine and expression semantics

$$\boxed{\varsigma \mid \mathcal{T}\mathcal{S} \rightarrow \varsigma' \mid \mathcal{T}\mathcal{S}'}$$

$$\text{ForkView}(\mathcal{V}) ::= (\emptyset, \emptyset, \mathcal{V}.cur, \mathcal{V}.cur)$$

$$\text{OT-STEP} \quad \frac{\mathcal{T}\mathcal{S}(\pi) = (e, V) \quad (\mathcal{M}, \mathcal{N}, V) \mid e \xrightarrow{\varepsilon, [e_{f,0}, \dots, e_{f,n}]} (\mathcal{M}', \mathcal{N}', V') \mid e' \quad \{\rho_0 \dots \rho_n\} \cap \text{dom} \mathcal{T}\mathcal{S} = \emptyset}{(\mathcal{M}, \mathcal{N}) \mid \mathcal{T}\mathcal{S} \rightarrow (\mathcal{M}', \mathcal{N}') \mid \mathcal{T}\mathcal{S} [\pi \leftarrow (e', V')] [\rho_i \leftarrow (e_{f,i}, \text{ForkView}(V')) \mid i \in [< n]]}$$

Figure 11: Threadpool semantics.

2 Correspondence of ORC11 to RC11

The memory model of ORC11 is modeled after Lahav et al. [2017] (referred to as “RC11” from now on) without SC accesses and SC fences. It is worth noting that the memory model of ORC11 is more conservative and declares more programs racy than RC11. To prove this, we show that any program that is racy under RC11 is also considered racy by ORC11. We make this claim more precise below.

The race detector in ORC11 (and the one in the intermediate OGS machine) is stronger, *i.e.*, detects more races, than RC11. In particular, ORC11 does not permit reducing a **CAS** expression with order **acq** in the presence of an unsynchronized non-atomic read *even* when the **CAS** itself synchronizes with the non-atomic read. In contrast, the self-synchronizing nature of **CAS** leads to RC11 accepting this particular behavior as non-racy.

To simplify the proof, we allow RC11 to take expression reduction steps that are disallowed in ORC11. In particular, the declarative semantics in RC11 may compare arbitrary values with each other, whereas ORC11 will get stuck in some of these cases (see Fig. 5). A potential theorem to prove would then be that ORC11 detects any RC11 race or gets stuck for other reasons. Fortunately, the race detector in ORC11 already models races as being stuck and so the theorem statement simply becomes: *Any program that is racy under RC11 will get stuck under ORC11* (see Theorem 1).

Definition 1 (Extended Order) The set of *extended orders* $ExtOrder$ is defined by

$$o \in ExtOrder := Order \uplus \{\mathbf{relacq}\}.$$

Note that $\mathbf{relacq} \sqsupseteq o$ for any (extended) order o . We define $o.w$ and $o.r$ s.t.

$$o.w, o.r := \begin{cases} \mathbf{rel}, \mathbf{acq} & \text{if } o = \mathbf{relacq} \\ \mathbf{rel}, \mathbf{rlx} & \text{if } o = \mathbf{rel} \\ \mathbf{rlx}, \mathbf{acq} & \text{if } o = \mathbf{acq} \\ \mathbf{rlx}, \mathbf{rlx} & \text{if } o = \mathbf{rlx} \\ \mathbf{na}, \mathbf{na} & \text{if } o = \mathbf{na} \end{cases}$$

Definition 2 (Labels) The set of *labels*, $Label$, is defined by the following (tagged) union of events:

$$\begin{aligned} \gamma \in Label := & \{R^o(\ell, v) \mid o \in Order, \ell \in Loc, v \in Val\} \\ & \cup \{W^o(\ell, v) \mid o \in Order, \ell \in Loc, v \in Val\} \\ & \cup \{U^o(\ell, v_r, v_w) \mid o \in ExtOrder, \ell \in Loc, v_r \in \text{codom}(\vdash \cdot =^? \cdot), v_w \in Val\} \\ & \cup \{F^o \mid o \in \{\mathbf{rel}, \mathbf{acq}\}\} \\ & \cup \{\mathbf{Fork}^\rho \mid \rho \in Thread\} \end{aligned}$$

We write $\gamma \sim \varepsilon$ when γ corresponds a memory event ε (mapping all labels except **Fork** to their corresponding counterparts in $MemEvent$).

2.1 Executions

An execution G is defined by:

1. a finite set of events $E \subseteq \mathbb{N}$. with events $E \supseteq E_0 := \{a_0^\ell \mid \ell \in \mathcal{L}\}$.
2. a labelling function $\text{lab} \in E \rightarrow \text{Label}$, with projections $\text{typ}, \text{mod}, \text{loc}, \text{val}_r, \text{val}_w$ where defined.
3. a function tid assigning a thread identifier to every event in E . We write E^π to denote the events in E with $\text{tid}(a) = \pi$.
4. a strict partial order $\text{sb} \subseteq E \times E$ which is total on E^π for every thread π . and which puts all events in E_0 before all other events.
5. a binary relation $\text{rf} \subseteq [\text{WU}]; =_{\text{loc}}; [\text{RU}]$ such that
 - (a) $\forall \langle a, b \rangle \in \text{rf}. \text{val}_w(a) = \text{val}_r(b)$
 - (b) $\forall b, \langle a_1, b \rangle \in \text{rf}, \langle a_2, b \rangle \in \text{rf}. a_1 = a_2$.
6. a family of strict total orders $\{\text{mo}_\ell\}_{\ell \in \mathcal{L}}$ and $\text{mo} := \uplus_{\ell \in \mathcal{L}} \text{mo}_\ell$.

2.1.1 Consistent Executions

Definition 3 (Completeness) An execution G is called *complete* if and only if for every $a \in R$ we have $\text{val}_r(a) = \text{⊥} \vee \exists b \in W_{\text{loc}(a)}. \langle b, a \rangle \in \text{rf}$. Note that this condition is weaker than in RC11 as it allows reads from uninitialized locations (signified by the value ⊥).

Definition 4 (Auxiliary relations)

$$\begin{aligned}
\text{rb} &:= \text{rf}^{-1}; \text{mo} \\
\text{eco} &:= (\text{rf} \cup \text{mo} \cup \text{rb})^+ \\
\text{rs} &:= [\text{WU}]; \text{sb}|_{=_{\text{loc}}}^?; [(\text{WU})^{\exists \text{rlx}}]; (\text{rf}; [U])^* \\
\text{asw} &:= [\text{Fork}_\rho]; (\text{sb}|_{\text{tid}=\rho}^?); [E^\rho] \\
\text{sw} &:= \text{asw} \cup ([E]^{\exists \text{rel}}]; ([F]; \text{sb})^?; \text{rs}; \text{rf}; [(\text{RU})^{\exists \text{rlx}}]; (\text{sb}; [F])^?; [E]^{\exists \text{acq}}] \\
\text{hb} &:= (\text{sb} \cup \text{sw})^+
\end{aligned}$$

Definition 5 (Consistency) An execution is called RC11-*consistent* (simply “consistent” from now on) if it is complete **and**

- $\text{hb}; \text{eco}^?$ is irreflexive (COHERENCE)
- $\text{sb} \cup \text{rf}$ is acyclic (NO-THIN-AIR)

This definition does not include RC11’s SC axiom.

2.2 Declarative Semantics

The following definitions are taken from Kaiser et al. [2017] (“iGPS”) and, if necessary, adapted to our setting. Below we define threadpool reduction that generates *traces*. Note that we circumvent checks (such as those for legal comparisons) in the expression reduction by providing existentially quantified memory \mathcal{M} and local view V .

$$\begin{array}{c}
\text{TRACE-RED-SILENT} \\
\frac{\mathcal{M}, \mathcal{V} \vdash \mathcal{TS}(\pi) \rightarrow e, []}{\mathcal{TS} \xrightarrow{\varepsilon}^{\pi} \mathcal{TS}[\pi \mapsto e]}
\end{array}
\qquad
\begin{array}{c}
\text{TRACE-RED-MEM} \\
\frac{\gamma \sim \varepsilon \quad \mathcal{M}, \mathcal{V} \vdash \mathcal{TS}(\pi) \xrightarrow{\varepsilon} e, []}{\mathcal{TS} \xrightarrow{\gamma}^{\pi} \mathcal{TS}[\pi \mapsto e]}
\end{array}$$

$$\begin{array}{c}
\text{TRACE-RED-FORK} \\
\frac{\mathcal{V}(\pi) = (e, V) \quad \mathcal{M}, \mathcal{V} \vdash \mathcal{TS}(\pi) \rightarrow e', e_f \quad \rho \notin \text{dom}(\mathcal{TS})}{\mathcal{TS} \xrightarrow{\text{Fork}_{\rho}}^{\pi} \mathcal{TS}[\pi \mapsto e'] \uplus [\rho \mapsto e_f]}
\end{array}$$

We write $\mathcal{TS} \Rightarrow^{\pi} \mathcal{TS}'$ if $\mathcal{TS} \xrightarrow{x}^{\pi} \mathcal{TS}'$ for some transition label x ; $\mathcal{TS} \xrightarrow{\pi} \mathcal{TS}'$ if $\mathcal{TS} \xrightarrow{x}^{\pi} \mathcal{TS}'$ for some thread identifier π ; and $\mathcal{TS} \Rightarrow \mathcal{TS}'$ if $\mathcal{TS} \xrightarrow{x}^{\pi} \mathcal{TS}'$ for some transition label x and thread identifier π . A threadpool is called *final* if $\mathcal{TS}(\pi) \in \text{Val}$ for every $\pi \in \text{dom}(\mathcal{TS})$.

Definition 6 (Traces) A *trace* is a sequence of pairs $\langle \gamma_1, \pi_1 \rangle, \dots, \langle \gamma_n, \pi_n \rangle$. We say that $tr = \langle \gamma_1, \pi_1 \rangle, \dots, \langle \gamma_n, \pi_n \rangle$ is a trace of an expression e if

$$[0 \mapsto e] \xrightarrow{\varepsilon}^* \xrightarrow{\gamma_1}^{\pi_1} \xrightarrow{\varepsilon}^* \dots \xrightarrow{\varepsilon}^* \xrightarrow{\gamma_n}^{\pi_n} \xrightarrow{\varepsilon}^* \mathcal{TS}$$

for some thread π and threadpool \mathcal{TS} . When \mathcal{TS} is final, we call tr a *full* trace.

Definition 7 A trace $tr = \langle \gamma_1, \pi_1 \rangle, \dots, \langle \gamma_n, \pi_n \rangle$ induces partial order on indices $\text{sb}(tr)$, called *sequenced-before*, and a relation on indices $\text{asw}(tr)$, called *additional-synchronized-with*. They are defined by:

$$\frac{i < j \quad \pi_i = \pi_j}{\langle i, j \rangle \in \text{sb}(tr)}
\qquad
\frac{\langle i, j \rangle \in \text{sb}(tr) \quad \langle j, k \rangle \in \text{sb}(tr)}{\langle i, k \rangle \in \text{sb}(tr)}$$

$$\frac{i < j \quad \gamma_i = \text{Fork}_{\pi_j}}{\langle i, j \rangle \in \text{asw}(tr)}$$

Lemma 1 Let tr be a trace of an expression e . Then

- Any prefix of tr is also a trace of e .
- Any permutation tr' of tr with $\text{sb}(tr') = \text{sb}(tr)$ and $\text{asw}(tr') = \text{asw}(tr)$ is a trace of e .

Definition 8 An execution G *follows* a trace $tr = \langle \gamma_1, \pi_1 \rangle, \dots, \langle \gamma_n, \pi_n \rangle$ if:

- $E = \{a_1, \dots, a_n\}$ such that $\text{lab}(a_k) = \gamma_k$ and $\text{tid}(a_k) = \pi_k$ for every $1 \leq k < n$
- $\text{sb} = \{\langle a_i, a_j \rangle \mid \langle i, j \rangle \in \text{sb}(tr)\}$.

We call G an execution of expression e if G follows some trace of e .

Definition 9 (Conflict) Two events a and b are called *conflicting* in an execution G if $a, b \in E$, $\{\text{typ}(a), \text{typ}(b)\} \cap \{\mathbb{W}, \mathbb{U}\} \neq \emptyset$, $a \neq b$, and $\text{loc}(a) = \text{loc}(b)$.

Definition 10 (Races) A pair $\langle a, b \rangle$ is called a *race* in G if a and b are conflicting events in G , and $\langle a, b \rangle \notin \text{hb} \cup \text{hb}^{-1}$. An execution G is called *racy* if there is some race $\langle a, b \rangle$ in G with $\text{na} \in \{\text{mod}(a), \text{mod}(b)\}$.

Definition 11 (Bugginess) An execution G is buggy if it is racy. An expression e is buggy if some consistent execution of e is buggy.

2.3 Operational Graph Semantics (OGS)

We now introduce an operationalized account of RC11 (*OGS*, short for Operational Graph Semantics), in which we build up executions step by step. This serves as an important stepping stone towards a our correspondence proof with ORC11.

Definition 12 (Execution Extension: Memory Accesses) We write $G' \in \text{Add}(G, \pi, \rho, \gamma)$ if there exists an event a s.t.

- $G'.E = G.E \uplus \{a\}$, $G'.\text{tid} = G.\text{tid} \cup \{a \mapsto \rho\}$, $G'.\text{lab} = G.\text{lab} \cup \{a \mapsto \gamma\}$
- if $\rho \neq \pi$ then $\rho \notin \text{codom}(G.\text{tid})$
- $G'.\text{sb} = (G.\text{sb} \uplus (G.E^\rho \times \{a\}))^+$
- $G'.\text{rf} \supseteq G.\text{rf}$
- $G'.\text{mo} \supseteq G.\text{mo}$ and if $\gamma = W^{\text{na}}(_, _)$ then a is **mo**-maximal in G'

Definition 13 (Race Predicate) We define a predicate $\text{Race}(G, \pi)$ which holds for all memory events that would cause a data race in execution G . Note that this race detector models exactly the rules implement in ORC11. Thus, it detects more races than RC11 *but* only in (potentially non-buggy) executions following buggy expressions.

$$\frac{\text{RACE-I} \quad o \sqsupseteq \mathbf{rlx} \quad \gamma \in (\mathbf{RU})_\ell^o \quad \exists a \in W_\ell^{\text{na}}. \forall b \in E^\pi. \langle a, b \rangle \notin \mathbf{hb}^*}{\gamma \in \text{Race}(G, \pi)}$$

$$\frac{\text{RACE-II} \quad \gamma = \mathbf{R}^{\text{na}}(\ell, _) \quad \exists a \in (\mathbf{WU})_\ell. \forall b \in E^\pi. \langle a, b \rangle \notin \mathbf{hb}^*}{\gamma \in \text{Race}(G, \pi)}$$

$$\frac{\text{RACE-III} \quad \gamma = W^{\text{na}}(\ell, _) \quad \exists a \in (\mathbf{RWU})_\ell. \forall b \in E^\pi. \langle a, b \rangle \notin \mathbf{hb}^*}{\gamma \in \text{Race}(G, \pi)}$$

$$\frac{\text{RACE-IV} \quad o \sqsupseteq \mathbf{rlx} \quad \gamma = W_\ell^o \quad \exists a \in (\mathbf{RW})_\ell^{\text{na}}. \forall b \in E^\pi. \langle a, b \rangle \notin \mathbf{hb}^*}{\gamma \in \text{Race}(G, \pi)}$$

$$\frac{\text{RACE-V} \quad \gamma = U_\ell^o \quad \exists a \in (\mathbf{RW})_\ell^{\text{na}}. \forall b \in E^\pi. \langle a, b \rangle \notin \mathbf{hb}^*}{\gamma \in \text{Race}(G, \pi)}$$

Definition 14 (OGS Reductions)

$$\frac{\text{OGS-MEMORY-STEP} \quad \begin{array}{l} \gamma \in \{\mathbf{R}^o(\ell, v), \mathbf{W}^o(\ell, v), \mathbf{F}^o\} \\ \gamma \notin \text{Race}(G, \pi) \\ G' \in \text{Add}(G, \pi, \pi, \gamma) \\ G' \text{ is consistent} \end{array}}{G \xrightarrow{\gamma}^\pi G'}$$

$$\frac{\text{OGS-FORK} \quad \begin{array}{l} G' \in \text{Add}(G, \pi, \rho, \text{Fork}_\rho) \\ G' \text{ is consistent} \end{array}}{G \xrightarrow{\text{Fork}_\rho}^\pi G'}$$

$$\frac{\text{OGS-RACE} \quad \gamma \in \text{Race}(G, \pi)}{G \xrightarrow{\gamma}^\pi \perp_{\text{race}}}$$

We define combined machine and expression semantics for OGS. We once again allow expression reductions to proceed independent of the current state, thus capturing more behaviors than those allowed by ORC11.

$$\begin{array}{c}
\text{OGS-COMBRED-PURE} \\
\frac{\mathcal{M}, \mathcal{V} \vdash e \rightarrow e', es}{G \mid e \xrightarrow{\perp, es} G \mid e'} \\
\\
\text{OGS-COMBRED-EVENT} \\
\frac{\forall \varepsilon', \mathcal{M}'', \mathcal{V}'', e''. \mathcal{M}, \mathcal{V} \vdash e \xrightarrow{\varepsilon'} e', [] \implies \neg(G \xrightarrow{\varepsilon} \pi \perp_{\text{race}}) \quad \mathcal{M}, \mathcal{V} \vdash e^\varepsilon \rightarrow e', [] \quad G \xrightarrow{\varepsilon} \pi G'}{G \mid e \xrightarrow{\varepsilon, []} \pi G' \mid e'} \\
\\
\text{OGS-OT-STEP} \\
\frac{\mathcal{TS}(\pi) = e \quad G \mid e \xrightarrow{\varepsilon, [e_{f,0}, \dots, e_{f,n}]} \pi G' \mid e' \quad \{\rho_0 \dots \rho_n\} \cap \text{dom } \mathcal{TS} = \emptyset}{G \mid \mathcal{TS} \rightarrow G' \mid \mathcal{TS}[\pi \leftarrow (e', V')] [\rho_i \leftarrow e_{f,i} \mid i \in [< n]]}
\end{array}$$

We define G_0 to be an execution in which all locations are allocated with an initial value of 0.

Lemma 2 (Inclusion of Behaviors (I)) Let G be a non-buggy, consistent execution of expression e . Then there exists a trace $tr = \langle \varepsilon_1, \pi_1 \rangle \dots \langle \varepsilon_n, \pi_n \rangle$ of e such that $G_0 \xrightarrow{\varepsilon_1} \pi_1 \dots \xrightarrow{\varepsilon_n} \pi_n G \vee \exists j \leq n. G_0 \xrightarrow{\varepsilon_1} \pi_1 \dots \xrightarrow{\varepsilon_j} \pi_j \perp_{\text{race}}$.

Proof. As G is consistent, we have that $\mathbf{sb} \cup \mathbf{rf}$ is acyclic. Let a_1, \dots, a_n be an enumeration of \mathbf{E} that respects $(\mathbf{sb} \cup \mathbf{rf})^+$. For every $1 \leq i \leq n$, let $\pi_i := \text{tid}(a_i)$, $\varepsilon_i = \text{lab}(a_i)$, and $tr = \langle \varepsilon_1, \pi_1 \rangle \dots \langle \varepsilon_n, \pi_n \rangle$. Adding events a_1, \dots, a_n one-by-one we can thus establish either $G_0 \xrightarrow{\varepsilon_1} \pi_1 \dots \xrightarrow{\varepsilon_n} \pi_n G$, or—if in any step $j \leq n$ the race predicate detects a spurious race— $G_0 \xrightarrow{\varepsilon_1} \pi_1 \dots \xrightarrow{\varepsilon_j} \pi_j \perp_{\text{race}}$. \square

Lemma 3 (Inclusion of Behaviors (II)) Let e be a buggy expression. Then $G_0; [0 \mapsto e] \rightarrow^* \perp_{\text{race}}$.

Proof. We have that e is buggy and, thus, a consistent execution G which is buggy. Let a_1, \dots, a_n be an enumeration of \mathbf{E} that respects $\mathbf{sb} \cup \mathbf{rf}$. Let k be the minimal index such that $G \cap \{a_1, \dots, a_k\}$ is buggy, *i.e.*, racy.

We thus have that $G \cap \{a_1, \dots, a_k\}$ is racy. Let $j < k$ be the minimal index such that $\text{loc}(a_k) = \text{loc}(a_j)$, $\langle a_k, a_j \rangle \notin \mathbf{hb} \cup \mathbf{hb}^{-1}$, and one of the following holds:

- $a_k \in (\mathbf{WU})^{\exists \text{rlx}} \wedge a_j \in \mathbf{R}^{\text{na}} \vee a_k \in \mathbf{W}^{\text{na}}$
- $a_j \in (\mathbf{WU})^{\exists \text{rlx}} \wedge a_k \in \mathbf{R}^{\text{na}} \vee a_j \in \mathbf{W}^{\text{na}}$

Note that we have $\text{tid}(a_k) \neq \text{tid}(a_j)$, as otherwise these events would be related by $G.\mathbf{sb}$, and, thus $G.\mathbf{hb}$.

1. $a_k \in \mathbf{W}^{\text{na}}$. We define $B := \{a \in \mathbf{E} \mid \langle a, a_j \rangle \in G.\mathbf{hb} \vee \langle a, a_k \rangle \in G.\mathbf{hb}^*\}$ and $G' := G \cap B$. Note that G' is non-empty, consistent, and not buggy. By **Lemma 2**, we have that $G_0 \xrightarrow{\varepsilon_1} \pi_1 \dots \xrightarrow{\varepsilon_n} \pi_n G' \vee \exists j \leq n. G_0 \xrightarrow{\varepsilon_1} \pi_1 \dots \xrightarrow{\varepsilon_j} \pi_j \perp_{\text{race}}$ for some trace $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$ of e . In the latter case our proof is done. Otherwise we have $\langle a_k, a_j \rangle \notin G'.\mathbf{hb}$ and we show that $G' \xrightarrow{\text{lab}(a_j)} \text{tid}(a_j) \perp_{\text{race}}$. By **Definition 13** (using whichever case corresponds to $\text{lab}(a_j)$), it suffices to show that $\langle a_k, b \rangle \notin G'.\mathbf{hb}^*$ for all $b \in \mathbf{E}^{\text{tid}(a_j)}$. By way of contradiction, assume $b \in \mathbf{E}^{\text{tid}(a_j)}$ and $\langle a_k, b \rangle \in G'.\mathbf{hb}^*$. By definition of G' , we have $\langle b, a_j \rangle \in G.\mathbf{hb} \vee \langle b, a_k \rangle \in G.\mathbf{hb}^*$.

- (a) $\langle b, a_j \rangle \in G.\mathbf{hb}$. By transitivity, we have $\langle a_k, a_j \rangle \in G.\mathbf{hb}$, which contradicts our assumption.
- (b) $\langle b, a_k \rangle \in G.\mathbf{hb}^*$. From $\text{tid}(a_k) \neq \text{tid}(a_j)$ we have that $b \neq a_k$. Thus, $\langle b, a_k \rangle \in G.\mathbf{hb}$. By transitivity, we have $\langle b, b \rangle \in G.\mathbf{hb}$, which contradicts \mathbf{hb} 's irreflexivity.

As $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle, \langle \text{lab}(a_j), \text{tid}(a_j) \rangle$ is a valid trace for e , we have that $G_0; [0 \mapsto e] \rightarrow^* \perp_{\text{race}}$.

2. $a_j \in \mathbb{W}^{\text{na}}$ is symmetric to the case above.
3. $a_k \in (\mathbb{W}\mathbb{U})^{\exists \text{rlx}} \wedge a_j \in \mathbb{R}^{\text{na}}$. We define $B := \{a \in \mathbb{E} \mid \langle a, a_j \rangle \in G.\mathbf{hb} \vee \langle a, a_k \rangle \in G.\mathbf{hb}^*\}$ and $G' := G \cap B$. Note that G' is consistent and not buggy. By Lemma 2, we have that $G_0 \xrightarrow{\varepsilon_1, \pi_1} \dots \xrightarrow{\varepsilon_n, \pi_n} G' \vee \exists j \leq n. G_0 \xrightarrow{\varepsilon_1, \pi_1} \dots \xrightarrow{\varepsilon_j, \pi_j} \perp_{\text{race}}$ for some trace $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$ of e . In the latter case our proof is done. Otherwise we have $\langle a_k, a_j \rangle \notin G'.\mathbf{hb}$ and we show that $G' \xrightarrow{\text{lab}(a_j), \text{tid}(a_j)} \perp_{\text{race}}$.

By Definition 13, it suffices to show that $\langle a_k, b \rangle \notin G'.\mathbf{hb}^*$ for all $b \in \mathbb{E}^{\text{tid}(a_j)}$. By way of contradiction, assume $b \in \mathbb{E}^{\text{tid}(a_j)}$ and $\langle a_k, b \rangle \in G'.\mathbf{hb}^*$. By definition of G' , we have $\langle b, a_j \rangle \in G.\mathbf{hb} \vee \langle b, a_k \rangle \in G.\mathbf{hb}^*$.

- (a) $\langle b, a_j \rangle \in G.\mathbf{hb}$. By transitivity, we have $\langle a_k, a_j \rangle \in G.\mathbf{hb}$, which contradicts our assumption.
- (b) $\langle b, a_k \rangle \in G.\mathbf{hb}^*$. From $\text{tid}(a_k) \neq \text{tid}(a_j)$ we have that $b \neq a_k$. Thus, $\langle b, a_k \rangle \in G.\mathbf{hb}$. By transitivity, we have $\langle b, b \rangle \in G'.\mathbf{hb}$, which contradicts \mathbf{hb} 's irreflexivity.

As $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle, \langle \text{lab}(a_j), \text{tid}(a_j) \rangle$ is a valid trace for e , we have that $G_0; [0 \mapsto e] \rightarrow^* \perp_{\text{race}}$.

4. $a_j \in \mathbb{W}^{\exists \text{rlx}} \wedge a_k \in \mathbb{R}^{\text{na}}$. This case is symmetric to the one above.

□

2.4 OGS to ORC11

Definition 15 We define auxiliary relations $\{\mathbf{auxrel}\}_\ell$, $\mathbf{auxfrel}$, and \mathbf{auxacq} .

$$\begin{aligned} \mathbf{auxrel}_\ell &:= \mathbf{hb}; [\mathbb{W}_\ell^{\text{rel}}] \\ \mathbf{auxfrel} &:= \mathbf{hb}; [\mathbb{F}^{\text{rel}}] \\ \mathbf{auxacq} &:= \mathbf{hb}; ([\mathbb{E}^{\text{rel}}]; ([\mathbb{F}]; \mathbf{sb})^?; \mathbf{rs}; \mathbf{rf}; [\mathbb{R}^{\text{rlx}}])^? \end{aligned}$$

Definition 16 (Timestamp Assignment) A *timestamp assignment* for an execution graph G is a function $ts : \mathbb{W} \rightarrow \text{Time}$, that satisfies $ts(a) < ts(b)$ whenever $\langle a, b \rangle \in \mathbf{mo}$. Given a timestamp assignment ts for G and an event $a \in \mathbb{W}$, the view of an event a in G according to ts , denoted $\text{view}(a, G, ts)$, and the message induced by a in G according to ts , denoted $\text{msg}(a, G, ts)$, are given by:

$$\text{view}(a, G, ts) = \begin{cases} \lambda \ell. \max \{ts(b) \mid b \in \mathbb{W}_\ell, \langle b, a \rangle \in \mathbf{hb}^*\} & \text{if } \text{mod}(a) = \mathbf{rel} \\ \lambda \ell. \max \{ts(b) \mid b \in \mathbb{W}_\ell, \langle b, a \rangle \in (\mathbf{auxfrel} \cup \mathbf{auxrel}_\ell)^?; \mathbf{hb}^*\} & \text{if } \text{mod}(a) = \mathbf{rlx} \\ \perp & \text{otherwise} \end{cases}$$

$$\text{msg}(a, G, ts) = (\text{val}_w(a), \text{view}(a, G, ts))$$

In these definitions, we take \perp to be the maximum of an empty set.

Definition 17 (Event Injection) Let G be an execution and $a \in E$. We define an injection into natural numbers, written $\text{Inj}(G, a)$, as follows.

$$\text{Inj}(G, a) := \text{prime}(\text{tid}(a))^{|\{b \mid \langle b, a \rangle \in \text{sb}\}|}$$

where $\text{prime}(n)$ is the n^{th} prime number. Note that Inj is injective and that performing a machine step $G \rightarrow G'$ implies $\text{Inj}(G', a) = \text{Inj}(G, a)$ for any $a \in G$.

We write $\text{Inj}(G, X)$ for $\{\text{Inj}(G, a) \mid a \in X\}$. We also write $a \in Y$ for $\text{Inj}(G, a) \in Y$ if Y is defined as $\text{Inj}(G, X)$ for some X . (Note that this implies $a \in X$.)

Definition 18 Let G be an execution and ts be a timestamp assignment for G . We define the physical state $(\mathcal{M}_G^{ts}, \mathcal{V}_G^{ts}, \mathcal{N}_G^{ts})$ as follows.

- The memory is defined by $\mathcal{M}_G^{ts} := \lambda \ell. \lambda t. \begin{cases} \text{msg}(a, G, ts) & \text{if } \exists a. t = ts(a) \\ \perp & \text{otherwise} \end{cases}$
- The thread view \mathcal{V}_G^{ts} is defined by

$$\begin{aligned} \text{ThEvs}(X, S, R) &:= \{a \in S \mid \exists b \in X. \langle a, b \rangle \in R^*\} \\ t_{\max}(X, S, R) &:= \max \{ts(a) \mid a \in \text{ThEvs}(X, S, R)\} \\ V(X, R) &:= \lambda \ell. \{ \mathbf{w} := t_{\max}(X, W_\ell, R), \\ &\quad \mathbf{aw} := \{ts(a) \mid a \in \text{ThEvs}(X, W_\ell^{\exists r1x}, R)\}, \\ &\quad \mathbf{nr} := \text{Inj}(G, \text{ThEvs}(X, R_\ell^{\text{na}}, R)), \\ &\quad \mathbf{ar} := \text{Inj}(G, \text{ThEvs}(X, R_\ell^{\exists r1x}, R)) \\ &\quad \} \\ \mathcal{V}_G^{ts}(\pi) &:= \{ \mathbf{rel} := \lambda \ell'. V(E^\pi, \mathbf{auxrel}_{\ell'}), \\ &\quad \mathbf{frel} := V(E^\pi, \mathbf{auxfrel}), \\ &\quad \mathbf{cur} := V(E^\pi, \mathbf{hb}), \\ &\quad \mathbf{acq} := V(E^\pi, \mathbf{auxacq}) \\ &\quad \} \end{aligned}$$

- The global race detector state \mathcal{N}_G^{ts} is defined by

$$\begin{aligned} \mathcal{N}_G^{ts} &:= \lambda \ell. \{ \\ &\quad \mathbf{w} := t_{\max}(E, W_\ell^{\text{na}}, (=)), \\ &\quad \mathbf{aw} := \{ts(a) \mid a \in W_\ell^{\exists r1x}\}, \\ &\quad \mathbf{nr} := \text{Inj}(G, R_\ell^{\text{na}}), \\ &\quad \mathbf{ar} := \text{Inj}(G, R_\ell^{\exists r1x}) \\ &\quad \} \end{aligned}$$

In these definitions, we take \perp to be the maximum of an empty set.

We say that G *relates* to a physical state $(\mathcal{M}, \mathcal{V}, \mathcal{N})$, denoted $G \sim (\mathcal{M}, \mathcal{V}, \mathcal{N})$, if and only if $(\mathcal{M}_G^{ts}, \mathcal{V}_G^{ts}, \mathcal{N}^{ts}) = (\mathcal{M}, \mathcal{V}, \mathcal{N})$.

Definition 19 In the following, we lift ORC11's machine semantics to thread views such that

$$(\mathcal{M}, \mathcal{N}, \mathcal{V}) \xrightarrow{\varepsilon}^{\pi} (\mathcal{M}, \mathcal{N}', \mathcal{V}') := (\mathcal{M}, \mathcal{N}) \mid \mathcal{V}(\pi) \xrightarrow{\varepsilon} (\mathcal{M}', \mathcal{N}') \mid V' \wedge \mathcal{V}' = \mathcal{V}[\pi \leftarrow V']$$

Lemma 4 Suppose $G \xrightarrow{\gamma}^{\pi} G'$, $\gamma \sim \varepsilon$ and let ts' be a timestamp assignment for G' . Then $ts \upharpoonright_{G.W}$ is a timestamp assignment for G and $(\mathcal{M}_G^{ts}, \mathcal{V}_G^{ts}, \mathcal{N}_G^{ts}) \xrightarrow{\varepsilon}^{\pi} (\mathcal{M}_{G'}^{ts'}, \mathcal{V}_{G'}^{ts'}, \mathcal{N}_{G'}^{ts'})$.

Lemma 5 (Inclusion of Behaviors (I)) Suppose $G \xrightarrow{\gamma_1}^{\pi_1} \dots \xrightarrow{\gamma_n}^{\pi_n} G_n$ and $G \sim (\mathcal{M}_1, \mathcal{V}_1, \mathcal{N}_1)$. Then either

- there exist $\varepsilon_1 \dots \varepsilon_n$, $(\mathcal{M}_2, \mathcal{N}_2, \mathcal{V}_2) G_2 \dots (\mathcal{M}_n, \mathcal{N}_n, \mathcal{V}_n) G_n$ such that $\xrightarrow{\pi_1} [\varepsilon_1] \dots \xrightarrow{\varepsilon_n}^{\pi_n} (\mathcal{M}_n, \mathcal{N}_n, \mathcal{V}_n)$.
- or there exist $j < n$, $\varepsilon_1 \dots \varepsilon_{j+1}$, $(\mathcal{M}_2, \mathcal{N}_2, \mathcal{V}_2) G_2 \dots (\mathcal{M}_j, \mathcal{N}_j, \mathcal{V}_j) G_j$ such that $\xrightarrow{\pi_1} [\varepsilon_1] \dots \xrightarrow{\varepsilon_j}^{\pi_j} (\mathcal{M}_{j-1}, \mathcal{N}_j, \mathcal{V}_j) \wedge \neg \left((\mathcal{M}_j, \mathcal{N}_j, \mathcal{V}_j) \xrightarrow{\varepsilon_{j+1}}^{\pi_{j+1}} _ \right)$.

Lemma 6 (Inclusion of Behaviors (II)) Let G be a consistent execution that is not buggy, $G \xrightarrow{\gamma}^{\pi} \perp_{\text{race}}$, $G \sim (\mathcal{M}, \mathcal{V}, \mathcal{N})$, and $\gamma \sim \varepsilon$. Then $\neg \left((\mathcal{M}, \mathcal{N}) \mid \mathcal{V}(\pi) \xrightarrow{\varepsilon} _ \right)$

Proof. Let ts be the timestamp assignment implied by $G \sim (\mathcal{M}, \mathcal{V}, \mathcal{N})$. We consider the following cases.

1. $\gamma \in \{\mathbf{R}^o(\ell, _), \mathbf{U}^o(\ell, _, _)\} \wedge o \sqsupseteq \mathbf{rlx} \wedge \exists a \in \mathbf{W}_{\ell}^{\text{na}}. \forall b \in \mathbf{E}^{\pi}. \langle a, b \rangle \notin \mathbf{hb}^*$.

We show $\neg(\mathcal{M}, \mathcal{N}, \mathcal{V}(\pi) \vdash \text{RaceFree}(\langle \text{Read}, \ell, _, o \rangle))$. It suffices to show that $\mathcal{V}(\pi). \text{cur}(\ell) < \mathcal{N}(\ell)$. Let $a_m \in \mathbf{W}_{\ell}^{\text{na}}$ be the **mo**-maximal non-atomic write event on ℓ (which implies $ts(a_m) \geq ts(a)$). Then $\mathcal{N}(\ell) = ts(a_m)$. It thus suffices to show that $\mathcal{V}(\pi). \text{cur}(\ell) < ts(a_m)$. By way of contradiction, assume that $\mathcal{V}(\pi). \text{cur}(\ell) \geq ts(a_m)$. Then, there exists $c \in \mathbf{W}_{\ell}$ and $b \in \mathbf{E}^{\pi}$ s.t. $\langle c, b \rangle \in \mathbf{hb}^* \wedge ts(c) \geq ts(a_m)$. From $\langle a, a_m \rangle \in \mathbf{mo}^*$, COHERENCE, and G being non-racy we have that $\langle a, a_m \rangle \in \mathbf{hb}^*$. As G is non-racy, we also have $c = a_m \vee \langle a_m, c \rangle \in \mathbf{hb} \vee \langle c, a_m \rangle \in \mathbf{hb}$.

- (a) $c = a_m$. We have $\langle a_m, b \rangle \in \mathbf{hb}^*$. Then, by transitivity, we have $\langle a, b \rangle \in \mathbf{hb}^*$ which contradicts our initial assumption.
- (b) $\langle a_m, c \rangle \in \mathbf{hb}$. By transitivity, we have $\langle a_m, b \rangle \in \mathbf{hb}^*$, and, thus, $\langle a, b \rangle \in \mathbf{hb}^*$. This contradicts our initial assumption.
- (c) $\langle c, a_m \rangle \in \mathbf{hb} \wedge c \neq a$. By COHERENCE, we have $\langle a_m, c \rangle \notin \mathbf{mo}$ and, thus, $\langle c, a_m \rangle \in \mathbf{mo}$. This contradicts $ts(c) \geq ts(a_m)$.

2. $\gamma \in \mathbf{R}^{\text{na}}(\ell, _) \wedge \exists a \in (\mathbf{WU})_{\ell}. \forall b \in \mathbf{E}^{\pi}. \langle a, b \rangle \notin \mathbf{hb}^*$.

We show $\neg(\mathcal{M}, \mathcal{N}, \mathcal{V}(\pi) \vdash \text{RaceFree}(\langle \text{Read}, \ell, _, \mathbf{na} \rangle))$. It suffices to show that there exists t' , $(v', V') = \mathcal{M}(\ell)(t')$ s.t. $\mathcal{V}(\pi). \text{cur}(\ell) < t' \vee \mathcal{N}(\ell). \text{aw} \not\sqsubseteq \mathcal{V}(\pi). \text{cur}(\ell). \text{aw}$.

We choose $t' = ts(a)$ and $(v', V') := \text{msg}(a, G, ts)$. It suffices to show $\mathcal{V}(\pi). \text{cur}(\ell) < ts(a) \vee \mathcal{N}(\ell). \text{aw} \not\sqsubseteq \mathcal{V}(\pi). \text{cur}(\ell). \text{aw}$. There exists $c \in \mathbf{W}_{\ell}$ and $b \in \mathbf{E}^{\pi}$ s.t. $\langle c, b \rangle \in \mathbf{hb}^*$ and $\mathcal{V}(\pi). \text{cur}(\ell) = ts(c)$.

We consider two cases:

- (a) $\text{mod}(a) = \mathbf{na}$. We show $ts(c) < ts(a)$. By way of contradiction, assume $ts(c) \geq ts(a)$. We have $c \neq a$ as otherwise $\langle a, b \rangle \in \mathbf{hb}^*$, contradicting our assumption. Thus we have $ts(c) > ts(a)$ and $\langle a, c \rangle \in \mathbf{mo}$. From G being non-racy, COHERENCE, and $\langle a, c \rangle \in \mathbf{mo}$ we have that $\langle a, c \rangle \in \mathbf{hb}$. By transitivity, $\langle a, b \rangle \in \mathbf{hb}^*$, which contradicts our assumption.
- (b) $\text{mod}(a) = \mathbf{rlx}$. We show $\mathcal{N}(\ell).\text{aw} \not\sqsubseteq \mathcal{V}(\pi).\text{cur}(\ell).\text{aw}$. By way of contradiction, assume that $\mathcal{N}(\ell).\text{aw} \sqsubseteq \mathcal{V}(\pi).\text{cur}(\ell).\text{aw}$. We have $a \in \mathcal{N}(\ell).\text{aw}$ and, thus, $a \in \mathcal{V}(\pi).\text{cur}(\ell).\text{aw}$. Hence, there exists $b' \in E^\pi$ s.t. $\langle a, b' \rangle \in \mathbf{hb}^*$, which contradicts our assumption.

3. $\gamma = \mathbf{W}^{\mathbf{na}}(\ell, v) \wedge \exists a \in (\mathbf{RWU})_\ell. \forall b \in E^\pi. \langle a, b \rangle \notin \mathbf{hb}^*$.

We show that $\neg(\mathcal{M}, \mathcal{N}, \mathcal{V}(\pi) \vdash \text{RaceFree}(\langle \text{Write}, \ell, _, \mathbf{na} \rangle))$. There exists $c \in W_\ell$ and $b \in E^\pi$ s.t. $\langle c, b \rangle \in \mathbf{hb}^* \wedge \mathcal{V}(\pi).\text{cur}(\ell) = ts(c)$. We also have $a \neq c$ as that would imply $\langle a, b \rangle \in \mathbf{hb}^*$, contradicting our assumption.

We consider the following cases.

- (a) $a \in W_\ell^{\mathbf{na}}$. We show that there exists t' , $(v', V') = \mathcal{M}(\ell)(t')$ s.t. $ts(c) < t'$. We choose $t' := ts(a)$ and $(v', V') := \text{msg}(a, G, ts)$. It suffices to show $ts(c) < ts(a)$. As G is non-racy, we have $\langle a, c \rangle \in \mathbf{hb} \vee \langle c, a \rangle \in \mathbf{hb}$. The former implies, by transitivity, that $\langle a, b \rangle \in \mathbf{hb}^*$, which would contradict our assumption. Thus, $\langle c, a \rangle \in \mathbf{hb}$. As $a \neq b$ we derive $\langle c, a \rangle \in \mathbf{mo}$ from COHERENCE and, thus, $ts(c) < ts(a)$.
- (b) $a \in (\mathbf{WU})_\ell^{\mathbf{rlx}}$. We show that $\mathcal{N}(\ell).\text{aw} \not\sqsubseteq \mathcal{V}(\pi).\text{cur}(\ell).\text{aw}$. By way of contradiction, assume that $\mathcal{N}(\ell).\text{aw} \sqsubseteq \mathcal{V}(\pi).\text{cur}(\ell).\text{aw}$. We have $a \in \mathcal{N}(\ell).\text{aw}$ and, thus, $a \in \mathcal{V}(\pi).\text{cur}(\ell).\text{aw}$. Hence, there exists $b' \in E^\pi$ s.t. $\langle a, b' \rangle \in \mathbf{hb}^*$, which contradicts our assumption.
- (c) $a \in R_\ell$. We show that $\mathcal{N}(\ell).\text{nr} \not\sqsubseteq \mathcal{V}(\pi).\text{cur}(\ell).\text{nr} \vee \mathcal{N}(\ell).\text{ar} \not\sqsubseteq \mathcal{V}(\pi).\text{cur}(\ell).\text{ar}$. We have $a \in \mathcal{N}(\ell).\text{nr} \vee a \in \mathcal{N}(\ell).\text{ar}$. By way of contradiction, assume that $\mathcal{N}(\ell).\text{nr} \sqsubseteq \mathcal{V}(\pi).\text{cur}(\ell).\text{nr} \wedge \mathcal{N}(\ell).\text{ar} \sqsubseteq \mathcal{V}(\pi).\text{cur}(\ell).\text{ar}$. Then $a \in \mathcal{V}(\pi).\text{cur}(\ell).\text{nr} \cup \mathcal{V}(\pi).\text{cur}(\ell).\text{ar}$. Hence, there exists $b' \in E^\pi$ s.t. $\langle a, b' \rangle \in \mathbf{hb}$. This contradicts our assumption.

4. $\gamma = \mathbf{W}^o(\ell, _) \wedge o \sqsupseteq \mathbf{rlx} \wedge \exists a \in (\mathbf{RW})_\ell^{\mathbf{na}}. \forall b \in E^\pi. \langle a, b \rangle \notin \mathbf{hb}^*$.

We show $\neg(\mathcal{M}, \mathcal{N}, \mathcal{V}(\pi) \vdash \text{RaceFree}(\langle \text{Write}, \ell, _, o \rangle))$. We consider $a \in R^{\mathbf{na}}$ and $a \in W^{\mathbf{na}}$ separately.

(a) $a \in R^{\mathbf{na}}$. It suffices to show that $\mathcal{N}(\ell).\text{nr} \not\sqsubseteq \mathcal{V}(\pi).\text{cur}(\ell).\text{nr}$.

By way of contradiction, assume that $\mathcal{N}(\ell).\text{nr} \sqsubseteq \mathcal{V}(\pi).\text{cur}(\ell).\text{nr}$. We have $a \in \mathcal{N}(\ell).\text{nr}$ and, thus, $a \in \mathcal{V}(\pi).\text{cur}(\ell).\text{nr}$. This implies that there exists $b' \in E^\pi$ s.t. $\langle a, b' \rangle \in \mathbf{hb}^*$, which contradicts our assumption.

(b) $a \in W^{\mathbf{na}}$. It suffices to show that $V(\pi).\text{cur}(\ell) < \mathcal{N}(\ell)$.

Let $a_m \in W_\ell^{\mathbf{na}}$ be the \mathbf{mo} -maximal non-atomic write event on ℓ (which implies $ts(a_m) \geq ts(a)$). Then $\mathcal{N}(\ell) = ts(a_m)$. It thus suffices to show that $V(\pi).\text{cur}(\ell) < ts(a_m)$. By way of contradiction, assume that $V(\pi).\text{cur}(\ell) \geq ts(a_m)$. Then, there exists $c \in W_\ell$ and $b \in E^\pi$ s.t. $\langle c, b \rangle \in \mathbf{hb}^* \wedge ts(c) \geq ts(a_m)$. From $\langle a, a_m \rangle \in \mathbf{mo}^*$, COHERENCE, and G being non-racy we have that $\langle a, a_m \rangle \in \mathbf{hb}^*$. As G is non-racy, we also have $c = a_m \vee \langle a_m, c \rangle \in \mathbf{hb} \vee \langle c, a_m \rangle \in \mathbf{hb}$.

- i. $c = a_m$. We have $\langle a_m, b \rangle \in \mathbf{hb}^*$. Then, by transitivity, we have $\langle a, b \rangle \in \mathbf{hb}^*$ which contradicts our initial assumption.

- ii. $\langle a_m, c \rangle \in \mathbf{hb}$. By transitivity, we have $\langle a_m, b \rangle \in \mathbf{hb}^*$, and, thus, $\langle a, b \rangle \in \mathbf{hb}^*$. This contradicts our initial assumption.
- iii. $\langle c, a_m \rangle \in \mathbf{hb} \wedge c \neq a$. By COHERENCE, we have $\langle a_m, c \rangle \notin \mathbf{mo}$ and, thus, $\langle c, a_m \rangle \in \mathbf{mo}$. This contradicts $ts(c) \geq ts(a_m)$.

5. $\gamma = \mathsf{U}^\circ(\ell, _, _) \wedge \exists a \in (\mathbf{RW})_\ell^{\mathbf{na}}. \forall b \in \mathsf{E}^\pi. \langle a, b \rangle \notin \mathbf{hb}^*$.

We show that performing the “write” part of the update event leads to a race in ORC11, *i.e.*, $\neg(\mathcal{M}, \mathcal{N}, \mathcal{V}(\pi) \vdash \text{RaceFree}(\langle \text{Write}, \ell, _, o.w \rangle))$. We consider $a \in \mathbf{R}^{\mathbf{na}}$ and $a \in \mathbf{W}^{\mathbf{na}}$ separately.

(a) $a \in \mathbf{R}^{\mathbf{na}}$. We show that $\mathcal{N}(\ell).\text{nr} \not\sqsubseteq \mathcal{V}(\pi).\text{cur}(\ell).\text{nr}$.

By way of contradiction, assume that $\mathcal{N}(\ell).\text{nr} \sqsubseteq \mathcal{V}(\pi).\text{cur}(\ell).\text{nr}$. We have $a \in \mathcal{N}(\ell).\text{nr}$ and, thus, $a \in \mathcal{V}(\pi).\text{cur}(\ell).\text{nr}$. This implies that there exists $b' \in \mathsf{E}^\pi$ s.t. $\langle a, b' \rangle \in \mathbf{hb}^*$, which contradicts our assumption.

(b) $a \in \mathbf{W}^{\mathbf{na}}$. It suffices to show that $V(\pi).\text{cur}(\ell) < \mathcal{N}(\ell)$.

Let $a_m \in \mathbf{W}_\ell^{\mathbf{na}}$ be the **mo**-maximal non-atomic write event on ℓ (which implies $ts(a_m) \geq ts(a)$). Then $\mathcal{N}(\ell) = ts(a_m)$. It thus suffices to show that $V(\pi).\text{cur}(\ell) < ts(a_m)$. By way of contradiction, assume that $V(\pi).\text{cur}(\ell) \geq ts(a_m)$. Then, there exists $c \in \mathbf{W}_\ell$ and $b \in \mathsf{E}^\pi$ s.t. $\langle c, b \rangle \in \mathbf{hb}^* \wedge ts(c) \geq ts(a_m)$. From $\langle a, a_m \rangle \in \mathbf{mo}^*$, COHERENCE, and G being non-racy we have that $\langle a, a_m \rangle \in \mathbf{hb}^*$. As G is non-racy, we also have $c = a_m \vee \langle a_m, c \rangle \in \mathbf{hb} \vee \langle c, a_m \rangle \in \mathbf{hb}$.

- i. $c = a_m$. We have $\langle a_m, b \rangle \in \mathbf{hb}^*$. Then, by transitivity, we have $\langle a, b \rangle \in \mathbf{hb}^*$ which contradicts our initial assumption.
- ii. $\langle a_m, c \rangle \in \mathbf{hb}$. By transitivity, we have $\langle a_m, b \rangle \in \mathbf{hb}^*$, and, thus, $\langle a, b \rangle \in \mathbf{hb}^*$. This contradicts our initial assumption.
- iii. $\langle c, a_m \rangle \in \mathbf{hb} \wedge c \neq a$. By COHERENCE, we have $\langle a_m, c \rangle \notin \mathbf{mo}$ and, thus, $\langle c, a_m \rangle \in \mathbf{mo}$. This contradicts $ts(c) \geq ts(a_m)$.

□

Lemma 7 Suppose $G \xrightarrow{\gamma_1} \pi_1 \dots \xrightarrow{\gamma_n} \pi_n \perp_{\text{race}}$ and $G \sim (\mathcal{M}_1, \mathcal{V}_1, \mathcal{N}_1)$. Then there exist $0 \leq j < n$, $\varepsilon_1 \dots \varepsilon_{j+1}$, $(\mathcal{M}_2, \mathcal{N}_2, \mathcal{V}_2) G_2 \dots (\mathcal{M}_j, \mathcal{N}_j, \mathcal{V}_j) G_j$ such that $\xrightarrow{\pi_1} [\varepsilon_1] \dots \xrightarrow{\varepsilon_j} \pi_j (\mathcal{M}_{j-1}, \mathcal{N}_j, \mathcal{V}_j) \wedge \neg((\mathcal{M}_j, \mathcal{N}_j, \mathcal{V}_j) \xrightarrow{\varepsilon_{j+1}} \pi_{j+1} _)$.

Proof. Follows from Lemma 5 and Lemma 6. □

Definition 20 (Initial State) We define the initial physical state \mathcal{M}_0 , global race detector state \mathcal{N}_0 as well as an initial thread view V_0 as follows.

$$\begin{aligned} \mathcal{M}_0 &:= \lambda \ell. \lambda t. \begin{cases} (0, \perp) & \text{if } t = 0 \\ \perp & \text{otherwise} \end{cases} \\ \mathcal{N}_0 &:= \lambda \ell. 0 \\ \text{auxview} &:= \lambda \ell. \{ w := 0, aw := \emptyset, nr := \emptyset, ar := \emptyset, \} \\ V_0 &:= \{ \text{rel} := \lambda \ell. \perp, \text{frel} := \perp, \text{cur} := \text{auxview}, \text{acq} := \text{auxview}, \} \end{aligned}$$

Theorem 1 (ORC11: Racy Programs Get Stuck) Suppose e is buggy. Then $(\mathcal{M}_0, \mathcal{N}_0) \mid [0 \mapsto (e, V_0)] \rightarrow^* (\mathcal{M}', \mathcal{N}') \mid \mathcal{TS}'$ such that $\neg((\mathcal{M}', \mathcal{N}') \mid _)$.

Proof. Follows from [Lemma 3](#) and [Lemma 7](#). □

3 Lifetime Logic for Views

The lifetime logic in RustBelt needs to be adapted to a logic like GPFSL where assertions depend on views.

3.1 Proof Rules

Splitting ownership in time. The lifetime logic adds a built-in notion of *lifetimes*, and the notion of “owning P borrowed for lifetime κ ”, written $\&_{\text{full}}^{\kappa} P$.

The rule [LFTL-BEGIN](#) is used to create a new lifetime. At this point, we obtain the token $[\kappa]_1$ which asserts that *we own the lifetime κ* : We know that the lifetime is still running, and we can end it any time by applying the view shift we got. Now, it turns out that we may want multiple parties to be able to witness that κ is ongoing, so we need to be able to split this assertion: $[\kappa]_q$ denotes ownership of the fraction q of κ . Lifetimes can be *intersected* using the \sqcap operator.

We also obtain an update to end the new lifetime again. This makes use of the “update that takes a step”, defined as follows:

$$P \multimap_{\varepsilon_1}^{\varepsilon_2} Q := P \multimap \varepsilon_1 \Rightarrow \varepsilon_2 \triangleright \varepsilon_2 \Rightarrow \varepsilon_1 Q$$

The core operation of the lifetime logic is *borrowing* an assertion P at a given lifetime. Using [LFTL-BORROW](#), P is split into ownership of P during the lifetime κ (the full borrow), and ownership when κ died (a view shift that lets us “inherit” P from κ). In some sense, we are *splitting ownership along the time axis*: The justification for the separating conjunction is the fact that a lifetime is never both ongoing and has already ended at the same time. Thus, the two parts that we split P into can be treated as disjoint resources: They govern the same part of the (logical and physical) state, but they do so at different points in time.

When a lifetime ends, full borrows at that lifetime are not worth anything any more, a fact that is witnessed by [LFTL-BOR-FAKE](#).

Borrowed assertions can still be split and merged, as shown by [LFTL-BOR-SEP](#). To get access to a borrowed assertion, we use [LFTL-BOR-ACC-STRONG](#). The rule is quite a mouthful, so it is worth looking at the following simpler (derived) version:

$$\langle \&_{\text{full}}^{\kappa} P * [\kappa]_q \Leftrightarrow \triangleright P \rangle_{\mathcal{M}_{\text{ift}}} \tag{1}$$

This lets us *open* full borrows ($\&_{\text{full}}^{\kappa} P$) if we can prove that the lifetime is still ongoing, which we do by presenting any fraction of the lifetime token. We obtain $\triangleright P$, but lose access to that token for as long as the full borrow is open, which ensures that we do not end the lifetime while the full borrow is open. Once we re-established $\triangleright P$, we can *close* the full borrow again the get our token back.

The full rule [LFTL-BOR-ACC-STRONG](#) actually lets us close not just with $\triangleright P$, but with any $\triangleright Q$ if we can show that Q entails P through a view shift. Furthermore, that view shift is only actually tun when the lifetime ends, which is witnessed by providing the appropriate token ($[\dagger\kappa]$).

Figure 12: Lifetime logic assertions and proof rules

Notation	Meaning	Timeless	Persistent
$[\kappa]_q$	Fraction q of lifetime token for κ : Witnessing that the lifetime is still ongoing	Yes	No
$[\dagger\kappa]$	Witness confirming that the lifetime κ is dead (<i>i.e.</i> , it has ended)	Yes	Yes
$\&_{\text{full}}^\kappa P$	Ownership of the <i>full borrow</i> of P for κ	No	No
$\&_i^\kappa P$	There is an <i>indexed borrow</i> named i of P for κ	No	Yes
$[\text{Bor} : i]$	Ownership of the indexed borrow i	Yes	No
$\&_{\text{at}}^{\kappa/0} P$	Internal atomic persistent borrow of P for κ	No	Yes

Lifetimes. Lifetimes κ form a cancellable PCM with intersection as the operation (\sqcap) and unit ε .

$$\kappa \sqsubseteq \kappa' := \square \forall q. \langle [\kappa]_q \Leftrightarrow q'. [\kappa']_{q'} \rangle_{\mathcal{N}_{\text{lft}}}$$

Lifetime creation and end.

LFTL-BEGIN $\text{True} \Rightarrow_{\mathcal{N}_{\text{lft}}} \exists \kappa. [\kappa]_1 * \square([\kappa]_1 \multimap_{\emptyset}^{\mathcal{N}_{\text{lft}}} [\dagger\kappa])$	LFTL-TOK-FRACT $[\kappa]_{q+q'} \Leftrightarrow [\kappa]_q * [\kappa]_{q'}$	LFTL-TOK-FRACT-OBJ $[\kappa]_{q+q'} \Rightarrow [\kappa]_q * \langle \text{obj} \rangle [\kappa]_{q'}$
LFTL-TOK-COMP $[\kappa \sqcap \kappa']_q \Leftrightarrow [\kappa]_q * [\kappa']_q$	LFTL-TOK-UNIT $\text{True} \Rightarrow [\varepsilon]_q$	LFTL-NOT-OWN-END $[\kappa]_q * [\dagger\kappa] \Rightarrow \text{False}$
	LFTL-END-COMP $[\dagger\kappa \sqcap \kappa'] \Leftrightarrow [\dagger\kappa] \vee [\dagger\kappa']$	
	LFTL-END-UNIT $[\dagger\varepsilon] \Rightarrow \text{False}$	

Creating full borrows and using them.

LFTL-BORROW $\triangleright P \Rightarrow_{\mathcal{N}_{\text{lft}}} \&_{\text{full}}^\kappa P * ([\dagger\kappa] \multimap_{\mathcal{N}_{\text{lft}}} \triangleright P)$	LFTL-BOR-SEP $\&_{\text{full}}^\kappa (P * Q) \Leftrightarrow_{\mathcal{N}_{\text{lft}}} \&_{\text{full}}^\kappa P * \&_{\text{full}}^\kappa Q$
	LFTL-BOR-FAKE $\langle \text{subj} \rangle [\dagger\kappa] \Rightarrow_{\mathcal{N}_{\text{lft}}} \&_{\text{full}}^\kappa P$
LFTL-BOR-ACC-STRONG $\&_{\text{full}}^\kappa P * [\kappa]_q \Rightarrow_{\mathcal{N}_{\text{lft}}} \exists \kappa'. \kappa \sqsubseteq \kappa' * \triangleright P * \left(\forall Q. \triangleright (\triangleright Q * \langle \text{subj} \rangle [\dagger\kappa'] \multimap_{\emptyset} \triangleright P) * \triangleright Q \multimap_{\mathcal{N}_{\text{lft}}} \&_{\text{full}}^{\kappa'} Q * [\kappa]_q \right)$	
LFTL-BOR-ACC-ATOMIC-STRONG $\&_{\text{full}}^\kappa P \stackrel{\mathcal{N}_{\text{lft}}}{\Rightarrow} \emptyset \left(\exists P' \kappa'. \kappa \sqsubseteq \kappa' * \triangleright ([P'] \wedge P) * \left(\forall Q. \triangleright (\triangleright Q * \langle \text{subj} \rangle [\dagger\kappa'] \multimap_{\emptyset} \triangleright P) * \triangleright ([P'] \wedge Q) \stackrel{\emptyset}{\multimap}_{\mathcal{N}_{\text{lft}}} \&_{\text{full}}^{\kappa'} Q \right) \right) \vee \left(\exists \kappa'. \kappa \sqsubseteq \kappa' * \langle \text{subj} \rangle [\dagger\kappa'] * \stackrel{\emptyset}{\multimap}_{\mathcal{N}_{\text{lft}}} \text{True} \right)$	

Figure 13: Lifetime logic assertions and proof rules, continued

Indexed borrows.

$$\begin{array}{c}
\text{LFTL-BOR-IDX} \\
& \&_{\text{full}}^{\kappa} P \Leftrightarrow \exists i. \&_i^{\kappa} P * [\text{Bor} : i] \\
\\
\text{LFTL-IDX-ACC} \\
& \&_i^{\kappa} P(V_{\text{tok}}) * [\text{Bor} : i](V_{\text{bor}}) * [\kappa]_q(V_{\text{tok}}) \Rightarrow_{\mathcal{N}_{\text{ift}}} \exists V. V \sqsubseteq V_{\text{tok}} \sqcup V_{\text{bor}} * \triangleright P(V) * \\
& \quad \left(\forall V'_{\text{tok}}. V_{\text{tok}} \sqsubseteq V'_{\text{tok}} * \triangleright P(V'_{\text{tok}} \sqcup V) \Rightarrow_{\mathcal{N}_{\text{ift}}} * [\text{Bor} : i](V'_{\text{tok}} \sqcup V) * [\kappa]_q(V'_{\text{tok}}) \right) \\
\\
\text{LFTL-IDX-BOR-UNNEST} \\
& \&_i^{\kappa} P * \&_{\text{full}}^{\kappa'}([\text{Bor} : i]) \Rightarrow_{\mathcal{N}_{\text{ift}}} \&_{\text{full}}^{\kappa \sqcap \kappa'} P \\
\\
\text{LFTL-IDX-SHORTEN} \\
& \frac{\kappa' \sqsubseteq \kappa}{\&_i^{\kappa} P \Rightarrow \&_i^{\kappa'} P} \\
\\
\text{LFTL-IDX-BOR-IFF} \\
& \frac{\triangleright \square(P \Leftrightarrow Q)}{\&_i^{\kappa} P \Rightarrow \&_i^{\kappa} Q}
\end{array}$$

Internal persistent atomic borrows.

$$\begin{array}{c}
\text{LFTL-BOR-IN-AT} \\
& \&_{\text{full}}^{\kappa} P \Rightarrow_{\mathcal{N}_{\text{ift}}} \&_{\text{at}}^{\kappa/0} P \\
\\
\text{LFTL-IN-AT-ACC} \\
& \&_{\text{at}}^{\kappa/0} P \vdash \langle [\kappa]_q \Leftrightarrow V_b. \triangleright [P]_{\sqcup V_b} \rangle_{\mathcal{N}_{\text{ift}}}^{\emptyset} \\
\\
\text{LFTL-IN-AT-IFF} \\
& \frac{\triangleright \square(P \Leftrightarrow Q)}{\&_{\text{at}}^{\kappa/0} P \Rightarrow \&_{\text{at}}^{\kappa/0} Q} \\
\\
\text{LFTL-IN-AT-SHORTEN} \\
& \frac{\kappa' \sqsubseteq \kappa}{\&_{\text{at}}^{\kappa/0} P \Rightarrow \&_{\text{at}}^{\kappa'/0} P}
\end{array}$$

Finally, the rule **LFTL-BOR-ACC-ATOMIC-STRONG** provides a way to access a full borrow *without* having a proof that the lifetime is still ongoing.

A closer look at lifetimes. Before we go on talking about the lifetime logic rules, we have to become more concrete about what a *lifetime* κ is. Lifetimes κ form a partial commutative monoid with unit ε . We will also refer to the composition operation (\sqcap) as *intersection* of lifetimes. Moreover, the PCM has to be *cancellable*, which means that the composition function is injective.

Furthermore, we define the following inclusion relation on lifetimes:

$$\kappa \sqsubseteq \kappa' := \square \left(\forall q. \langle [\kappa]_q \Leftrightarrow q'. [\kappa']_{q'} \rangle_{\mathcal{N}_{\text{ift}}} \right)$$

This says that κ is dynamically shorter than κ' if, given any fraction the token for κ , we can produce some fraction of the token for κ' . It is easy to show that this inclusion interacts as expected with lifetime intersection (**LFTL-INCL-ISECT**).

Indexed borrows. While the proof rules given so far bring us pretty far, it turns out that for some of the advanced reasoning we need to do for Rust, they do not suffice. As we start to build more complicated protocols involving full borrows, the fact that $\&_{\text{full}}^{\kappa} P$ is neither timeless nor persistent really becomes a problem.

For this reason, the logic provides a way to *decompose* a full borrow into timeless and persistent pieces (the borrow token and the indexed borrow, respectively), which are tied together by an *index* i (**LFTL-BOR-IDX**). Indexed borrows can be opened using **LFTL-IDX-ACC**, but they cannot be strengthened, reborrowed or split. Furthermore, indexed borrows can be *shortened* (**LFTL-IDX-SHORTEN**) following the dynamic lifetime inclusion $\kappa' \sqsubseteq \kappa$.

Indexed borrows are used to state the rule **LFTL-IDX-BOR-UNNEST**, which will be used later to prove two important derived rules: unnesting and reborrowing.

Internal atomic persistent borrows. They are a primitive form of atomic persistent borrow (see the paragraph below about atomic persistent borrows). They have the same opening and closing rules as atomic persistent borrows, but use \mathcal{N}_{ift} as namespace, which could not be used with atomic persistent borrows.

Internally, they are implemented in a very similar fashion as atomic persistent borrows. The reason we need them is that they are used for implementing fractured borrows, which are in turn used for creating dynamic lifetime inclusion, and this cannot afford using a different mask as \mathcal{N}_{ift} .

3.2 Derived Forms of Borrowing

Fig. 14 shows some rules that can be derived from the basic rules discussed in the previous subsection.

Furthermore, we introduce in **Fig. 15** some derived forms of borrowing – that is, assertions that share are somewhat like $\&_{\text{full}}^{\kappa} P$, but not exactly.

Reborrowing. Two The rule **LFTL-REBORROW** lets us *reborrow* a $\&_{\text{full}}^{\kappa} P$, which means that we can pick some statically shorter lifetime $\kappa' \sqsubseteq \kappa$ and obtain P borrowed at κ' . When κ' ends, we can get our original full borrow back.

The rule **LFTL-BOR-UNNEST** is related. It deals with the case that we have a full borrow of a full borrow ($\&_{\text{full}}^{\kappa'} \&_{\text{full}}^{\kappa} P$). If we have already opened that full borrow and stripped a way the \triangleright added

$$\begin{array}{c}
\text{LFTL-INCL-ISECT} \\
\frac{\kappa \sqcap \kappa' \sqsubseteq \kappa}{\text{LFTL-REBORROW}} \\
\kappa' \sqsubseteq \kappa \vdash \&_{\text{full}}^{\kappa} P \Rightarrow_{\mathcal{N}_{\text{ift}}} \&_{\text{full}}^{\kappa'} P * ([\dagger \kappa'] \Rightarrow_{\mathcal{N}_{\text{ift}}} *_{\mathcal{N}_{\text{ift}}} \&_{\text{full}}^{\kappa} P) \\
\text{LFTL-BOR-ACC-CONS} \\
\&_{\text{full}}^{\kappa} P * [\kappa]_q \Rightarrow_{\mathcal{N}_{\text{ift}}} \triangleright P * \forall Q. \triangleright (\triangleright Q \Rightarrow_{\mathcal{N}_{\text{ift}}} *_{\emptyset} \triangleright P) * \triangleright Q \Rightarrow_{\mathcal{N}_{\text{ift}}} *_{\mathcal{N}_{\text{ift}}} \&_{\text{full}}^{\kappa} Q * [\kappa]_q \\
\text{LFTL-BOR-ACC} \\
\langle [\kappa]_q * \&_{\text{full}}^{\kappa} P \Leftrightarrow \triangleright P \rangle_{\mathcal{N}_{\text{ift}}} \\
\text{LFTL-INCL-GLB} \\
\frac{\kappa \sqsubseteq \kappa' \quad \kappa \sqsubseteq \kappa''}{\kappa \sqsubseteq \kappa' \sqcap \kappa''} \\
\text{LFTL-FRACT-LINCL} \\
\frac{\&_{\text{frac}}^{\kappa} q'. [\kappa']_{q \cdot q'}}{\&_{\text{full}}^{\kappa} (\&_{\text{full}}^{\kappa'} P) \Rightarrow_{\mathcal{N}_{\text{ift}}} \&_{\text{full}}^{\kappa} P} \\
\text{LFTL-BOR-UNNEST} \\
\frac{\&_{\text{full}}^{\kappa} P \Rightarrow \&_{\text{full}}^{\kappa'} P}{\&_{\text{full}}^{\kappa} (\&_{\text{full}}^{\kappa'} P) \Rightarrow_{\mathcal{N}_{\text{ift}}} \&_{\text{full}}^{\kappa \sqcap \kappa'} P} \\
\text{LFTL-BOR-FREEZE} \\
\frac{\tau \text{ inhabited}}{(\&_{\text{full}}^{\kappa} \exists x : \tau. P) \Rightarrow_{\mathcal{N}_{\text{ift}}} \exists x : \tau. \&_{\text{full}}^{\kappa} P} \\
\text{LFTL-BOR-IFF} \\
\frac{\triangleright \square (P \Leftrightarrow Q)}{\&_{\text{full}}^{\kappa} P \Rightarrow \&_{\text{full}}^{\kappa} Q}
\end{array}$$

Figure 14: Lifetime logic derived rules

by opening, then we can use **LFTL-BOR-UNNEST** to “unnest” the full borrow in the sense that we end up with a full borrow at the intersected lifetime ($\&_{\text{full}}^{\kappa' \sqcap \kappa} P$).

Both of these rules are *derived* from **LFTL-IDX-BOR-UNNEST**.

Persistent borrows. Persistent borrows are a persistent version of borrows. This means that many parties are allowed to get access to its content. In order to avoid reentrant accesses, we can use *two* different mechanisms, giving rise to two flavors of persistent borrows.

Similarly to invariants in Iris, the first possible mechanism is to force only atomic accesses. We then get *atomic persistent borrows*, which are essentially like invariant in Iris with the additional quirk that the invariant is only maintained for the duration of the lifetime of the borrow. They can be defined as follows:

$$\&_{\text{at}}^{\kappa/\mathcal{N}} P := \exists i. \&_i^{\kappa} P * \mathcal{N} \# \mathcal{N}_{\text{ift}} * [\text{Bor} : i]^{\mathcal{N}}$$

The other possible mechanism is to restrict the persistent borrow to be used in a threaded manner, by using the mechanism of *non-atomic invariants* described in the Iris documentation (and can be adapted to the GPFSL logic with the same rules). The persistent borrows of this other flavor are called *non-atomic persistent borrows*. They can be defined by:

$$\&_{\text{na}}^{\kappa/p.\mathcal{N}} P := \exists i. \&_i^{\kappa} P * \text{NaInv}^{p.\mathcal{N}}([\text{Bor} : i])$$

Fractured borrows. A *fractured borrow* is a borrow of a permission $\Phi(q)$ that can be *fractured*, *i.e.*, decomposed according to a fraction:

$$\Phi(q_1 + q_2) \Leftrightarrow \Phi(q_1) * \Phi(q_2)$$

Intuitively, it should be possible to share such a borrow, and still obtain some fraction of Φ via a non-atomic accessor, *i.e.*, $\Phi(q)$ can actually be kept around for non-atomic expressions. This is because even if other threads are concurrently accessing the borrow, they will always leave *some* fraction of Φ in the borrow.

Fractured borrows are particularly interesting for giving rise to dynamic lifetime inclusion (**LFTL-FRACT-LINCL**).

Notation	Meaning	Timeless	Persistent
$\&_{\text{at}}^{\kappa/\mathcal{N}} P$	There is a <i>atomic persistent borrow</i> of P for κ in namespace \mathcal{N}	No	Yes
$\&_{\text{frac}}^{\kappa} \lambda q. P$	There is a <i>fractured borrow</i> of $\lambda q. P$ for κ	No	Yes
$\&_{\text{na}}^{\kappa/p.\mathcal{N}} P$	There is a <i>non-atomic persistent borrow</i> of P for κ in non-atomic invariant pool p , namespace \mathcal{N}	No	Yes

Atomic persistent borrows

$$\begin{array}{c}
\text{LFTL-BOR-AT} \\
\mathcal{N} \# \mathcal{N}_{\text{ift}} \vdash \&_{\text{full}}^{\kappa} P \Rightarrow_{\mathcal{N}_{\text{ift}}} \&_{\text{at}}^{\kappa/\mathcal{N}} P \\
\text{LFTL-AT-SHORTEN} \\
\frac{\kappa' \sqsubseteq \kappa}{\&_{\text{at}}^{\kappa/\mathcal{N}} P \Rightarrow \&_{\text{at}}^{\kappa'/\mathcal{N}} P}
\end{array}
\qquad
\begin{array}{c}
\text{LFTL-AT-ACC} \\
\&_{\text{at}}^{\kappa/\mathcal{N}} P \vdash \langle [\kappa]_q \Leftrightarrow V_b. \triangleright [P]_{\sqcup V_b} \rangle_{\mathcal{N}_{\text{ift}}, \mathcal{N}} \\
\text{LFTL-AT-IFF} \\
\frac{\triangleright \Box (P \Leftrightarrow Q)}{\&_{\text{at}}^{\kappa/\mathcal{N}} P \Rightarrow \&_{\text{at}}^{\kappa/\mathcal{N}} Q}
\end{array}$$

Non-atomic persistent borrows

$$\begin{array}{c}
\text{LFTL-BOR-NA} \\
\&_{\text{full}}^{\kappa} P \Rightarrow_{\mathcal{N}} \&_{\text{na}}^{\kappa/p.\mathcal{N}} P \\
\text{LFTL-NA-SHORTEN} \\
\frac{\kappa' \sqsubseteq \kappa}{\&_{\text{na}}^{\kappa/p.\mathcal{N}} P \Rightarrow \&_{\text{na}}^{\kappa'/p.\mathcal{N}} P}
\end{array}
\qquad
\begin{array}{c}
\text{LFTL-NA-ACC} \\
\&_{\text{na}}^{\kappa/p.\mathcal{N}} P \vdash \langle [\kappa]_q * [\text{Na} : p.\mathcal{N}] \Leftrightarrow \triangleright P \rangle_{\mathcal{N}_{\text{ift}}, \mathcal{N}} \\
\text{LFTL-NA-IFF} \\
\frac{\triangleright \Box (P \Leftrightarrow Q)}{\&_{\text{na}}^{\kappa/p.\mathcal{N}} P \Rightarrow \&_{\text{na}}^{\kappa/p.\mathcal{N}} Q}
\end{array}$$

Fractured borrows

$$\begin{array}{c}
\text{LFTL-BOR-FRACTURE} \\
\frac{\forall q_1, q_2. \Phi(q_1 + q_2) \Leftrightarrow \Phi(q_1) * \Phi(q_2)}{\&_{\text{full}}^{\kappa} \Phi(1) \Rightarrow_{\mathcal{N}_{\text{ift}}} \&_{\text{frac}}^{\kappa} \Phi} \\
\text{LFTL-FRACT-SHORTEN} \\
\frac{\kappa' \sqsubseteq \kappa}{\&_{\text{frac}}^{\kappa} \Phi \Rightarrow \&_{\text{frac}}^{\kappa'} \Phi}
\end{array}
\qquad
\begin{array}{c}
\text{LFTL-FRACT-ACC} \\
\&_{\text{frac}}^{\kappa} \Phi \vdash \langle [\kappa]_q \Leftrightarrow q'. \triangleright \Phi(q') \rangle_{\mathcal{N}_{\text{ift}}} \\
\text{LFTL-FRACT-IFF} \\
\frac{\triangleright \Box (\forall q. \Phi(q) \Leftrightarrow \Psi(q))}{\&_{\text{frac}}^{\kappa} \Phi \Rightarrow \&_{\text{frac}}^{\kappa} \Psi}
\end{array}$$

Figure 15: Lifetime logic derived forms

4 GPFSL

GPFSL is an extension of iGPS (Kaiser et al. [2017]) that adopts the fence modalities from FSL (Doko and Vafeiadis [2016, 2017]). Fig. 16 lists the rules for traditional points-to assertions (non-atomics). Fig. 17 lists the rules for fork and fences.

GPFSL also combines GPS single-location protocols and iGPS single-write protocols with atomic borrows (Fig. 18, Fig. 19, Fig. 21, Fig. 22, Fig. 23, Fig. 24). These protocols are used to verify `Mutex`, `RwLock`.

GPFSL also develops protocols based on *view-dependent cancellable* invariants, which is a simpler variant of the lifetime logic. The protocols differ slightly from the atomic-borrows-based versions. Some of them are given in Fig. 25 and Fig. 26. These protocols are used to verify `Arc<T>`, `thread::spawn`, and `rayon::join`.

$$\begin{array}{c}
\text{NA-FRAC-AGREE} \\
\ell \xrightarrow{q} v * \ell \xrightarrow{q'} v' \Leftrightarrow \ell \xrightarrow{q+q'} v * v = v' \\
\\
\text{NA-FREEABLE-COMBINE} \\
\uparrow_q^m \ell * \uparrow_{q'}^{m'} \ell + m \Leftrightarrow \uparrow_{q+q'}^{m+m'} \ell \\
\\
\text{NA-ALLOC} \\
\{\text{True}\} \mathbf{alloc}(n) \{ \ell. \exists \bar{v}. \ell \mapsto \bar{v} * |\bar{v}| = n * \uparrow_1^n \ell \} \\
\\
\text{NA-FREE} \\
\{ \ell \mapsto \bar{v} * \uparrow_1^{|\bar{v}|} \ell \} \mathbf{free}(|\bar{v}|, v) \{\text{True}\} \\
\\
\text{NA-READ} \\
\{ \ell \xrightarrow{q} v \} * \ell \{ v'. v' = v * \ell \xrightarrow{q} v \} \\
\\
\text{NA-WRITE} \\
\{ \ell \mapsto v \} \ell := w \{ \ell \mapsto w \} \\
\\
\text{NA-MEMCPY} \\
\frac{|\bar{v}_1| = |\bar{v}_2| = n}{\{ \ell_1 \mapsto \bar{v}_1 * \ell_2 \xrightarrow{q} \bar{v}_2 \} \ell_1 :=_n * \ell_2 \{ \ell_1 \mapsto \bar{v}_2 * \ell_2 \xrightarrow{q} \bar{v}_2 \}}
\end{array}$$

Figure 16: Non-atomics rules.

$$\begin{array}{c}
\text{FORK} \\
\frac{\forall \rho. \{P\} e \text{ in } \rho \{\text{True}\}}{\{P\} \mathbf{fork} \{e\} \{\text{True}\}} \\
\\
\text{REL-FENCE} \\
\{P\} \mathbf{fence}_{\text{rel}} \text{ in } \pi \{ \Delta_\pi P \} \\
\\
\text{ACQ-FENCE} \\
\{ \nabla_\pi P \} \mathbf{fence}_{\text{acq}} \text{ in } \pi \{P\}
\end{array}$$

Figure 17: Fork and fences rules.

$$\begin{array}{c} \text{ATBOR-N-PERSISTENT} \\ \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}} \Rightarrow \square \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}} \end{array} \qquad \begin{array}{c} \text{ATBOR-N-LOCAL} \\ \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}} \Rightarrow \mathcal{R}(\ell, t, s, v, \mathcal{I}) \end{array}$$

$$\begin{array}{c} \text{ATBOR-N-LOCAL-JOIN} \\ \mathcal{R}(\ell, t, s, v, \mathcal{I}) * \&^\kappa \boxed{\ell : (t', s', v') \mid \mathcal{I}} \Rightarrow \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}} \end{array}$$

$$\begin{array}{c} \text{ATBOR-N-INIT} \\ [\kappa]_q * \&^\kappa_{\text{full}} (\exists v. \ell \mapsto v * P(v)) * (\forall t, v. \triangleright P(v) * \triangleright \mathcal{I}_w(\ell, t, s, v)) * \\ (\square \forall t, s, v. \triangleright \mathcal{I}_w(\ell, t, s, v) \Rightarrow \triangleright P(v)) \Rightarrow * \\ [\kappa]_q * \exists t, v. \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}} \end{array}$$

$$\begin{array}{c} \text{ATBOR-N-RLX-READ} \\ \forall t' \sqsupseteq t, s' \sqsupseteq s, v'. \mathcal{I}_r(\ell, t', s', v') \Rightarrow \mathcal{I}_r(\ell, t', s', v') * P(t', s', v') \\ \forall t' \sqsupseteq t, s' \sqsupseteq s, v'. \mathcal{I}_w(\ell, t', s', v') \Rightarrow \mathcal{I}_w(\ell, t', s', v') * P(t', s', v') \\ \hline \{[\kappa]_q * \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}\}^{*\text{rlx}} \ell \text{ in } \pi \left\{ v'. [\kappa]_q * \exists t' \sqsupseteq t, s' \sqsupseteq s. \&^\kappa \boxed{\ell : (t', s', v') \mid \mathcal{I}} * \nabla_\pi P(t', s', v') \right\} \end{array}$$

$$\begin{array}{c} \text{ATBOR-N-ACQ-READ} \\ \forall t', s', v'. \mathcal{I}_r(\ell, t', s', v') \Rightarrow \mathcal{I}_r(\ell, t', s', v') * P(t', s', v') \\ \forall t', s', v'. \mathcal{I}_w(\ell, t', s', v') \Rightarrow \mathcal{I}_w(\ell, t', s', v') * P(t', s', v') \\ \hline \{[\kappa]_q * \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}\}^{*\text{acq}} \ell \left\{ v'. [\kappa]_q * \exists t' \sqsupseteq t, s' \sqsupseteq s. \&^\kappa \boxed{\ell : (t', s', v') \mid \mathcal{I}} * P(t', s', v') \right\} \end{array}$$

$$\begin{array}{c} \text{ATBOR-N-RLX-WRITE} \\ \{[\kappa]_q * \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}} * \Delta_\pi (\forall t' > t. \mathcal{R}(\ell, t', s', v', \mathcal{I}) \Rightarrow \mathcal{I}_w(\ell, t', s', v'))\} \\ \ell :=_{\text{rlx}} v' \text{ in } \pi \\ \{[\kappa]_q * \&^\kappa \boxed{\ell : (t', s', v') \mid \mathcal{I}}\} \end{array}$$

$$\begin{array}{c} \text{ATBOR-N-REL-WRITE} \\ \{[\kappa]_q * \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}} * (\forall t' > t. \mathcal{R}(\ell, t', s', v', \mathcal{I}) \Rightarrow \mathcal{I}_w(\ell, t', s', v'))\} \\ \ell :=_{\text{rel}} v' \text{ in } \pi \\ \{[\kappa]_q * \&^\kappa \boxed{\ell : (t', s', v') \mid \mathcal{I}}\} \end{array}$$

Figure 18: Atomic-borrow-based normal GPFSL protocols.

$$\Delta_{\pi}^{?o_w} P := (\text{if } o_w = \mathbf{rel} \text{ then } P \text{ else } \Delta_{\pi} P)$$

$$\nabla_{\pi}^{?o_r} P := (\text{if } o_r = \mathbf{acq} \text{ then } P \text{ else } \nabla_{\pi} P)$$

ATBOR-N-CAS

$$\frac{\begin{array}{c} o_f, o_r \in \{\mathbf{rlx}, \mathbf{acq}\} \quad o_w \in \{\mathbf{rlx}, \mathbf{rel}\} \\ \forall t' \sqsupseteq t, s' \sqsupseteq s, v'. \mathcal{I}_w(\ell, t', s', v') \vee \mathcal{I}_r(\ell, t', s', v') \Rightarrow (\vdash v_r =^? v') \\ \forall t' \sqsupseteq t, s' \sqsupseteq s, v'. (\vdash v' \neq v_r) \Rightarrow \mathcal{I}_r(\ell, t', s', v') \Rightarrow \mathcal{I}_r(\ell, t', s', v') * R(t', s', v') \\ \forall t' \sqsupseteq t, s' \sqsupseteq s, v'. (\vdash v' \neq v_r) \Rightarrow \mathcal{I}_w(\ell, t', s', v') \Rightarrow \mathcal{I}_w(\ell, t', s', v') * R(t', s', v') \\ \forall t' \sqsupseteq t, s' \sqsupseteq s. \triangleright \mathcal{I}_w(\ell, t', s', v_r) \Rightarrow \triangleright Q_1(t', s') * \triangleright Q_2(t', s') \\ \forall t' \sqsupseteq t, s' \sqsupseteq s. P * \triangleright Q_2(t', s') \Rightarrow \exists s'' \sqsupseteq s'. \forall t'' > t. \triangleright \mathcal{R}(\ell, t'', s'', v_w, \mathcal{I}) \\ \Rightarrow (\langle \text{obj} \rangle \triangleright Q_1(t', s') \Rightarrow \triangleright \mathcal{I}_m(\ell, t', s', v_r)) * \boxtimes (Q(t'', s'') * \mathcal{I}_w(\ell, t'', s'', v_w)) \end{array}}{\left\{ [\kappa]_q * \&^{\kappa} \boxed{\ell : (t, s, v)} \mathcal{I} * \Delta_{\pi}^{?o_w} P \right\}} \\ \mathbf{CAS}(\ell, v_r, v_w, o_f, o_r, o_w) \text{ in } \pi \\ \left\{ \begin{array}{l} b = 1 * \exists t' > t. \&^{\kappa} \boxed{\ell : (t', s', v_w)} \mathcal{I} * \nabla_{\pi}^{?o_r} Q(t'', s'') \\ b. [\kappa]_q * \exists s' \sqsupseteq s. \\ \vee b = 0 * \Delta_{\pi}^{?o_w} P * \exists t' \geq t, v'. (\vdash v' \neq v_r) * \&^{\kappa} \boxed{\ell : (t', s', v')} \mathcal{I} * \nabla_{\pi}^{?o_f} R(t', s', v') \end{array} \right\}$$

Figure 19: CAS rule for atomic-borrow-based normal GPFSL protocols.

<p>SW-WRITER-LOCAL-EXCLUSIVE</p> $\mathcal{W}(\ell, t, s, v, \mathcal{I}) * \mathcal{W}(\ell, t, s, v, \mathcal{I}) \Rightarrow \mathbf{False}$	<p>SW-LOCAL-WRITER-READER</p> $\mathcal{W}(\ell, t, s, v, \mathcal{I}) \Rightarrow \mathcal{R}(\ell, t, s, v, \mathcal{I})$
<p>SW-CREADERS-LOCAL-JOIN</p> $\mathcal{R}_{\text{shr}}^q(\ell, t, s, v, \mathcal{I}) * \mathcal{R}_{\text{shr}}^{q'}(\ell, t', s', v', \mathcal{I}) \Rightarrow \mathcal{R}_{\text{shr}}^{q+q'}(\ell, t, s, v, \mathcal{I})$	
<p>SW-CREADERS-LOCAL-SPLIT</p> $\mathcal{R}_{\text{shr}}^{q+q'}(\ell, t, s, v, \mathcal{I}) \Rightarrow \mathcal{R}_{\text{shr}}^q(\ell, t, s, v, \mathcal{I}) * \mathcal{R}_{\text{shr}}^{q'}(\ell, t, s, v, \mathcal{I})$	
<p>SW-CWRITER-LOCAL-EXCLUSIVE</p> $\mathcal{W}_{\text{shr}}(\ell, t, s, v, \mathcal{I}) * \mathcal{W}_{\text{shr}}(\ell, t', s', v', \mathcal{I}) \Rightarrow \mathbf{False}$	
<p>SW-SHARE-LOCAL-CWRITER</p> $\mathcal{W}(\ell, t, s, v, \mathcal{I}) \Rightarrow \mathcal{W}_{\text{shr}}(\ell, t, s, v, \mathcal{I}) * \mathcal{R}_{\text{shr}}^1(\ell, t, s, v, \mathcal{I})$	<p>SW-READER-CREADER-LOCAL</p> $\mathcal{R}_{\text{shr}}^q(\ell, t, s, v, \mathcal{I}) \Rightarrow \mathcal{R}(\ell, t, s, v, \mathcal{I})$
<p>SW-CW-LOCAL-EXCLUSIVE</p> $\mathcal{W}(\ell, t, s, v, \mathcal{I}) * \mathcal{W}_{\text{shr}}(\ell, t', s', v', \mathcal{I}) \Rightarrow \mathbf{False}$	<p>SW-CR-LOCAL-EXCLUSIVE</p> $\mathcal{W}(\ell, t, s, v, \mathcal{I}) * \mathcal{R}_{\text{shr}}^q(\ell, t', s', v', \mathcal{I}) \Rightarrow \mathbf{False}$

Figure 20: Local assertions of single-writer GPFSL protocols.

$$\begin{array}{l}
\text{ATBOR-SW-READER-PERSISTENT} \\
& \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_R \Rightarrow \square \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_R \\
& \text{ATBOR-SW-READER-LOCAL} \\
& \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_R \Rightarrow \mathcal{R}(\ell, t, s, v, \mathcal{I}) \\
& \text{ATBOR-SW-READER-LOCAL-JOIN} \\
& \mathcal{R}(\ell, t, s, v, \mathcal{I}) * \&^\kappa \boxed{\ell : (t', s', v') \mid \mathcal{I}}_R \Rightarrow \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_R \\
& \text{ATBOR-SW-WRITER-LOCAL} \\
& \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_W \Rightarrow \mathcal{W}(\ell, t, s, v, \mathcal{I}) \\
& \text{ATBOR-SW-WRITER-LOCAL-JOIN} \\
& \mathcal{W}(\ell, t, s, v, \mathcal{I}) * \&^\kappa \boxed{\ell : (t', s', v') \mid \mathcal{I}}_R \Rightarrow \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_W \\
& \text{ATBOR-SW-CREADER-LOCAL} \\
& \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_{CR}^q \Rightarrow \mathcal{R}_{\text{shr}}^q(\ell, t, s, v, \mathcal{I}) \\
& \text{ATBOR-SW-CREADER-LOCAL-JOIN} \\
& \mathcal{R}_{\text{shr}}^q(\ell, t, s, v, \mathcal{I}) * \&^\kappa \boxed{\ell : (t', s', v') \mid \mathcal{I}}_R \Rightarrow \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_{CR}^q \\
& \text{ATBOR-SW-CWRITER-LOCAL} \\
& \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_{CW} \Rightarrow \mathcal{W}_{\text{shr}}(\ell, t, s, v, \mathcal{I}) \\
& \text{ATBOR-SW-CWRITER-LOCAL-JOIN} \\
& \mathcal{W}_{\text{shr}}(\ell, t, s, v, \mathcal{I}) * \&^\kappa \boxed{\ell : (t', s', v') \mid \mathcal{I}}_R \Rightarrow \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_{CW} \\
& \text{ATBOR-SW-UNSHARE-LOCAL-CWRITER} \\
& [\kappa]_q * \mathcal{W}_{\text{shr}}(\ell, t, s, v, \mathcal{I}) * \&^\kappa \boxed{\ell : (t', s', v') \mid \mathcal{I}}_{CR}^1 \Rightarrow [\kappa]_q * \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_W
\end{array}$$

Figure 21: Atomic-borrow-based single-writer GPFSL protocols (1).

$$\begin{array}{l}
\text{ATBOR-SW-INIT} \\
[\kappa]_q * \&_{\text{full}}^\kappa (\exists v. \ell \mapsto v * P(v)) \multimap (\forall t, v. \triangleright P(v) \multimap \mathcal{W}(\ell, t, s, v, \mathcal{I}) \Rightarrow \triangleright \mathcal{I}_w(\ell, t, s, v) * Q(t, v)) \multimap \\
(\square \forall t, s, v. \triangleright \mathcal{I}_w(\ell, t, s, v) \Rightarrow \triangleright P(v)) \equiv \star \\
[\kappa]_q * \exists t, v. \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_R * Q(t, v)
\end{array}$$

$$\begin{array}{l}
\text{ATBOR-SW-READ} \\
o \in \{\mathbf{rlx}, \mathbf{acq}\} \\
\forall t' \sqsupseteq t, s' \sqsupseteq s, v'. \mathcal{I}_r(\ell, t', s', v') \Rightarrow \mathcal{I}_r(\ell, t', s', v') * P(t', s', v') \\
\forall t' \sqsupseteq t, s' \sqsupseteq s, v'. \mathcal{I}_w(\ell, t', s', v') \Rightarrow \mathcal{I}_w(\ell, t', s', v') * P(t', s', v') \\
\forall t' \sqsupseteq t, s' \sqsupseteq s, v'. \mathcal{I}_m(\ell, t', s', v') \Rightarrow \mathcal{I}_m(\ell, t', s', v') * P(t', s', v') \\
\hline
\{[\kappa]_q * \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_R\} \\
\text{}^{*o}\ell \text{ in } \pi \\
\{v'. [\kappa]_q * \exists t' \sqsupseteq t, s' \sqsupseteq s. \&^\kappa \boxed{\ell : (t', s', v') \mid \mathcal{I}}_R * \nabla_\pi^{?o} P(t', s', v')\}
\end{array}$$

$$\begin{array}{l}
\text{ATBOR-SW-EXCLUSIVE-READ} \\
o \in \{\mathbf{rlx}, \mathbf{acq}\} \quad \mathcal{I}_w(\ell, t, s, v) \Rightarrow \mathcal{I}_w(\ell, t, s, v) * P \\
\hline
\{[\kappa]_q * \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_W * \text{}^{*o}\ell \text{ in } \pi \{v. [\kappa]_q * \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_W * \nabla_\pi^{?o} P\}\}
\end{array}$$

$$\begin{array}{l}
\text{ATBOR-SW-CREADER-READ} \\
o \in \{\mathbf{rlx}, \mathbf{acq}\} \\
\forall t' \sqsupseteq t, s' \sqsupseteq s, v'. \mathcal{I}_r(\ell, t', s', v') \Rightarrow \mathcal{I}_r(\ell, t, s, v) * P(t', s', v') \\
\forall t' \sqsupseteq t, s' \sqsupseteq s, v'. \mathcal{I}_w(\ell, t', s', v') \Rightarrow \mathcal{I}_w(\ell, t', s', v') * P(t', s', v') \\
\hline
\{[\kappa]_{q_0} * \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_{CR}^q\} \\
\text{}^{*o}\ell \text{ in } \pi \\
\{v'. [\kappa]_{q_0} * \exists t' \sqsupseteq t, s' \sqsupseteq s. \&^\kappa \boxed{\ell : (t', s', v') \mid \mathcal{I}}_{CR}^q * \nabla_\pi^{?o} P(t', s', v')\}
\end{array}$$

Figure 22: Atomic-borrow-based single-writer GPFSL protocols (2).

$$\begin{array}{c}
\text{ATBOR-SW-WRITE} \\
o \in \{\mathbf{rlx}, \mathbf{rel}\} \quad s \sqsubseteq s' \quad \triangleright \langle \mathit{obj} \rangle (\mathcal{I}_w(\ell, t, s, v) \Rightarrow \mathcal{I}_m(\ell, t, s, v) * Q) \\
\hline
\left\{ [\kappa]_q * \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_W * \Delta_{\pi}^{\circ} (\forall t' > t. \mathcal{R}(\ell, t', s', v', \mathcal{I}) \Rightarrow \mathcal{I}_w(\ell, t', s', v')) \right\} \\
\ell :=_o v' \text{ in } \pi \\
\left\{ [\kappa]_q * \&^\kappa \boxed{\ell : (t', s', v') \mid \mathcal{I}}_W * Q \right\}
\end{array}$$

$$\begin{array}{c}
\text{ATBOR-SW-REL-WRITE} \\
s \sqsubseteq s' \quad \triangleright \langle \mathit{obj} \rangle (\mathcal{I}_w(\ell, t, s, v) \Rightarrow Q_1 * Q_2) \\
\hline
\left\{ [\kappa]_q * \&^\kappa \boxed{\ell : (t, s, v) \mid \mathcal{I}}_W * \right. \\
\left. \triangleright (\forall t' > t. \mathcal{W}(\ell, t', s', v', \mathcal{I}) \rightarrow Q_2 \Rightarrow * ((\mathit{obj}) (Q_1 \Rightarrow * \mathcal{I}_m(\ell, t, s, v)) * \Rightarrow (\mathcal{I}_w(\ell, t', s', v') * Q(t')))) \right\} \\
\ell :=_{\mathbf{rel}} v' \\
\left\{ [\kappa]_q * \exists t' > t. \&^\kappa \boxed{\ell : (t', s', v') \mid \mathcal{I}}_R * Q(t') \right\}
\end{array}$$

Figure 23: Atomic-borrow-based single-writer GPFSL protocols (3).

$b ? P := \text{if } b \text{ then } P \text{ else True} \quad b ? P : Q := \text{if } b \text{ then } P \text{ else } Q$

ATBOR-SW-CREADER-CAS

$$\begin{array}{c}
o_f, o_r \in \{\mathbf{rlx}, \mathbf{acq}\} \quad o_w \in \{\mathbf{rlx}, \mathbf{rel}\} \\
\forall t' \sqsupseteq t, s' \sqsupseteq s, v'. \mathcal{I}_w(\ell, t', s', v') \vee \mathcal{I}_r(\ell, t', s', v') \Rightarrow (\vdash v_r =? v') \\
\forall t' \sqsupseteq t, s' \sqsupseteq s, v'. (\vdash v' \neq v_r) \Rightarrow \mathcal{I}_r(\ell, t', s', v') \Rightarrow \mathcal{I}_r(\ell, t', s', v') * R(t', s', v') \\
\forall t' \sqsupseteq t, s' \sqsupseteq s, v'. (\vdash v' \neq v_r) \Rightarrow \mathcal{I}_w(\ell, t', s', v') \Rightarrow \mathcal{I}_w(\ell, t', s', v') * R(t', s', v') \\
\forall t' \sqsupseteq t, s' \sqsupseteq s. \triangleright \mathcal{I}_w(\ell, t', s', v_r) \Rightarrow \triangleright Q_1(t', s') * \triangleright Q_2(t', s') \\
\left(\begin{array}{l}
\forall t' \sqsupseteq t, s' \sqsupseteq s. P * \triangleright Q_2(t', s') \Rightarrow \triangleright \mathcal{W}_{\text{shr}}(\ell, t', s', v_r) * \exists s'' \sqsupseteq s'. \\
\forall t'' > t. \triangleright \mathcal{W}_{\text{shr}}(\ell, t'', s'', v_w, \mathcal{I}) * b_{\text{drop}} ? \triangleright \mathcal{R}_{\text{shr}}^q(\ell, t'', s'', v_w, \mathcal{I}) \\
\Rightarrow (\langle \text{obj} \rangle (\triangleright Q_1(t', s') \Rightarrow \triangleright \mathcal{I}_m(\ell, t', s', v_r))) * \boxtimes (Q(t'', s'') * \mathcal{I}_w(\ell, t'', s'', v_w))
\end{array} \right) \\
\hline
\{[\kappa]_{q_0} * \&^\kappa \boxed{\ell : (t, s, v)} \boxed{\mathcal{I}}_{CR}^q * \Delta_\pi^{?o_w} P\} \\
\mathbf{CAS}(\ell, v_r, v_w, o_f, o_r, o_w) \text{ in } \pi \\
\left(\begin{array}{l}
b = 1 * \exists t' > t. \left(b_{\text{drop}} ? \&^\kappa \boxed{\ell : (t', s', v_w)} \boxed{\mathcal{I}}_R : \&^\kappa \boxed{\ell : (t', s', v_w)} \boxed{\mathcal{I}}_{CR}^q \right) * \\
\quad \nabla_\pi^{?o_r} Q(t', s') \\
b. [\kappa]_{q_0} * \exists s' \sqsupseteq s. \\
\quad \vee b = 0 * \Delta_\pi^{?o_w} P * \exists t' \geq t, v'. (\vdash v' \neq v_r) * \&^\kappa \boxed{\ell : (t', s', v')} \boxed{\mathcal{I}}_{CR}^q * \\
\quad \nabla_\pi^{?o_f} R(t', s', v')
\end{array} \right)
\end{array}$$

ATBOR-SW-READER-CAS

$$\begin{array}{c}
o_f, o_r \in \{\mathbf{rlx}, \mathbf{acq}\} \quad o_w \in \{\mathbf{rlx}, \mathbf{rel}\} \\
\forall t' \sqsupseteq t, s' \sqsupseteq s, v'. \mathcal{I}_w(\ell, t', s', v') \vee \mathcal{I}_r(\ell, t', s', v') \vee \mathcal{I}_m(\ell, t', s', v') \Rightarrow (\vdash v_r =? v') \\
\forall t' \sqsupseteq t, s' \sqsupseteq s, v'. (\vdash v' \neq v_r) \Rightarrow \mathcal{I}_r(\ell, t', s', v') \Rightarrow \mathcal{I}_r(\ell, t', s', v') * R(t', s', v') \\
\forall t' \sqsupseteq t, s' \sqsupseteq s, v'. (\vdash v' \neq v_r) \Rightarrow \mathcal{I}_w(\ell, t', s', v') \Rightarrow \mathcal{I}_w(\ell, t', s', v') * R(t', s', v') \\
\forall t' \sqsupseteq t, s' \sqsupseteq s, v'. (\vdash v' \neq v_r) \Rightarrow \mathcal{I}_m(\ell, t', s', v') \Rightarrow \mathcal{I}_m(\ell, t', s', v') * R(t', s', v') \\
\forall t' \sqsupseteq t, s' \sqsupseteq s. \triangleright (P * \mathcal{I}_m(\ell, t', s', v_r) \Rightarrow \text{False}) \\
\forall t' \sqsupseteq t, s' \sqsupseteq s. \triangleright \mathcal{I}_w(\ell, t', s', v_r) \Rightarrow \triangleright Q_1(t', s') * \triangleright Q_2(t', s') \\
\left(\begin{array}{l}
\forall t' \sqsupseteq t, s' \sqsupseteq s. P * \triangleright Q_2(t', s') \Rightarrow \triangleright \mathcal{W}_{\text{shr}}(\ell, t', s', v_r) * \exists s'' \sqsupseteq s'. \forall t'' > t. \triangleright \mathcal{W}_{\text{shr}}(\ell, t'', s'', v_w, \mathcal{I}) \\
\Rightarrow (\langle \text{obj} \rangle (\triangleright Q_1(t', s') \Rightarrow \triangleright \mathcal{I}_m(\ell, t', s', v_r))) * \boxtimes (Q(t'', s'') * \mathcal{I}_w(\ell, t'', s'', v_w))
\end{array} \right) \\
\hline
\{[\kappa]_q * \&^\kappa \boxed{\ell : (t, s, v)} \boxed{\mathcal{I}}_R * \Delta_\pi^{?o_w} P\} \\
\mathbf{CAS}(\ell, v_r, v_w, o_f, o_r, o_w) \text{ in } \pi \\
\left(\begin{array}{l}
b = 1 * \exists t' > t. \&^\kappa \boxed{\ell : (t', s', v_w)} \boxed{\mathcal{I}}_R * \nabla_\pi^{?o_r} Q(t', s') \\
b. [\kappa]_q * \exists s' \sqsupseteq s. \vee b = 0 * \Delta_\pi^{?o_w} P * \exists t' \geq t, v'. (\vdash v' \neq v_r) * \&^\kappa \boxed{\ell : (t', s', v')} \boxed{\mathcal{I}}_R * \\
\quad \nabla_\pi^{?o_f} R(t', s', v')
\end{array} \right)
\end{array}$$

Figure 24: Atomic-borrow-based single-writer GPFSL protocols (4).

$$\begin{array}{c}
\text{VIEWINV-SW-READER-PERSISTENT} \\
{}^t[\bar{\ell} : (\bar{t}, \bar{s}, \bar{v}) | \bar{\mathcal{I}}]_R \Rightarrow \square {}^t[\bar{\ell} : (\bar{t}, \bar{s}, \bar{v}) | \bar{\mathcal{I}}]_R \\
\text{VIEWINV-SW-READER-LOCAL} \\
{}^t[\bar{\ell} : (\bar{t}, \bar{s}, \bar{v}) | \bar{\mathcal{I}}]_R \Rightarrow \mathcal{R}(\ell, t, s, v, \mathcal{I}) \\
\text{VIEWINV-SW-READER-LOCAL-JOIN} \\
\mathcal{R}(\ell, t, s, v, \mathcal{I}) * {}^t[\bar{\ell} : (\bar{t}', \bar{s}', \bar{v}') | \bar{\mathcal{I}}]_R \Rightarrow {}^t[\bar{\ell} : (\bar{t}, \bar{s}, \bar{v}) | \bar{\mathcal{I}}]_R \\
\text{VIEWINV-SW-WRITER-LOCAL} \\
{}^t[\bar{\ell} : (\bar{t}, \bar{s}, \bar{v}) | \bar{\mathcal{I}}]_{\mathcal{W}} \Rightarrow \mathcal{W}(\ell, t, s, v, \mathcal{I}) \\
\text{VIEWINV-SW-WRITER-LOCAL-JOIN} \\
\mathcal{W}(\ell, t, s, v, \mathcal{I}) * {}^t[\bar{\ell} : (\bar{t}', \bar{s}', \bar{v}') | \bar{\mathcal{I}}]_R \Rightarrow {}^t[\bar{\ell} : (\bar{t}, \bar{s}, \bar{v}) | \bar{\mathcal{I}}]_{\mathcal{W}} \\
\text{VIEWINV-SW-CREADER-LOCAL} \\
{}^t[\bar{\ell} : (\bar{t}, \bar{s}, \bar{v}) | \bar{\mathcal{I}}]_{CR}^q \Rightarrow \mathcal{R}_{\text{shr}}^q(\ell, t, s, v, \mathcal{I}) \\
\text{VIEWINV-SW-CREADER-LOCAL-JOIN} \\
\mathcal{R}_{\text{shr}}^q(\ell, t, s, v, \mathcal{I}) * {}^t[\bar{\ell} : (\bar{t}', \bar{s}', \bar{v}') | \bar{\mathcal{I}}]_R \Rightarrow {}^t[\bar{\ell} : (\bar{t}, \bar{s}, \bar{v}) | \bar{\mathcal{I}}]_{CR}^q \\
\text{VIEWINV-SW-CWRITER-LOCAL} \\
{}^t[\bar{\ell} : (\bar{t}, \bar{s}, \bar{v}) | \bar{\mathcal{I}}]_{CW} \Rightarrow \mathcal{W}_{\text{shr}}(\ell, t, s, v, \mathcal{I}) \\
\text{VIEWINV-SW-CWRITER-LOCAL-JOIN} \\
\mathcal{W}_{\text{shr}}(\ell, t, s, v, \mathcal{I}) * {}^t[\bar{\ell} : (\bar{t}', \bar{s}', \bar{v}') | \bar{\mathcal{I}}]_R \Rightarrow {}^t[\bar{\ell} : (\bar{t}, \bar{s}, \bar{v}) | \bar{\mathcal{I}}]_{CW} \\
\text{VIEWINV-SW-UNSHARE-LOCAL-CWRITER} \\
[l]_q * \mathcal{W}_{\text{shr}}(\ell, t, s, v, \mathcal{I}) * {}^t[\bar{\ell} : (\bar{t}', \bar{s}', \bar{v}') | \bar{\mathcal{I}}]_{CR}^1 \Rightarrow [l]_q * {}^t[\bar{\ell} : (\bar{t}, \bar{s}, \bar{v}) | \bar{\mathcal{I}}]_{\mathcal{W}}
\end{array}$$

Figure 25: View-invariant-based single-writer GPFSL protocols (1).

$$\begin{array}{c}
\text{VIEWINV-SW-INIT} \\
\ell \mapsto v \text{ } * (\forall t, t', v. \triangleright \mathcal{I}_w(\ell, t, s, v)) \equiv * \exists t, t', v. [l]_1 * {}^t[\bar{\ell} : (\bar{t}, \bar{s}, \bar{v}) | \bar{\mathcal{I}}]_{\mathcal{W}} \\
\text{VIEWINV-SW-REL-WRITE} \\
\frac{s \sqsubseteq s' \quad \triangleright \langle \text{obj} \rangle (\mathcal{I}_w(\ell, t, s, v) \Rightarrow Q_1 * Q_2)}{\left\{ \begin{array}{l} [l]_q * {}^t[\bar{\ell} : (\bar{t}, \bar{s}, \bar{v}) | \bar{\mathcal{I}}]_{\mathcal{W}} * \\ \triangleright \left(\forall t' > t. \mathcal{W}(\ell, t', s', v', \mathcal{I}) \text{ } * Q_2 \text{ } * [l]_q \equiv * \left(\langle \text{obj} \rangle (Q_1 \equiv * \mathcal{I}_m(\ell, t, s, v)) * \models (\mathcal{I}_w(\ell, t', s', v') * Q(t')) \right) \right) \\ \ell :=_{\text{rel}} v' \\ \{ \exists t' > t. {}^t[\bar{\ell} : (\bar{t}', \bar{s}', \bar{v}') | \bar{\mathcal{I}}]_R * Q(t') \} \end{array} \right\}}
\end{array}$$

Figure 26: View-invariant-based single-writer GPFSL protocols (2).

$$\mathcal{R}(\ell, t, s, v, \mathcal{I}) \Rightarrow \Box \mathcal{R}(\ell, t, s, v, \mathcal{I})$$

$$\frac{\begin{array}{l} \langle obj \rangle \forall v', t' \geq t, s' \sqsupseteq s. \mathcal{I}_r(\ell, t', s', v') \Rightarrow \mathcal{I}_r(\ell, t', s', v') * P(v') \\ \langle obj \rangle \forall v', t' \geq t, s' \sqsupseteq s. \mathcal{I}_w(\ell, t', s', v') \Rightarrow \mathcal{I}_w(\ell, t', s', v') * P(v') \\ \langle obj \rangle \forall v', t' \geq t, s' \sqsupseteq s. \mathcal{I}_m(\ell, t', s', v') \Rightarrow \mathcal{I}_m(\ell, t', s', v') * P(v') \end{array}}{\{\mathcal{R}(\ell, t, s, v, \mathcal{I}) * [\triangleright \mathcal{GS}(\ell, \mathcal{I})]_V\}} \\ \text{*rlx } \ell \text{ in } \pi \\ \{v'. \nabla_\pi P(v') * t \leq t' * s \sqsubseteq s' * \mathcal{R}(\ell, t', s', v', \mathcal{I}_r) * [\triangleright \mathcal{GS}(\ell, \mathcal{I})]_V\}$$

$$\frac{\begin{array}{l} \langle obj \rangle \forall v', t' \geq t, s' \sqsupseteq s. \mathcal{I}_r(\ell, t', s', v') \Rightarrow \mathcal{I}_r(\ell, t', s', v') * P(v') \\ \langle obj \rangle \forall v', t' \geq t, s' \sqsupseteq s. \mathcal{I}_w(\ell, t', s', v') \Rightarrow \mathcal{I}_w(\ell, t', s', v') * P(v') \\ \langle obj \rangle \forall v', t' \geq t, s' \sqsupseteq s. \mathcal{I}_m(\ell, t', s', v') \Rightarrow \mathcal{I}_m(\ell, t', s', v') * P(v') \end{array}}{\{\mathcal{R}(\ell, t, s, v, \mathcal{I}) * [\triangleright \mathcal{GS}(\ell, \mathcal{I})]_V\}} \\ \text{*acq } \ell \\ \{v'. P(v') * t \leq t' * s \sqsubseteq s' * \mathcal{R}(\ell, t', s', v', \mathcal{I}_r) * [\triangleright \mathcal{GS}(\ell, \mathcal{I})]_V\}$$

$$\begin{array}{l} \mathcal{W}(\ell, \mathcal{I}) * \mathcal{W}(\ell, \mathcal{I}) \Rightarrow \text{False} \quad \{\mathcal{W}(\ell)\} \ell :=_{\text{rlx}} w \{\text{True}\} \quad \{\mathcal{W}(\ell)\} \ell :=_{\text{rel}} w \{\text{True}\} \\ \{\mathcal{R}(\ell)\} \mathbf{CAS}(\ell, v_1, v_2, o_f, o_r, o_w) \{\text{True}\} \quad \{\mathcal{R}_{\text{shr}}^q(\ell)\} \text{*rlx } \ell \{\text{True}\} \quad \{\mathcal{R}_{\text{shr}}^q(\ell)\} \text{*acq } \ell \{\text{True}\} \\ \{\mathcal{R}_{\text{shr}}^q(\ell)\} \mathbf{CAS}(\ell, v_1, v_2, o_f, o_r, o_w) \{\mathcal{R}_{\text{shr}}^q(\ell)\} \end{array}$$

Figure 27: Intermediate-level rules for GPS single-writers.

5 Case Study: `Arc`

Using the lifetime logic for views (§3), the GPFSL logic (§4), and the combination of them, we successfully ported RustBelt’s semantic typing proofs of the concurrent libraries `Arc`, `Mutex`, `RwLock`, `thread::spawn`, and `rayon::join`. Of these seemingly simple libraries, `Arc`, `Mutex`, and `RwLock` encapsulate very subtle correctness conditions, since they support APIs for “extensive” resource reclamation. Our definition of “extensive” reclamation is twofold. First, by extensive reclamation we refer to the ability to fully recover all resources for deallocation at the end, which previous logics were not able to demonstrate. For example, FSL++ (Doko and Vafeiadis [2017]) verifies an `Arc` implementation that can reclaim the content (the ownership of the `T` in `Arc<T>`), but not the internal reference counter of the `Arc` itself. Meanwhile, iGPS (Kaiser et al. [2017]) is able to reclaim the ownership of the atomic integer that implements a spin lock, but is not able to reclaim the shared resource that the lock is protecting in the case the lock is unlocked. Second, by extensive reclamation, we refer to the ability to *temporarily* obtain full ownership of shared resources. For example, `Arc<T>`, `Mutex<T>`, and `RwLock<T>` have APIs that allow a thread to temporarily gain full access to the underlying content `T` when the thread can prove that it is the *unique owner* of the `Arc`, `Mutex`, or `RwLock`. Verifying this kind of temporary reclamation has never been attempted before. To the best of our knowledge, these are the very first formal, machine-checked proofs of relaxed memory algorithms that perform extensive resource reclamation.

In this section, we discuss at a high level the technical challenges in verifying `Arc` with its powerful yet delicate temporary resource reclamation scheme. This attempt to verify the full APIs of `Arc` has led us to the discovery of a data race in one of its temporary reclamation API—namely, `Arc::get_mut`. The data race is due to the use of a relaxed access that does not provide sufficient synchronization, which was not discovered by the original SC RustBelt verification because all atomic accesses had been strengthened to sequential consistency.

5.1 The Full APIs of `Arc`

`Arc<T>`, short for *Atomically Reference Counted*, is used to share atomically a value of type `T`, whose mutation is disabled by default. In order to mutate `T`, programmers need to instantiate `Arc` with wrappers that support thread-safe mutability, for example `Arc<Mutex<T>>` or `Arc<RwLock<T>>`. The following Rust example instantiates `Arc` with an atomic integer `AtomicUsize` and demonstrates how `Arc` is typically used:

```
1 let arc1 = Arc::new(AtomicUsize::new(5));
2 let arc2 = Arc::clone(&arc1);
3 thread::spawn(move || {
4     println!("child thread: {:?}", arc2.fetch_add(1, Ordering::Relaxed));
5 });
6 println!("main thread: {:?}", arc1.fetch_add(2, Ordering::Relaxed));
```

In line 1 in the main thread, a new `Arc` value `arc1` is created to govern an atomic integer allocated in shared memory. The `Arc`’s internal counter for the number of references to the content is set to 1. An `Arc` pointer acts almost like its underlying content, so in line 6 we can call `fetch_add` on `arc1` as if on the atomic integer itself. To share the content with the child thread, we create another `arc2` by `clone`-ing `arc1` (line 2), which effectively increments the internal counter to 2: there are now 2 pointers sharing the atomic integer. When the `Arc` pointers go out of scope (line 4 and 6), their destructors—the `drop` function—are called and the internal counter is decremented accordingly. The last call of `drop` will additionally deallocate the underlying content and the internal counter itself.

Arc	Weak
<code>new: fn(T) -> Arc<T></code>	<code>new: fn() -> Weak<T></code>
<code>deref: fn(&Arc<T>) -> &T</code>	
<code>clone: fn(&Arc<T>) -> Arc<T></code>	<code>clone: fn(&Weak<T>) -> Weak<T></code>
<code>downgrade: fn(&Arc<T>) -> Weak<T></code>	<code>upgrade: fn(&Weak<T>) -> Option<Arc<T>></code>
<code>drop: fn(Arc<T>) -> ()</code>	<code>drop: fn(Weak<T>) -> ()</code>
<code>get_mut: fn(&mut Arc<T>) -> Option<&mut T></code>	
<code>make_mut: fn(&mut Arc<T>) -> &mut T</code>	

Figure 28: An excerpt of Rust’s `Arc<T>` and `Weak<T>` APIs.

As expected, to allow concurrent updates by multiple threads, the internal counter is implemented with an atomic integer.

Fig. 28 gives an excerpt of the most interesting APIs from the `Arc` library. The four functions `Arc::new`, `Arc::deref`, `Arc::clone`, and `Arc::drop` were successfully verified by FSL++ [Doko and Vafeiadis \[2017\]](#). The full APIs of Rust’s `Arc<T>`, however, was not attempted in the relaxed memory setting before. Our verification must therefore tackle two extra sets of behaviors, presented as the following two main challenges.

`Arc<T>` has a subordinate type `Weak<T>` The first challenge involves a type called `Weak<T>`. `Weak<T>` itself is a variant of `Arc<T>`: it has a counter to count how many `Weak` pointers are existing, and also has the similar `clone` and `drop` functions (Fig. 28). However, `Weak<T>` is not tied to the underlying value of type `T`: while owning an `Arc<T>` guarantees that the value is still available, owning a `Weak<T>` does not prevent the underlying value to be reclaimed. Therefore, in order to access the underlying value with a `Weak` pointer, one first has to call `Weak::upgrade` to obtain an `Arc`. `Weak::upgrade` can fail when the value has already been reclaimed, that is when there is no `Arc` pointer left. A `Weak` pointer are typically created by calling `Arc::downgrade` on a shared reference of `Arc`.

The challenge for verifying `Arc` and `Weak` in a relaxed memory model is that they involve two tightly coupled atomic locations. As multi-location invariants are in general unsound for relaxed memory models, we need to use separate GPFSL single-location invariants for each counter and at the same time maintain their relation. This is a known challenge, as has been observed by GPS ([Turon et al. \[2014\]](#)). The general solution is to construct ghost states and/or ghost tokens to encode the relation between the locations and prevent their invariants from breaking the relation. We were able to design ghost states for the two counters of `Arc` and `Weak` without much difficulty.

`Arc<T>` supports temporary borrows of the underlying content The second challenge involves the support to temporarily reclaim full ownership of the underlying content when the thread knows it is the unique `Arc` and `Weak` pointers. The functions `Arc::get_mut` and `Arc::make_mut` provide these capabilities. In particular, `get_mut` (Fig. 29) checks that the thread owns the unique `Arc` and `Weak` pointers in two steps (in `is_unique`):

- First, in line 3 (Fig. 29), it acquires a “lock” on the `Weak` counter to make sure that there is no other `Weak` pointers, so that there cannot be other `Arc` pointers created by `upgrade`-ing.

```

1 fn is_unique(&mut self) -> bool {
2     // lock the weak pointer count if we appear to be the sole weak pointer holder.
3     if self.inner().weak.compare_exchange(1, usize::MAX, Acquire, Relaxed).is_ok() {
4         let unique = self.inner().strong.load(Relaxed) == 1;
5
6         self.inner().weak.store(1, Release); // release the lock
7         unique
8     } else { false }
9 }
10 fn get_mut(this: &mut Self) -> Option<&mut T> {
11     if this.is_unique() {
12         unsafe { Some(&mut this.ptr.as_mut().data) }
13     } else { None }
14 }
15 fn drop(&mut self) {
16     if self.inner().strong.fetch_sub(1, Release) != 1 {
17         return;
18     } ...
19 }

```

Figure 29: Rust’s implementation (excerpt) of `Arc::get_mut` and `Arc::drop`.

- Second, in line 4, it reads the `Arc` counter and check if the value read from the counter is 1.

If the value read in line 4 is 1, `get_mut` concludes that thread owns the unique `Arc` and `Weak` pointers, and gives the thread temporary full access to the underlying content with type `&mut T` (line 12). Otherwise, it releases the lock on the `Weak` counter (line 6) and fails. `Arc::make_mut` also follows the similar pattern of `Arc::get_mut`, but the targets are reversed: it first acquires a “lock” on the `Arc` counter and then reads the `Weak` counter.

It is not obvious at all how to justify to the correctness of `Arc::get_mut` and `Arc::make_mut`. Unfortunately here we cannot discuss the proofs in detail here. A superficial explanation would be that this is the combined result of multiple synchronization points scattered in the code of `Arc` and `Weak`. In other words, their correctness depends on tracking global information across multiple functions of `Arc` and `Weak` (`clone`’s, `drop`’s, `downgrade`, and `upgrade`). For example, in the case of `Arc::get_mut`, the combined information lets the thread know that it has observed all the creation of any possible `Arc` pointers, either by `clone`-ing or `upgrade`-ing. Since every creation of an `Arc` pointer must increment the `Arc` counter by 1, this means that the thread has observed all the increments to the `Arc` counter. So when thread reads 1 from the `Arc` counter, it must be the case that all other `Arc` pointers have been `drop`-ed (each `drop` decrements the counter by 1), and the thread must be owning the last `Arc` pointer. It is then safe to “trade” the ownership of the `Arc` pointer to get temporary access to the underlying content.

For the verification, we design several ghost states to track the history of actions by the `Arc` and `Weak` pointers, and ultimately accumulate their effects to achieve the guarantee needed in `get_mut` and `take_mut`. The ghost states and their derived rules show that, at the read in the check’s second step (line 4, Fig. 29), the thread has collected synchronized, sufficient resources to trade for the full ownership of the content *and* the counters. Such a reasoning has never been attempted before in relaxed memory logics.

5.2 Insufficient Synchronization in `get_mut`

Unfortunately, our setup was not strong enough to verify `Arc` and `Weak` without change. The two reads of the counters in the second check of `get_mut` and `make_mut` were `rlx` in the original code (line 4, Fig. 29), and we had to strengthen them both to `acq` in order to make the verification go through. The reason is that, while we managed to temporarily get the full resources out by a read, the `rlx` reads do not give us those resources in the current view (they are under a ∇ modality). While we conjecture that a `rlx` read in `make_mut` is in fact sufficient, a `rlx` read in `get_mut` turned out to be insufficient and we have reported the bug and the fix has been merged into Rust codebase. The following example invokes a data race when using `get_mut`:

```
1 let mut arc1 = Arc::new(0);
2 let     arc2 = Arc::clone(&arc1);
3 thread::spawn( move || { let _ : u32 = *arc2; /* drop(arc2); */ } );
4 loop { match Arc::get_mut(&mut arc1) {
5     None => {}
6     Some(m) => { *m = 1u32; return; }}}
```

In this example there are two non-atomic operations: the read of the underlying integer in line 3 (child thread) and the write to the same integer in line 6 (parent thread). The read should be safe because the child thread owns `arc2`, thus the underlying integer is shared and *immutable*. The write should be safe because `get_mut` guarantees that the parent thread owns the unique `Arc` pointer (`arc1`) and should temporarily gain full access to the non-atomic integer. This can only happen after the child thread finishes and `arc2` has been dropped. However, the two non-atomic operations constitute a data race by C11 standard, because neither one happens-before the other. More specifically, in line 3 of the child thread, when `arc2` goes out of scope, it will be destructed by `Arc::drop`, which uses a release (`rel`) RMW (see the code at line 16, Fig. 29). This release RMW will be read by `get_mut` (line 4, Fig. 29) in the parent thread (line 4). If this read had been `acq`, then there would have been a release-acquire synchronization between the release RMW of `drop` and the acquire read of `get_mut`, and the non-atomic read of the child thread would have been guaranteed to happen-before the non-atomic write of the parent thread. However, the read was `rlx`, thus no happen-before relationship can be established between the two non-atomic operations.

References

- M. Doko and V. Vafeiadis. A program logic for C11 memory fences. In *VMCAI*, LNCS, pages 413–430. Springer, 2016.
- M. Doko and V. Vafeiadis. Tackling real-life relaxed concurrency with FSL++. In *ESOP*, 2017.
- R. Jung, J.-H. Jourdan, R. Krebbers, and D. Dreyer. RustBelt: Securing the foundations of the Rust programming language – technical appendix and coq development, 2017. <https://plv.mpi-sws.org/rustbelt/pop118/>.
- J.-O. Kaiser, H.-H. Dang, D. Dreyer, O. Lahav, and V. Vafeiadis. Strong logic for weak memory: Reasoning about release-acquire consistency in Iris. In *ECOOP*, LIPIcs, pages 17:1–17:29, 2017.
- O. Lahav, V. Vafeiadis, J. Kang, C.-K. Hur, and D. Dreyer. Repairing sequential consistency in C/C++11. In *PLDI*, 2017.

A. Turon, V. Vafeiadis, and D. Dreyer. GPS: Navigating weak memory with ghosts, protocols, and separation. In *OOPSLA*, pages 691–707. ACM, 2014.