

RustBelt: Securing the Foundations of the Rust Programming Language – Technical appendix

July 7, 2017

Contents

1	Syntax	2
1.1	Grammar	2
1.2	Operational semantics	2
1.3	Continuation-passing-style let-normal programs	3
1.4	Type System	6
1.4.1	Well-formedness	7
1.4.2	Size, Copy, Send, Sync	8
1.4.3	Lifetime context judgments	9
1.4.4	Type Inclusion	9
1.4.5	Well-typed functions and steps	11
2	Some examples	14
3	λ_{Rust} in Iris	21
4	Lifetime logic	22
4.1	Proof rules	22
4.2	Derived forms of borrowing	24
4.3	Model	27
5	λ_{Rust} model	34
5.1	Types	34
5.1.1	Owned pointers and mutable references	35
5.1.2	Shared references	35
5.1.3	Compound types: Sums and products	38
5.1.4	Copy, Send, Sync	38
5.2	Type and continuation contexts	38
5.3	Lifetime contexts and judgments	39
5.4	Judgments	39
5.5	Theorems	39

1 Syntax

1.1 Grammar

λ_{Rust} is a lambda calculus with natural numbers and state, with explicit deallocation, and with a primitive operation to copy regions of memory. Products and sums are not values, the only exist in their heap representation and are manipulated there, or on a per-field basis.

The grammar of the language is given in [Figure 1](#). e is the grammatical class of expressions and v represents values. z is any integer, while n and i are natural numbers. Typically, i is used as an index into something (*e.g.*, a list). ℓ is a heap location, their structure will be defined later. x, f are all program variables, the second usually denoting a function. \bar{x} is a list of x (and similar for other metavariables), which can be constructed as $[x_1, x_2, \dots]$. To simplify the formalization of the semantics, we also introduce a notion of *evaluation contexts* K . Memory accesses are annotated with a *memory order* o , which is either non-atomic (**na**) or sequentially consistent (**sc**). The purpose of this is that the program gets stuck when there are races involving non-atomic accesses. Thus, by proving safety of a program, we show data-race freedom. (The order **na** is just an internal implementation detail.)

This semantics is in some sense too definite, compared to Rust: It fixes the memory representation of products and sums, and it allows mutation of any location at any time.

1.2 Operational semantics

A location $\ell = (i, n)$ consists of a *block* i and an offset into the block n . Allocation and deallocation is always performed on entire blocks. Address arithmetic works within a block: $\ell + m$ increments the offset, and leaves the block number untouched.

A memory h is a finite partial map from locations to pairs of values and lock states:

$$Mem := \mathbb{N} \times \mathbb{N} \xrightarrow{\text{fin}} LockSt \times Val$$

The lock states encode a per-location reader-writer-lock that serves to detect data races. Notice that **reading0** corresponds to the lock being unlocked. Non-atomic accesses require the location to be locked; they will always be immediately preceded by the operation to acquire that location's lock. Atomic accesses, on the other hand, fail if there is a conflicting lock, *i.e.*, all accesses fail if the write lock is held, and write accesses fail if the read lock is held. Putting all these pieces together, we have shown that a program that is *safe* (*i.e.*, cannot get stuck in any execution) is free of data-races: We can never reach a state such that two different threads will, if they get a chance to take a step now, perform *conflicting* memory accesses—*i.e.*, accesses to the same location, at least one of which is non-atomic, and at least one of which is a write.

To define the behavior of equality tests and CAS, we employ a helper judgment $v_1 \vdash v_2 = b$ for $b \in \mathbb{B}$ saying whether v_1 and v_2 can compare (in)equal. In particular, comparing two different locations of which at least one is not allocated is *non-deterministic*: they can compare equal or unequal. For CAS, this means that the CAS could either succeed or fail. The purpose of **O-CAS-STUCK** is to make sure that if such a non-deterministic CAS races with a non-atomic access to the same location, the program is stuck. To this end, if such a situation is detected, we step to the stuck state $0(0)$.

We use the notation $[<n]$ to denote the set $\{m \mid m < n\}$, and $[\geq m, <n]$ to denote $\{m' \mid m \leq m' < n\}$. The notation $h[\ell \leftarrow v]$ denotes the map h updated with location ℓ to map to v . $h[\ell \leftarrow v \mid x \in T]$ does the update for every $x \in T$, where ℓ and v may depend on x .

$z \in \mathbb{Z}$	
$Expr \ni e ::= v \mid x$	$Ctx \ni K ::= \bullet$
$ e.e \mid e + e \mid e - e \mid e \leq e \mid e == e$	$ K.e \mid v.K \mid K + e \mid v + K \mid K - e \mid v - K$
$ e(\bar{e})$	$ K \leq e \mid v \leq K \mid K == e \mid v == K$
$ *^o e$	$ K(\bar{e}) \mid v(\bar{v} ++ [K] ++ \bar{e})$
$ e_1 :=_o e_2$	$ *^o K \mid K :=_o e \mid v :=_o K$
$ \mathbf{CAS}(e_0, e_1, e_2)$	$ \mathbf{CAS}(K, e_1, e_2)$
$ \mathbf{alloc}(e)$	$ \mathbf{CAS}(v_0, K, e_2)$
$ \mathbf{free}(e_1, e_2)$	$ \mathbf{CAS}(v_0, v_1, K)$
$ \mathbf{case } e \mathbf{ of } \bar{e}$	$ \mathbf{alloc}(K)$
$ \mathbf{fork} \{ e \}$	$ \mathbf{free}(K, e_2)$
$Val \ni v ::= () \mid \ell \mid z \mid \mathbf{rec } f(\bar{x}) := e$	$ \mathbf{free}(e_1, K)$
$Loc \ni \ell ::= (i, n)$	$ \mathbf{case } K \mathbf{ of } \bar{e}$
$Order \ni o ::= \mathbf{sc} \mid \mathbf{na} \mid \mathbf{na}'$	
$LockSt \ni \pi ::= \mathbf{writing} \mid \mathbf{reading } n$	

Figure 1: Language syntax.

Note how **O-ALLOC** picks nondeterministic contents for all the fresh allocated cells.

We give the semantics as a small-step reduction relation of machine states, which encompass the current memory and term, written $h \mid e$. Notably, this semantics defines more behaviors than Rust does: The representation of product and sum types is fixed, uninitialized allocated locations have deterministic reads, and it is possible to implement interior mutability without *UnsafeCell*.

1.3 Continuation-passing-style let-normal programs

In this section, we define the surface language that will be used for type-checking. To support control flow operators such as **return** or **break**, the type system will enforce program to be in continuation-passing style and in let-normal form. Before we come to the details of this, we need to define some derived constructions that will be primitive in the surface language. In the following, we use f for program variables that are used as functions, and k for program variables used as continuations.

There are some operations that are not primitive to the language in a syntactic sense, but that come with primitive typing rules: Copying a range of memory, and initializing a sum. These operations are pervasively used, but cannot be typed in the type system. So we define them here as derived forms, together with some useful syntactic sugar for sequencing, non-atomic accesses and conditionals.

Finally, we define the operational behavior of the coercions that start and end lifetimes. They don't actually do anything, but they take a physical step – which we need to make the proofs go through.

$$\begin{array}{c}
\boxed{h \vdash v_1 = v_2} \\
\\
\frac{}{h \vdash z = z} \quad \frac{}{h \vdash \ell = \ell} \quad \frac{\ell_1 \notin \text{dom}(h) \vee \ell_2 \notin \text{dom}(h)}{h \vdash \ell_1 = \ell_2} \\
\boxed{h \vdash v_1 \neq v_2} \\
\\
\frac{z_1 \neq z_2}{h \vdash z_1 \neq z_2} \quad \frac{\ell_1 \neq \ell_2}{h \vdash \ell_1 \neq \ell_2} \\
\\
\boxed{h \mid e \rightarrow h' \mid e'_1, e'_2?} \\
\\
\text{O-ECTX} \quad \frac{h \mid e \rightarrow h \mid e'_1, e'_2?}{h \mid K[e] \rightarrow h \mid K[e'_1], e'_2?} \quad \text{O-PROJ} \quad \frac{}{h \mid \ell.n \rightarrow h \mid \ell + n} \quad \text{O-ADD} \quad \frac{z_1 + z_2 = z'}{h \mid z_1 + z_2 \rightarrow h \mid z'} \quad \text{O-SUB} \quad \frac{z_1 - z_2 = z'}{h \mid z_1 - z_2 \rightarrow h \mid z'} \\
\\
\text{O-LE-TRUE} \quad \frac{z_1 \leq z_2}{h \mid z_1 \leq z_2 \rightarrow h \mid 1} \quad \text{O-LE-FALSE} \quad \frac{z_1 > z_2}{h \mid z_1 \leq z_2 \rightarrow h \mid 0} \quad \text{O-EQ-TRUE} \quad \frac{h \vdash v_1 = v_2}{h \mid v_1 == v_2 \rightarrow h \mid 1} \quad \text{O-EQ-FALSE} \quad \frac{h \vdash v_1 \neq v_2}{h \mid v_1 == v_2 \rightarrow h \mid 0} \\
\\
\text{O-ALLOC} \quad \frac{n > 0 \quad \ell = (i, n') \quad \{i\} \times \mathbb{N} \# \text{dom}(h) \quad h' = h[\ell + m \leftarrow (\text{reading } 0, \bar{v}_m) \mid m \in [< n]]}{h \mid \text{alloc}(n) \rightarrow h' \mid \ell} \\
\\
\text{O-FREE} \quad \frac{n > 0 \quad \ell = (i, n') \quad \text{dom}(h) \cap \{i\} \times \mathbb{N} = \{i\} \times ([\geq n', < n' + n]) \quad h' = h[\ell + m \leftarrow \perp \mid m \in [< n]]}{(h \mid \text{free}(n, \ell)) \rightarrow (h' \mid ())} \\
\\
\text{O-DEREF-SC} \quad \frac{h(\ell) = (\text{reading } n, v)}{h \mid *^{\text{sc}}\ell \rightarrow h \mid v} \quad \text{O-DEREF-NA} \quad \frac{h(\ell) = (\text{reading } n, v)}{(h \mid *^{\text{na}}\ell) \rightarrow (h[\ell \leftarrow (\text{reading } n + 1, v)] \mid *^{\text{na}}\ell)} \\
\\
\text{O-DEREF-NA}' \quad \frac{h(\ell) = (\text{reading } n + 1, v)}{(h \mid *^{\text{na}}\ell) \rightarrow (h[\ell \leftarrow (\text{reading } n, v)] \mid v)} \quad \text{O-ASSIGN-SC} \quad \frac{h(\ell) = (\text{reading } 0, v')}{(h \mid \ell :=_{\text{sc}} v) \rightarrow (h[\ell \leftarrow (\text{reading } 0, v)] \mid ())} \\
\\
\text{O-ASSIGN-NA} \quad \frac{h(\ell) = (\text{reading } 0, v')}{(h \mid \ell :=_{\text{na}} v) \rightarrow (h[\ell \leftarrow (\text{writing}, v')] \mid \ell :=_{\text{na}}, v)} \quad \text{O-ASSIGN-NA}' \quad \frac{h(\ell) = (\text{writing}, v')}{(h \mid \ell :=_{\text{na}}, v) \rightarrow (h[\ell \leftarrow (\text{reading } 0, v)] \mid ())} \\
\\
\text{O-CAS-FAIL} \quad \frac{h(\ell) = (\text{reading } n, v') \quad h \vdash v' \neq v_1}{(h \mid \text{CAS}(\ell, v_1, v_2)) \rightarrow (h \mid 0)} \quad \text{O-CAS-SUC} \quad \frac{h(\ell) = (\text{reading } 0, v') \quad h \vdash v' = v_1}{(h \mid \text{CAS}(\ell, v_1, v_2)) \rightarrow (h[\ell \leftarrow (\text{reading } 0, z_2)] \mid 1)} \\
\\
\text{O-CAS-STUCK} \quad \frac{h(\ell) = (\text{reading } n, v') \quad n > 0 \quad h \vdash v' = v_1}{(h \mid \text{CAS}(\ell, v_1, v_2)) \rightarrow (h \mid 0())} \quad \text{O-CASE} \quad \frac{}{(h \mid \text{case } i \text{ of } \bar{e}) \rightarrow (h \mid \bar{e}_i)} \\
\\
\text{O-APP} \quad \frac{}{(h \mid (\text{rec } f(\bar{x}) := e)(\bar{v})) \rightarrow (h \mid e[\text{rec } f(\bar{x}) := e/f, \bar{v}/\bar{x}])} \quad \text{O-FORK} \quad \frac{}{h \mid \text{fork } \{e\} \rightarrow h \mid (), e}
\end{array}$$

Figure 2: Operational semantics.

```

funrec  $f(\bar{x})$  ret  $k := e := \text{rec } f([k] \uparrow \bar{x}) := e$ 
    let  $x = e$  in  $e' := (\text{rec\_}([x]) := e')(e)$ 
     $e'; e := \text{let\_} = e' \text{ in } e$ 
letcont  $k(\bar{x}) := e$  in  $e' := \text{let } k = (\text{rec } k(\bar{x}) := e) \text{ in } e'$ 
    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  $e_2 := \text{case } e_0 \text{ of } [e_1, e_2]$ 

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

    skip := let  $x = ()$  in  $x$ 
    newlft := ()
    endlft := skip

```

We distinguish between three classes of expressions: *function bodies* F consist of *instructions* I that operate on *paths* p . The **letcall** operator makes it possible to call *other functions*, passing the remainder of the current function as a continuation. This is in contrast to calling a continuation by just jumping there. We the purpose of λ_{Rust} , we are not concerned with whole programs, but only

with the individual functions of a program.

$$\begin{aligned}
Path \ni p &::= x \mid p.n \\
Instr \ni I &::= \mathbf{false} \mid \mathbf{true} \mid z \mid \mathbf{funrec} f(\bar{x}) \mathbf{ret} k := F \mid p \mid p_1 + p_2 \mid p_1 - p_2 \mid p_1 \leq p_2 \\
&\quad \mid \mathbf{new}(n) \mid \mathbf{delete}(n, p) \mid *p \mid p_1 := p_2 \mid p \stackrel{\text{inj } i}{=} () \mid p_1 \stackrel{\text{inj } i}{=} p_2 \mid p_1 :=_n *p_2 \mid p_1 \stackrel{\text{inj } i}{=}_n *p_2 \\
FuncBody \ni F &::= \mathbf{let} x = I \mathbf{in} F \mid \mathbf{letcont} k(\bar{x}) := F_1 \mathbf{in} F_2 \mid \mathbf{newlft}; F \mid \mathbf{endlft}; F \\
&\quad \mid \mathbf{if} p \mathbf{then} F_1 \mathbf{else} F_2 \mid \mathbf{case} *p \mathbf{of} \bar{F} \mid \mathbf{jump} k(\bar{p}) \mid \mathbf{call} f(\bar{p}) \mathbf{ret} k
\end{aligned}$$

1.4 Type System

The key concept of the λ_{Rust} type system is the notion of a *lifetime*. Essentially, a lifetime represents a part of the program execution.

Programs are type-checked under five contexts: the *variable context* Γ contains all binders. It assigns variables to their *sort* σ (either program variable $x : \mathbf{val}$, a lifetime $\alpha : \mathbf{lft}$ or a type $T : \mathbf{type}$). Furthermore, there is a context for *external lifetimes* \mathbf{E} , a context for *local lifetimes* \mathbf{L} , a context assigning *types* to variables \mathbf{T} and a context managing *continuations* \mathbf{K} .

$$\begin{aligned}
Sort \ni \sigma &::= \mathbf{val} \mid \mathbf{lft} \mid \mathbf{type} \\
\Gamma &::= \emptyset \mid \Gamma, X : \sigma \\
Lft \ni \kappa &::= \alpha \mid \mathbf{static} \\
\mathbf{E} &::= \emptyset \mid \mathbf{E}, \kappa \sqsubseteq_e \kappa' \\
\mathbf{L} &::= \emptyset \mid \mathbf{L}, \kappa \sqsubseteq_1 \bar{\kappa} \\
Mod \ni \mu &::= \mathbf{mut} \mid \mathbf{shr} \\
Type \ni \tau &::= T \mid \mathbf{bool} \mid \mathbf{int} \mid \downarrow_n \\
&\quad \mid \mathbf{own}_n \tau \mid \&_{\mu}^{\kappa} \tau \mid \Sigma \bar{\tau} \mid \Pi \bar{\tau} \mid \forall \bar{\alpha}. \mathbf{fn}(\bar{f} : \mathbf{E}; \bar{\tau}) \rightarrow \tau \mid \mu T. \tau \\
\mathbf{T} &::= \emptyset \mid \mathbf{T}, p \triangleleft \tau \mid \mathbf{T}, p \triangleleft^{\dagger \kappa} \tau \\
\mathbf{K} &::= \emptyset \mid \mathbf{K}, k \triangleleft \mathbf{cont}(\mathbf{L}; \bar{x}. \mathbf{T})
\end{aligned}$$

Types describe not just single values, but regions of memory – for now, you can think of them as making statements about *lists* of values. The type system introduces a distinction between functions and continuations, matching the program grammar given in §1.3. In particular, continuations are not types. We use this distinction to control that continuations are not passed to other functions or otherwise leaked from their context. We will use function types for Rust-level functions, and continuation types for the basic blocks of a Rust function.

The types of the form $\mathbf{own}_n \tau$ represent *owned pointers*. The index n gives the size of the block that this object has been allocated in; this is important because only entire blocks can be deallocated.

A particularly interesting type is a *reference* $\&_{\mu}^{\kappa} \tau$, also called (*temporarily*) *borrowed pointer*. References are qualified by a *modifier*, which is either **mut** (mutable, *i.e.*, unique) and **shr** (shared). This pointer is only valid as long as its *lifetime* is still active, which can be proven by showing that the lifetime is in the lifetime context \mathbf{L} . Functions can be polymorphic over lifetimes.

Finally, the type system supports *recursive types*, with the restriction (enforced by well-formedness) that the recursive occurrence is below a *pointer type*.

The type context can contain “normal” type assignments ($p \triangleleft \tau$) and type assignments that are *blocked by a lifetime*: $p \triangleleft^{\dagger\kappa} \tau$ means that we can only use this type assignment again when κ has ended.

1.4.1 Well-formedness

The well-formedness judgments (\vdash_{wf}) document the binding structure of our grammar. There should be no surprises. In the other judgments defined later, we always implicitly assume well-formedness and we also will frequently leave the variable context Γ implicit.

Well-formed paths

$$\boxed{\Gamma \vdash_{\text{wf}} p}$$

$$\frac{x : \mathbf{val} \in \Gamma}{\Gamma \vdash_{\text{wf}} x} \qquad \frac{\Gamma \vdash_{\text{wf}} p}{\Gamma \vdash_{\text{wf}} p.n}$$

Well-formed lifetimes

$$\boxed{\Gamma \vdash_{\text{wf}} \kappa}$$

$$\frac{\alpha : \mathbf{lft} \in \Gamma}{\Gamma \vdash_{\text{wf}} \alpha} \qquad \Gamma \vdash_{\text{wf}} \mathbf{static}$$

Well-formed external lifetime contexts

$$\boxed{\Gamma \vdash_{\text{wf}} \mathbf{E}}$$

$$\Gamma \vdash_{\text{wf}} \emptyset \qquad \frac{\Gamma \vdash_{\text{wf}} \mathbf{L} \quad \Gamma \vdash_{\text{wf}} \kappa \quad \Gamma \vdash_{\text{wf}} \kappa'}{\Gamma \vdash_{\text{wf}} \mathbf{L}, \kappa \sqsubseteq_e \kappa'}$$

Well-formed local lifetime contexts

$$\boxed{\Gamma \vdash_{\text{wf}} \mathbf{L}}$$

$$\Gamma \vdash_{\text{wf}} \emptyset \qquad \frac{\Gamma \vdash_{\text{wf}} \mathbf{L} \quad \Gamma \vdash_{\text{wf}} \kappa \quad \forall \kappa' \in \bar{\kappa}. \Gamma \vdash_{\text{wf}} \kappa'}{\Gamma \vdash_{\text{wf}} \mathbf{L}, \kappa \sqsubseteq_l \bar{\kappa}}$$

Well-formed types

$$\boxed{\Gamma \vdash_{\text{wf}} \tau}$$

$$\begin{array}{c} \frac{T : \mathbf{type} \in \Gamma}{\Gamma \vdash_{\text{wf}} T} \quad \Gamma \vdash_{\text{wf}} \mathbf{bool} \quad \Gamma \vdash_{\text{wf}} \mathbf{int} \quad \Gamma \vdash_{\text{wf}} \not\downarrow_n \quad \frac{\Gamma \vdash_{\text{wf}} \tau}{\Gamma \vdash_{\text{wf}} \mathbf{own}_n \tau} \quad \frac{\Gamma \vdash_{\text{wf}} \kappa \quad \Gamma \vdash_{\text{wf}} \tau}{\Gamma \vdash_{\text{wf}} \&_{\mu}^{\kappa} \tau} \\[10pt] \frac{\forall i. \Gamma \vdash_{\text{wf}} \bar{\tau}_i}{\Gamma \vdash_{\text{wf}} \Pi \bar{\tau}} \quad \frac{\forall i. \Gamma \vdash_{\text{wf}} \bar{\tau}_i}{\Gamma \vdash_{\text{wf}} \Sigma \bar{\tau}} \quad \frac{\Gamma, \bar{\alpha}, \bar{f} : \mathbf{lft} \vdash_{\text{wf}} \mathbf{E} \quad \forall i. \Gamma, \bar{\alpha} : \mathbf{lft} \vdash_{\text{wf}} \bar{\tau}_i \quad \Gamma, \bar{\alpha} : \mathbf{lft} \vdash_{\text{wf}} \tau}{\Gamma \vdash_{\text{wf}} \forall \bar{\alpha}. \mathbf{fn}(\bar{f} : \mathbf{E}; \bar{\tau}) \rightarrow \tau} \\[10pt] \frac{\Gamma, T : \mathbf{type} \vdash_{\text{wf}} \tau \quad T \text{ is guarded by pointer types in } \tau}{\Gamma \vdash_{\text{wf}} \mu T. \tau} \end{array}$$

Well-formed type contexts

$$\boxed{\Gamma \vdash_{\text{wf}} \mathbf{T}}$$

$$\Gamma \vdash_{\text{wf}} \emptyset \qquad \frac{\Gamma \vdash_{\text{wf}} \mathbf{T} \quad \Gamma \vdash_{\text{wf}} p \quad \Gamma \vdash_{\text{wf}} \tau}{\Gamma \vdash_{\text{wf}} \mathbf{T}, p \triangleleft \tau} \qquad \frac{\Gamma \vdash_{\text{wf}} \mathbf{T} \quad \Gamma \vdash_{\text{wf}} p \quad \Gamma \vdash_{\text{wf}} \kappa \quad \Gamma \vdash_{\text{wf}} \tau}{\Gamma \vdash_{\text{wf}} \mathbf{T}, p \triangleleft^{\dagger\kappa} \tau}$$

Well-formed continuation contexts

$\Gamma \vdash_{\text{wf}} \mathbf{K}$

$$\frac{\Gamma \vdash_{\text{wf}} \emptyset \quad \Gamma \vdash_{\text{wf}} \mathbf{T} \quad \Gamma \vdash_{\text{wf}} k \quad \Gamma \vdash_{\text{wf}} \mathbf{L} \quad \Gamma, \bar{x} : \mathbf{val} \vdash_{\text{wf}} \mathbf{T}}{\Gamma \vdash_{\text{wf}} \mathbf{T}, k \triangleleft \mathbf{cont}(\mathbf{L}; \bar{x}. \mathbf{T})}$$

1.4.2 Size, Copy, Send, Sync

The *size* of a type says how many memory locations a type spans. It is defined as follows:

$$\begin{aligned} \text{size}(\mathbf{bool}) &:= 1 & \text{size}(\mathbf{own}_n \tau) &:= 1 \\ \text{size}(\mathbf{int}) &:= 1 & \text{size}(\&_{\mu}^{\kappa} \tau) &:= 1 \\ \text{size}(\&_n) &:= n & \text{size}(\Pi \bar{\tau}) &:= \sum_i \text{size}(\bar{\tau}_i) \\ & & \text{size}(\Sigma \bar{\tau}) &:= 1 + \max_i \text{size}(\bar{\tau}_i) \\ \text{size}(\mu T. \tau) &:= \text{size}(\tau) & \text{size}(\forall \bar{\alpha}. \mathbf{fn}(\mathbf{F} : \mathbf{E}; \bar{\tau}) \rightarrow \tau) &:= 1 \end{aligned}$$

Notice that there is no case for type variables: since well-formed recursive types always have their recursive occurrence below a pointer type, the size of a recursive type does not depend on the size of the recursive occurrence.

Some types are *copyable*, which means they can be used arbitrarily often. This is expressed by the following judgment. Notice that in proving a recursive type to be copyable, you can assume the type variable T to be copy.

Copy types

$\tau \text{ copy}$

$$\begin{array}{c} \mathbf{bool} \text{ copy} \quad \mathbf{int} \text{ copy} \quad \&_n \text{ copy} \quad \&_{\text{shr}}^{\kappa} \tau \text{ copy} \quad \frac{\forall i. \tau_i \text{ copy}}{\Pi \bar{\tau} \text{ copy}} \quad \frac{\forall i. \tau_i \text{ copy}}{\Sigma \bar{\tau} \text{ copy}} \\ \\ (\forall \bar{\alpha}. \mathbf{fn}(\mathbf{F} : \mathbf{E}; \bar{\tau}) \rightarrow \tau) \text{ copy} \quad T \text{ copy} \quad \frac{\tau \text{ copy}}{\mu T. \tau \text{ copy}} \end{array}$$

A type is *send* if ownership can be transferred to another thread. It is *sync* if shared instances of the type can be transferred to another thread.

Send types

$\tau \text{ send}$

$$\begin{array}{c} \mathbf{bool} \text{ send} \quad \mathbf{int} \text{ send} \quad \&_n \text{ send} \quad \frac{\tau \text{ send}}{\mathbf{own}_n \tau \text{ send}} \quad \frac{\tau \text{ send}}{\&_{\text{mut}}^{\kappa} \tau \text{ send}} \quad \frac{\tau \text{ sync}}{\&_{\text{shr}}^{\kappa} \tau \text{ send}} \quad \frac{\forall i. \tau_i \text{ send}}{\Pi \bar{\tau} \text{ send}} \\ \\ \frac{\forall i. \tau_i \text{ send}}{\Sigma \bar{\tau} \text{ send}} \quad (\forall \bar{\alpha}. \mathbf{fn}(\mathbf{F} : \mathbf{E}; \bar{\tau}) \rightarrow \tau) \text{ send} \quad T \text{ send} \quad \frac{\tau \text{ send}}{\mu T. \tau \text{ send}} \end{array}$$

Sync types

$$\begin{array}{c}
\boxed{\tau \text{ sync}} \\
\text{bool sync} \quad \text{int sync} \quad \downarrow_n \text{ sync} \quad \frac{\tau \text{ sync}}{\text{own}_n \tau \text{ sync}} \quad \frac{\tau \text{ sync}}{\&_{\text{mut}}^\kappa \tau \text{ sync}} \quad \frac{\tau \text{ sync}}{\&_{\text{shr}}^\kappa \tau \text{ sync}} \quad \frac{\forall i. \tau_i \text{ sync}}{\Pi \bar{\tau} \text{ sync}} \\
\\
\frac{\forall i. \tau_i \text{ sync}}{\Sigma \bar{\tau} \text{ sync}} \quad (\forall \bar{\alpha}. \mathbf{fn}(\mathbf{f} : \mathbf{E}; \bar{\tau}) \rightarrow \tau) \text{ sync} \quad T \text{ sync} \quad \frac{\tau \text{ sync}}{\mu T. \tau \text{ sync}}
\end{array}$$

1.4.3 Lifetime context judgments

The following judgments express various properties of lifetime contexts.

Lifetime inclusion

$$\begin{array}{c}
\boxed{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa_1 \sqsubseteq \kappa_2} \\
\\
\mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \mathbf{static} \quad \frac{\kappa \sqsubseteq_1 \bar{\kappa} \in \mathbf{L} \quad \kappa' \in \bar{\kappa}}{\mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \kappa'} \quad \frac{\kappa \sqsubseteq_e \kappa' \in \mathbf{E}}{\mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \kappa'} \quad \mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \kappa \\
\\
\frac{\mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \kappa' \quad \mathbf{E}; \mathbf{L} \vdash \kappa' \sqsubseteq \kappa''}{\mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \kappa''}
\end{array}$$

Lifetime liveness

$$\begin{array}{c}
\boxed{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive}} \\
\\
\mathbf{E}; \mathbf{L} \vdash \mathbf{static} \text{ alive} \quad \frac{\kappa \sqsubseteq_1 \bar{\kappa} \in \mathbf{L} \quad \forall i. \mathbf{E}; \mathbf{L} \vdash \bar{\kappa}_i \text{ alive}}{\mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive}} \quad \frac{\mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive} \quad \mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \kappa'}{\mathbf{E}; \mathbf{L} \vdash \kappa' \text{ alive}}
\end{array}$$

Local lifetime context inclusion

$$\boxed{\Gamma \vdash \mathbf{L}_1 \Rightarrow \mathbf{L}_2}$$

$$\frac{\mathbf{L}' \text{ is a permutation of } \mathbf{L}}{\mathbf{L} \Rightarrow \mathbf{L}'}$$

External lifetime context satisfiability

$$\boxed{\Gamma \mid \mathbf{E}_1; \mathbf{L}_1 \vdash \mathbf{E}_2}$$

$$\mathbf{E}_1; \mathbf{L}_1 \vdash \emptyset \quad \frac{\mathbf{E}_1; \mathbf{L}_1 \vdash \kappa \sqsubseteq \kappa' \quad \mathbf{E}_1; \mathbf{L}_1 \vdash \mathbf{E}_2}{\mathbf{E}_1; \mathbf{L}_1 \vdash \mathbf{E}_2, \kappa \sqsubseteq_e \kappa'}$$

1.4.4 Type Inclusion

The main subtyping supported in Rust is lifetime inclusion and (un)folding recursive types. Furthermore, products of uninitialized types are equivalent to one large uninitialized type. Finally, type constructors have structural rules witnessing covariance and contravariance of type constructors. This is reflected in our subtyping relation. Some of the rules state an equivalence (\Rightarrow), which is meant as sugar for mutual inclusion.

Subtyping

$$\boxed{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \tau_1 \Rightarrow \tau_2}$$

$$\begin{array}{c}
\text{T-REFL} \\
\mathbf{E}; \mathbf{L} \vdash \tau \Rightarrow \tau
\end{array}
\quad
\frac{\text{T-TRANS} \quad \mathbf{E}; \mathbf{L} \vdash \tau \Rightarrow \tau' \quad \mathbf{E}; \mathbf{L} \vdash \tau' \Rightarrow \tau''}{\mathbf{E}; \mathbf{L} \vdash \tau \Rightarrow \tau''}
\quad
\frac{\text{T-BOR-LFT} \quad \mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \kappa'}{\mathbf{E}; \mathbf{L} \vdash \&_{\mu}^{\kappa'} \tau \Rightarrow \&_{\mu}^{\kappa} \tau}$$

$$\frac{\text{T-UNINIT-PROD} \quad \mathbf{E}; \mathbf{L} \vdash \not\sqsubseteq_{\Sigma \bar{n}} \Leftrightarrow \Pi_{\not\sqsubseteq n}}{\mathbf{E}; \mathbf{L} \vdash \mu T_1. \tau_1 \Rightarrow \mu T_2. \tau_2}
\quad
\frac{\text{T-REC} \quad \forall \tau'_1, \tau'_2. (\mathbf{E}; \mathbf{L} \vdash \tau'_1 \Rightarrow \tau'_2) \Rightarrow (\mathbf{E}; \mathbf{L} \vdash \tau_1[\tau'_1/T_1] \Rightarrow \tau_2[\tau'_2/T_2])}{\mathbf{E}; \mathbf{L} \vdash \mu T_1. \tau_1 \Rightarrow \mu T_2. \tau_2}$$

$$\frac{\text{T-REC-UNFOLD} \quad \mathbf{E}; \mathbf{L} \vdash \mu T. \tau \Leftrightarrow \tau[\mu T. \tau/T]}{\mathbf{E}; \mathbf{L} \vdash \mu T. \tau \Leftrightarrow \tau[\mu T. \tau/T]}
\quad
\frac{\text{T-OWN} \quad \mathbf{E}; \mathbf{L} \vdash \tau_1 \Rightarrow \tau_2}{\mathbf{E}; \mathbf{L} \vdash \mathbf{own}_n \tau_1 \Rightarrow \mathbf{own}_n \tau_2}
\quad
\frac{\text{T-BOR-SHR} \quad \mathbf{E}; \mathbf{L} \vdash \tau_1 \Rightarrow \tau_2}{\mathbf{E}; \mathbf{L} \vdash \&_{\text{shr}}^{\kappa} \tau_1 \Rightarrow \&_{\text{shr}}^{\kappa} \tau_2}$$

$$\frac{\text{T-BOR-MUT} \quad \mathbf{E}; \mathbf{L} \vdash \tau_1 \Leftrightarrow \tau_2}{\mathbf{E}; \mathbf{L} \vdash \&_{\text{mut}}^{\kappa} \tau_1 \Leftrightarrow \&_{\text{mut}}^{\kappa} \tau_2}
\quad
\frac{\text{T-PROD} \quad \forall i. \mathbf{E}; \mathbf{L} \vdash \bar{\tau}_i \Rightarrow \bar{\tau}'_i}{\mathbf{E}; \mathbf{L} \vdash \Pi \bar{\tau} \Rightarrow \Pi \bar{\tau}'}
\quad
\frac{\text{T-SUM} \quad \forall i. \mathbf{E}; \mathbf{L} \vdash \bar{\tau}_i \Rightarrow \bar{\tau}'_i}{\mathbf{E}; \mathbf{L} \vdash \Sigma \bar{\tau} \Rightarrow \Sigma \bar{\tau}'}$$

$$\frac{\text{T-FN} \quad \Gamma, \bar{\alpha}', \text{fn} : \mathbf{lft} \mid \mathbf{E}', \mathbf{E}_0; \mathbf{L}_0 \vdash \mathbf{E}[\bar{\kappa}/\bar{\alpha}]}{\Gamma \mid \mathbf{E}_0; \mathbf{L}_0 \vdash \forall \bar{\alpha}. \text{fn}(\text{fn} : \mathbf{E}; \bar{\tau}) \rightarrow \tau \Rightarrow \forall \bar{\alpha}'. \text{fn}(\text{fn} : \mathbf{E}'; \bar{\tau}') \rightarrow \tau'}$$

Inclusion of type *contexts* does not just allow applying subtyping; there are also a few coercions supported by the type system. Most notably, a mutable reference can be coerced to a shared reference.

Type context inclusion

$$\boxed{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \mathbf{T}_1 \Rightarrow \mathbf{T}_2}$$

$$\frac{\text{C-PERM} \quad \mathbf{T}' \text{ is a permutation of } \mathbf{T}}{\mathbf{E}; \mathbf{L} \vdash \mathbf{T} \Rightarrow \mathbf{T}'}
\quad
\frac{\text{C-WEAKEN} \quad \mathbf{E}; \mathbf{L} \vdash \mathbf{T}, \mathbf{T}' \Rightarrow \mathbf{T}}{\mathbf{E}; \mathbf{L} \vdash \mathbf{T}, \mathbf{T}' \Rightarrow \mathbf{T}}
\quad
\frac{\text{C-FRAME} \quad \mathbf{E}; \mathbf{L} \vdash \mathbf{T}_1 \Rightarrow \mathbf{T}_2}{\mathbf{E}; \mathbf{L} \vdash \mathbf{T}', \mathbf{T}_1 \Rightarrow \mathbf{T}', \mathbf{T}_2}$$

$$\frac{\text{C-COPY} \quad \tau \text{ copy}}{\mathbf{E}; \mathbf{L} \vdash p \triangleleft \tau \Rightarrow p \triangleleft \tau, p \triangleleft \tau}
\quad
\frac{\text{C-SUBTYPE} \quad \mathbf{E}; \mathbf{L} \vdash \tau \Rightarrow \tau'}{\mathbf{E}; \mathbf{L} \vdash p \triangleleft \tau \Rightarrow p \triangleleft \tau'}
\quad
\frac{\text{C-SHARE} \quad \mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive}}{\mathbf{E}; \mathbf{L} \vdash p \triangleleft \&_{\text{mut}}^{\kappa} \tau \Rightarrow p \triangleleft \&_{\text{shr}}^{\kappa} \tau}$$

$$\frac{\text{C-SPLIT-OWN} \quad \bar{\tau} \neq [] \quad \forall i. m_i = \sum_{j < i} \text{size}(\bar{\tau}_j)}{\mathbf{E}; \mathbf{L} \vdash p \triangleleft \mathbf{own}_n \Pi \bar{\tau} \Rightarrow p.m \triangleleft \mathbf{own}_n \tau}
\quad
\frac{\text{C-SPLIT-BOR} \quad \bar{\tau} \neq [] \quad \forall i. m_i = \sum_{j < i} \text{size}(\bar{\tau}_j)}{\mathbf{E}; \mathbf{L} \vdash p \triangleleft \&_{\mu}^{\kappa} \Pi \bar{\tau} \Rightarrow p.m \triangleleft \&_{\mu}^{\kappa} \tau}$$

$$\frac{\text{C-BORROW} \quad \mathbf{E}; \mathbf{L} \vdash p \triangleleft \mathbf{own}_n \tau \Rightarrow p \triangleleft \&_{\text{mut}}^{\kappa} \tau, p \triangleleft \dagger^{\kappa} \mathbf{own}_n \tau}{\mathbf{E}; \mathbf{L} \vdash p \triangleleft \&_{\mu}^{\kappa} \tau \Rightarrow p \triangleleft \&_{\mu}^{\kappa'} \tau, p \triangleleft \dagger^{\kappa'} \&_{\mu}^{\kappa} \tau}
\quad
\frac{\text{C-REBORROW} \quad \mathbf{E}; \mathbf{L} \vdash \kappa' \sqsubseteq \kappa}{\mathbf{E}; \mathbf{L} \vdash p \triangleleft \&_{\mu}^{\kappa} \tau \Rightarrow p \triangleleft \&_{\mu}^{\kappa'} \tau, p \triangleleft \dagger^{\kappa'} \&_{\mu}^{\kappa} \tau}$$

The following judgment expresses that when κ ends, we can “unblock” the parts of the typing context that is blocked by κ .

Type context unblocking

$$\boxed{\Gamma \vdash \mathbf{T}_1 \Rightarrow^{\dagger\kappa} \mathbf{T}_2}$$

$$\begin{array}{c} \emptyset \Rightarrow^{\dagger\kappa} \emptyset \\ \frac{\mathbf{T}_1 \Rightarrow^{\dagger\kappa} \mathbf{T}_2}{\mathbf{T}_{1,p} \triangleleft \tau \Rightarrow^{\dagger\kappa} \mathbf{T}_{2,p} \triangleleft \tau} \quad \frac{\mathbf{T}_1 \Rightarrow^{\dagger\kappa} \mathbf{T}_2}{\mathbf{T}_{1,p} \triangleleft^{\dagger\kappa} \tau \Rightarrow^{\dagger\kappa} \mathbf{T}_{2,p} \triangleleft \tau} \\ \frac{\mathbf{T}_1 \Rightarrow^{\dagger\kappa} \mathbf{T}_2}{\mathbf{T}_{1,p} \triangleleft^{\dagger\kappa'} \tau \Rightarrow^{\dagger\kappa} \mathbf{T}_{2,p} \triangleleft^{\dagger\kappa'} \tau} \end{array}$$

Continuation context inclusion

$$\boxed{\Gamma \mid \mathbf{E} \vdash \mathbf{K}_1 \Rightarrow \mathbf{K}_2}$$

$$\begin{array}{c} \frac{\mathbf{K}' \text{ is a permutation of } \mathbf{K}}{\mathbf{E} \vdash \mathbf{K} \Rightarrow \mathbf{K}'} \quad \mathbf{E} \vdash \mathbf{K}, \mathbf{K}' \Rightarrow \mathbf{K} \\ \frac{\Gamma \mid \mathbf{E} \vdash \mathbf{K} \Rightarrow \mathbf{K}' \quad \Gamma, \bar{x} : \mathbf{val} \mid \mathbf{E}; \mathbf{L} \vdash \mathbf{T}' \asymp \mathbf{T}}{\Gamma \mid \mathbf{E} \vdash \mathbf{K}, k \triangleleft \mathbf{cont}(\mathbf{L}; \bar{x}. \mathbf{T}) \Rightarrow \mathbf{K}', k \triangleleft \mathbf{cont}(\mathbf{L}; \bar{x}. \mathbf{T}')} \end{array}$$

1.4.5 Well-typed functions and steps

Finally we come to the main typing judgment: $\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \vdash F$ says that F is a well-typed *function body* (as defined by the grammar in §1.3). This means that, under the assumptions described by the contexts, the function is safe to execute. (Functions are in CPS and hence do not return.)

The grammar dictates that a function consists of a bunch of continuations (representing basic blocks) that each consist of a sequence of *instructions*. Instructions *do* return and produce a value, so their typing judgment $\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{T}_1 \vdash I \dashv x. \mathbf{T}_2$ features two typing contexts: if \mathbf{T}_1 holds before the step is executed, then \mathbf{T}_2 holds after the step was executed.

We also have two of small helper judgments to express what loading from memory and storing to memory does to types. $\mathbf{E}; \mathbf{L} \vdash \tau_1 \multimap^\tau \tau_2$ says that we can write something of type τ to a location described by τ_1 , which will change the type of the location to τ_2 . Similarly, $\mathbf{E}; \mathbf{L} \vdash \tau_1 \multimap^\tau \tau_2$ says that when reading from a location of type τ_1 , we will read something of type τ and the type of the locations changes to τ_2 .

Well-typed functions

$$\boxed{\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \vdash F}$$

$$\begin{array}{c} \text{F-CONSEQUENCE} \\ \frac{\mathbf{L} \Rightarrow \mathbf{L}' \quad \mathbf{E}; \mathbf{L} \vdash \mathbf{T} \asymp \mathbf{T}' \quad \mathbf{E} \vdash \mathbf{K} \Rightarrow \mathbf{K}' \quad \mathbf{E}; \mathbf{L}' \mid \mathbf{K}'; \mathbf{T}' \vdash F}{\mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \vdash F} \\ \text{F-EQUALIZE} \quad \text{F-LET} \\ \frac{\mathbf{E}, \alpha \sqsubseteq_e \kappa, \kappa \sqsubseteq_e \alpha; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \vdash F}{\mathbf{E}; \mathbf{L}, \alpha \sqsubseteq_1 [\kappa] \mid \mathbf{K}; \mathbf{T} \vdash F} \quad \frac{\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{T}_1 \vdash I \dashv x. \mathbf{T}_2 \quad \Gamma, x : \mathbf{val} \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}_2, \mathbf{T} \vdash F}{\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}_1, \mathbf{T} \vdash \mathbf{let } x = I \mathbf{ in } F} \end{array}$$

$$\begin{array}{c}
\text{F-LETCONT} \\
\frac{\Gamma, k, \bar{x} : \mathbf{val} \mid \mathbf{E}; \mathbf{L}_1 \mid \mathbf{K}, k \triangleleft \mathbf{cont}(\mathbf{L}_1; \bar{x}. \mathbf{T}'); \mathbf{T}' \vdash F_1 \quad \Gamma, k : \mathbf{val} \mid \mathbf{E}; \mathbf{L}_2 \mid \mathbf{K}, k \triangleleft \mathbf{cont}(\mathbf{L}_1; \bar{x}. \mathbf{T}'); \mathbf{T} \vdash F_2}{\Gamma \mid \mathbf{E}; \mathbf{L}_2 \mid \mathbf{K}; \mathbf{T} \vdash \mathbf{letcont} \, k(\bar{x}) := F_1 \mathbf{in} \, F_2} \\
\\
\text{F-IF} \quad \frac{\mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \vdash F_1 \quad \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \vdash F_2}{\mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}, p \triangleleft \mathbf{bool} \vdash \mathbf{if} \, p \mathbf{then} \, F_1 \mathbf{else} \, F_2} \quad \text{F-JUMP} \quad \frac{\mathbf{E}; \mathbf{L} \vdash \mathbf{T} \Rightarrow \mathbf{T}'[\bar{y}/\bar{x}]}{\mathbf{E}; \mathbf{L} \mid k \triangleleft \mathbf{cont}(\mathbf{L}; \bar{x}. \mathbf{T}'); \mathbf{T} \vdash \mathbf{jump} \, k(\bar{y})} \\
\\
\text{F-CALL} \quad \frac{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \mathbf{T} \Rightarrow \bar{p} \triangleleft \mathbf{own} \, \bar{\tau}, \mathbf{T}' \quad \mathbf{E}; \mathbf{L} \vdash \bar{\kappa} \text{ alive} \quad \Gamma, \mathbf{f} : \mathbf{lft} \mid \mathbf{E}, \mathbf{f} \sqsubseteq_{\mathbf{e}} \bar{\kappa}; \mathbf{L} \vdash \mathbf{E}'}{\Gamma \mid \mathbf{E}; \mathbf{L} \mid k \triangleleft \mathbf{cont}(\mathbf{L}; y. y \triangleleft \mathbf{own} \, \tau, \mathbf{T}'); \mathbf{T}, \mathbf{f} \triangleleft \mathbf{fn}(\mathbf{f} : \mathbf{E}'; \bar{\tau}) \rightarrow \tau \vdash \mathbf{call} \, f(\bar{p}) \mathbf{ret} \, k} \\
\\
\text{F-NEWLFT} \quad \frac{\Gamma, \alpha : \mathbf{lft} \mid \mathbf{E}; \mathbf{L}, \alpha \sqsubseteq_1 \bar{\kappa} \mid \mathbf{K}; \mathbf{T} \vdash F}{\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \vdash \mathbf{newlft}; F} \quad \text{F-ENDLFT} \quad \frac{\mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}' \vdash F \quad \mathbf{T} \Rightarrow^{\dagger \kappa} \mathbf{T}'}{\mathbf{E}; \mathbf{L}, \kappa \sqsubseteq_1 \bar{\kappa} \mid \mathbf{K}; \mathbf{T} \vdash \mathbf{endlft}; F} \\
\\
\text{F-CASE-OWN} \quad \frac{\forall i. (\mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}, p.0 \triangleleft \mathbf{own}_n \not\triangleleft, p.1 \triangleleft \mathbf{own}_n \bar{\tau}_i, p.(1 + \text{size}(\bar{\tau}_i)) \triangleleft \mathbf{own}_n \not\triangleleft_{(\max_j \text{size}(\bar{\tau}_j)) - \text{size}(\bar{\tau}_i)} \vdash F_i) \vee (\mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}, p \triangleleft \mathbf{own}_n \Sigma \bar{\tau} \vdash F_i)}{\mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}, p \triangleleft \mathbf{own}_n \Sigma \bar{\tau} \vdash \mathbf{case}^* p \mathbf{of} \, \bar{F}} \\
\\
\text{F-CASE-BOR} \quad \frac{\mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive} \quad \forall i. (\mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}, p.1 \triangleleft \&_{\mu}^{\kappa} \tau_i \vdash F_i) \vee (\mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}, p \triangleleft \&_{\mu}^{\kappa} \Sigma \bar{\tau} \vdash F_i)}{\mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}, p \triangleleft \&_{\mu}^{\kappa} \Sigma \bar{\tau} \vdash \mathbf{case}^* p \mathbf{of} \, \bar{F}}
\end{array}$$

Type writing

$$\boxed{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \tau_1 \multimap^{\tau} \tau_2}$$

$$\begin{array}{c}
\text{TWRITE-OWN} \\
\frac{\text{size}(\tau) = \text{size}(\tau')}{\mathbf{E}; \mathbf{L} \vdash \mathbf{own}_n \tau' \multimap^{\tau} \mathbf{own}_n \tau}
\end{array}$$

$$\begin{array}{c}
\text{TWRITE-BOR} \\
\frac{\mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive}}{\mathbf{E}; \mathbf{L} \vdash \&_{\mathbf{mut}}^{\kappa} \tau \multimap^{\tau} \&_{\mathbf{mut}}^{\kappa} \tau}
\end{array}$$

Type reading

$$\boxed{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \tau_1 \multimap^{\tau} \tau_2}$$

$$\begin{array}{c}
\text{TREAD-OWN-COPY} \\
\frac{\tau \text{ copy}}{\mathbf{E}; \mathbf{L} \vdash \mathbf{own}_n \tau \multimap^{\tau} \mathbf{own}_n \tau}
\end{array}$$

$$\begin{array}{c}
\text{TREAD-OWN-MOVE} \\
\frac{n = \text{size}(\tau)}{\mathbf{E}; \mathbf{L} \vdash \mathbf{own}_m \tau \multimap^{\tau} \mathbf{own}_m \not\triangleleft n}
\end{array}$$

$$\begin{array}{c}
\text{TREAD-BOR} \\
\frac{\tau \text{ copy} \quad \mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive}}{\mathbf{E}; \mathbf{L} \vdash \&_{\mu}^{\kappa} \tau \multimap^{\tau} \&_{\mu}^{\kappa} \tau}
\end{array}$$

Well-typed instructions

$$\boxed{\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{T}_1 \vdash I \dashv x. \mathbf{T}_2}$$

$$\begin{array}{c}
\text{S-TRUE} \\
\mathbf{E}; \mathbf{L} \mid \emptyset \vdash \mathbf{true} \dashv x. x \triangleleft \mathbf{bool}
\end{array}$$

$$\begin{array}{c}
\text{S-FALSE} \\
\mathbf{E}; \mathbf{L} \mid \emptyset \vdash \mathbf{false} \dashv x. x \triangleleft \mathbf{bool}
\end{array}$$

$$\begin{array}{c}
\text{S-NUM} \\
\mathbf{E}; \mathbf{L} \mid \emptyset \vdash z \dashv x. x \triangleleft \mathbf{int}
\end{array}$$

$$\begin{array}{c}
\text{S-FN} \\
\frac{\Gamma, \bar{\alpha}, _F : \mathbf{lft}, f, \bar{x}, k : \mathbf{val} \mid \mathbf{E}, \mathbf{E}'; _F \sqsubseteq_1 [] \mid k \triangleleft \mathbf{cont}(_F \sqsubseteq_1 []; y. y \triangleleft \mathbf{own} \tau); \quad \bar{p} \triangleleft \bar{\tau}', \bar{x} \triangleleft \mathbf{own} \bar{\tau}, f \triangleleft \forall \bar{\alpha}. \mathbf{fn}(_F : \mathbf{E}; \bar{\tau}) \rightarrow \tau \vdash F}{\Gamma \mid \mathbf{E}'; \mathbf{L}' \mid \bar{p} \triangleleft \bar{\tau}' \vdash \mathbf{funrec} f(\bar{x}) \mathbf{ret} k := F \vdash f. f \triangleleft \forall \bar{\alpha}. \mathbf{fn}(_F : \mathbf{E}; \bar{\tau}) \rightarrow \tau} \\
\\
\begin{array}{cc}
\text{S-PATH} & \text{S-NAT-OP} \\
\mathbf{E}; \mathbf{L} \mid p \triangleleft \tau \vdash p \dashv x. x \triangleleft \tau & \mathbf{E}; \mathbf{L} \mid p_1 \triangleleft \mathbf{int}, p_2 \triangleleft \mathbf{int} \vdash p_1 \{+, -\} p_2 \dashv x. x \triangleleft \mathbf{int}
\end{array} \\
\\
\begin{array}{cc}
\text{S-NAT-LEQ} & \text{S-NEW} \\
\mathbf{E}; \mathbf{L} \mid p_1 \triangleleft \mathbf{int}, p_2 \triangleleft \mathbf{int} \vdash p_1 \leq p_2 \dashv x. x \triangleleft \mathbf{bool} & \mathbf{E}; \mathbf{L} \mid \emptyset \vdash \mathbf{new}(n) \dashv x. x \triangleleft \mathbf{own}_n \not\leq n
\end{array} \\
\\
\begin{array}{cc}
\text{S-DELETE} & \text{S-DEREF} \\
\frac{n = \mathbf{size}(\tau)}{\mathbf{E}; \mathbf{L} \mid p \triangleleft \mathbf{own}_n \tau \vdash \mathbf{delete}(n, p) \dashv \emptyset} & \frac{\mathbf{E}; \mathbf{L} \vdash \tau_1 \multimap^\tau \tau'_1 \quad \mathbf{size}(\tau) = 1}{\mathbf{E}; \mathbf{L} \mid p \triangleleft \tau_1 \vdash *p \dashv x. p \triangleleft \tau'_1, x \triangleleft \tau}
\end{array} \\
\\
\begin{array}{cc}
\text{S-DEREF-BOR-OWN} & \text{S-DEREF-BOR-BOR} \\
\frac{\mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive}}{\mathbf{E}; \mathbf{L} \mid p \triangleleft \&_\mu^\kappa \mathbf{own}_n \tau \vdash *p \dashv x. x \triangleleft \&_\mu^\kappa \tau} & \frac{\mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive} \quad \mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \kappa'}{\mathbf{E}; \mathbf{L} \mid p \triangleleft \&_\mu^\kappa \&_{\mathbf{mut}}^{\kappa'} \tau \vdash *p \dashv x. x \triangleleft \&_\mu^\kappa \tau}
\end{array} \\
\\
\begin{array}{cc}
\text{S-ASSGN} & \text{S-SUM-ASSGN-UNIT} \\
\frac{\mathbf{E}; \mathbf{L} \vdash \tau_1 \multimap^\tau \tau'_1}{\mathbf{E}; \mathbf{L} \mid p_1 \triangleleft \tau_1, p_2 \triangleleft \tau \vdash p_1 := p_2 \dashv p_1 \triangleleft \tau'_1} & \frac{\bar{\tau}_i = \Pi[] \quad \mathbf{E}; \mathbf{L} \vdash \tau_1 \multimap^{\Sigma \bar{\tau}} \tau'_1}{\mathbf{E}; \mathbf{L} \mid p \triangleleft \tau_1 \vdash p \stackrel{\text{inj } i}{=} () \dashv p \triangleleft \tau'_1}
\end{array} \\
\\
\text{S-SUM-ASSGN} \\
\frac{\bar{\tau}_i = \tau \quad \tau_1 \multimap^{\Sigma \bar{\tau}} \tau'_1}{\mathbf{E}; \mathbf{L} \mid p_1 \triangleleft \tau_1, p_2 \triangleleft \tau \vdash p_1 \stackrel{\text{inj } i}{=} p_2 \dashv p_1 \triangleleft \tau'_1} \\
\\
\text{S-MEMCPY} \\
\frac{\mathbf{size}(\tau) = n \quad \mathbf{E}; \mathbf{L} \vdash \tau_1 \multimap^\tau \tau'_1 \quad \mathbf{E}; \mathbf{L} \vdash \tau_2 \multimap^\tau \tau'_2}{\mathbf{E}; \mathbf{L} \mid p_1 \triangleleft \tau_1, p_2 \triangleleft \tau_2 \vdash p_1 :=_n *p_2 \dashv p_1 \triangleleft \tau'_1, p_2 \triangleleft \tau'_2} \\
\\
\text{S-SUM-MEMCPY} \\
\frac{\mathbf{size}(\tau) = n \quad \mathbf{E}; \mathbf{L} \vdash \tau_1 \multimap^{\Sigma \bar{\tau}} \tau'_1 \quad \mathbf{E}; \mathbf{L} \vdash \tau_2 \multimap^\tau \tau'_2 \quad \bar{\tau}_i = \tau}{\mathbf{E}; \mathbf{L} \mid p_1 \triangleleft \tau_1, p_2 \triangleleft \tau_2 \vdash p_1 \stackrel{\text{inj } i}{=} _n *p_2 \dashv p_1 \triangleleft \tau'_1, p_2 \triangleleft \tau'_2}
\end{array}$$

2 Some examples

This section contains some manually type-checked functions demonstrating how the type system looks like in action.

We write $\tau_1 \times \tau_2 \times \dots$ for $\Pi \bar{\tau}$, **unit** for $\Pi[]$, $\tau_1 + \tau_2 + \dots$ for $\Sigma \bar{\tau}$ and **!** for $\Sigma[]$. We use \downarrow as sugar for \downarrow_1 . Finally, **own** τ is short for **own**_{size(τ)} τ .

It turns out to be useful to have some syntactic sugar for calling a function and using its return value, for declaring continuations without writing code “backwards”, and for immediately initializing a fresh allocation.

$$\begin{aligned} \text{havecont } k \text{ in } F_1 \text{ wherecont } k(\bar{x}) := F_2 &:= \text{letcont } k(x) := F_2 \text{ in } F_1 \\ \text{letcall } x = f(\bar{p}) \text{ in } F &:= \text{letcont } k(x) := F \text{ in call } f(\bar{p}) \text{ ret } k \\ \text{letalloc } x : \tau := p \text{ in } F &:= \text{let } x = \text{new}(1) \text{ in } x := p \text{ in } F \\ \text{letalloc } x : \tau := *p \text{ in } F &:= \text{let } x = \text{new}(\tau) \text{ in } x :=_{\tau} *p \text{ in } F \end{aligned}$$

Notice that we sometimes use types as subscripts where the syntax expects a number. In this case, we implicitly refer to the size of the type.

The syntactic sugar above enjoys the typing rules below.

$$\begin{aligned} &\text{F-LETCALL} \\ &\frac{\Gamma, x : \mathbf{val} \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}, x \triangleleft \tau \vdash F \quad \mathbf{E}; \mathbf{L} \vdash \mathbf{E}'}{\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}, f \triangleleft \mathbf{fn}(\mathbf{E}'; \bar{\tau}) \rightarrow \tau, \bar{p} \triangleleft \bar{\tau} \vdash \text{letcall } x = f(\bar{p}) \text{ in } F} \\ &\text{F-LETALLOC-ASSGN} \\ &\frac{\Gamma, x : \mathbf{val} \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}, x \triangleleft \mathbf{own} \tau \vdash F}{\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}, p \triangleleft \tau \vdash \text{letalloc } x : \tau := p \text{ in } F} \\ &\text{F-LETALLOC-MEMCPY} \\ &\frac{\tau_1 \circ \tau \quad \tau_2 \quad \Gamma, x : \mathbf{val} \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}, x \triangleleft \mathbf{own} \tau, p \triangleleft \tau_2 \vdash F}{\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}, p \triangleleft \tau_1 \vdash \text{letalloc } x : \tau := *p \text{ in } F} \end{aligned}$$

Example 1: Mutable to shared reference, field reference.

```
1 struct Point { x: i32, y: i32 }
2 fn get_x<'a>(p: &'a mut Point) -> &'a i32 {
3     &(*p).x
4 }
```

The types translate as follows:

$$\begin{aligned} \text{Point} &:= \mathbf{int} \times \mathbf{int} \\ \text{get_x} &:= \forall \alpha. \mathbf{fn}(\alpha \text{ alive}; \mathbf{own} \ \&_{\text{mut}}^{\alpha} \text{Point}) \rightarrow \mathbf{own} \ \&_{\text{shr}}^{\alpha} \mathbf{int} \end{aligned}$$

All lifetime bounds (in particular, the fact that lifetime parameters are alive) all have to be made explicit. Furthermore, all variables and return values are passed via owned pointers.

The code itself translates to:

$$\begin{aligned} \text{funrec } \text{get_x}(p) \text{ ret } \text{ret} &:= \\ \text{let } p' = *p \text{ in letalloc } r : \&_{\text{shr}}^{\alpha} \mathbf{int} &:= p'.0 \text{ in} \\ \text{delete}(\&_{\text{mut}}^{\alpha} \text{Point}, p); \text{jump } \text{ret}(r) & \end{aligned}$$

I will usually try to make the Rust code and the λ_{Rust} code match up in terms of lines, so one line of code in the original function corresponds to one line of code in the translation. At the end, there will always be an additional line deallocating the stack frame and jumping to the return continuation.

Now we can typecheck the function body.

Context: α alive

```
{ret  $\triangleleft$  cont( $r$ .  $r \triangleleft$  own  $\&_{\text{shr}}^{\alpha}$  int);  $p \triangleleft$  own  $\&_{\text{mut}}^{\alpha}$  Point}
  { $p$ } let  $p' = *p$  in { $p' \triangleleft \&_{\text{mut}}^{\alpha}$  Point,  $p \triangleleft$  own  $\frac{1}{2}$ }
  { $p'$ } Split { $p'.0 \triangleleft \&_{\text{mut}}^{\alpha}$  int,  $p'.1 \triangleleft \&_{\text{mut}}^{\alpha}$  int}
  { $p'.0$ } letalloc  $r : \&_{\text{shr}}^{\alpha}$  int :=  $p'.0$  in { $r \triangleleft$  own  $\&_{\text{shr}}^{\alpha}$  int}
  { $p$ ,  $ret$ ,  $r$ } delete( $\&_{\text{mut}}^{\alpha}$  Point,  $p$ );  $ret(r)$ 
```

A quick note on our typing outline conventions: In the very first line, we state the external lifetime context E , which does not change during verification. The line below that state the remaining contexts. The next lines are indented, which means that they describe a *change* in the contexts – the items on the left are used, the items on the right are produced. Formally, such lines correspond to the application of one of the rules for well-typed programs. Used items will, in general, be considered gone because our contexts are substructural. However, if the item permits duplication (*e.g.*, for $x \triangleleft \tau$ for τ copy), the item will implicitly be duplicated. When specifying items of the form to be consumed, we will often shorten $X \triangleleft _$ to just X . This still uniquely identifies the used-up context item.

This choice of notation prevents repeating the same items over and over; however, the current context is fairly implicit: the context consists of all the items that are produced by any preceeding step, minus consumed non-duplicable items. When appropriate, we will repeat the entire current context in an un-indented line to keep things clearer.

Example 2: Copying a reference out of ownership

```
1 fn rebor(mut t1: Point, mut t2: Point) -> i32 {
2   let mut x = &mut t1;
3   let y = &mut (*x).y;
4   x = &mut t2;
5   *y
6 }
```

In the code above, dereferencing x *consumes* that pointer, so its permission is gone – it is now essentially uninitialized. This is why x is wrapped in curly braces: if we had just written $(*x).y$, Rust would instead have *re-borrowed* x so the program would not typecheck.

$$\text{rebor} := \text{fn}(\text{own Point}, \text{own Point}) \rightarrow \text{own int}$$

```

funrec rebor( $t1, t2$ ) ret  $ret :=$ 
  newlft;
  letalloc  $x : \&_{mut}^{\alpha} Point := t1$ ;
  let  $x' = *x$  in let  $y = x'.0$  in
   $x := t2$  in
  let  $y' = *y$  in letalloc  $r : int := y'$  in
  endlft; delete( $Point, t1$ ); delete( $Point, t2$ ); delete( $\&_{mut}^{\alpha} Point, x$ ); jump  $ret(r)$ 

{  $ret \triangleleft cont(r, r \triangleleft own\ int)$ ;  $t1 \triangleleft own\ Point, t2 \triangleleft own\ Point$  }
  newlft;
Local lifetimes:  $\alpha \sqsubseteq []$ 
  {  $t1, t2$  } Borrow at  $\alpha$  {  $t1 \triangleleft \&_{mut}^{\alpha} Point, t1 \triangleleft^{\dagger\alpha} own\ Point, t2 \triangleleft \&_{mut}^{\alpha} Point, t2 \triangleleft^{\dagger\alpha} own\ Point$  }
  {  $t1$  } letalloc  $x : \&_{mut}^{\alpha} Point := t1$ ; {  $x \triangleleft own\ \&_{mut}^{\alpha} Point$  }
  {  $x$  } let  $x' = *x$  in {  $x' \triangleleft \&_{mut}^{\alpha} Point, x \triangleleft own\ \frac{1}{2}$  }
  {  $x'$  } let  $y = x'.0$  in {  $y \triangleleft \&_{mut}^{\alpha} int$  }
  {  $x, t2$  }  $x := t2$  in {  $x \triangleleft own\ \&_{mut}^{\alpha} Point$  }
  {  $y$  } let  $y' = *y$  in {  $y' \triangleleft int, y$  }
  {  $y'$  } letalloc  $r : int := y'$  {  $r \triangleleft own\ int$  }
  {  $t1, t2$  } endlft; {  $t1 \triangleleft own\ Point, t2 \triangleleft own\ Point$  }
Local lifetimes:  $\emptyset$ 
  {  $t1, t2, x, ret, r$  } delete( $Point, t1$ ); delete( $Point, t2$ ); delete( $\&_{mut}^{\alpha} Point, x$ ); jump  $ret(r)$ 

```

Example 3: Borrowing from a borrowed box.

Now we can consider this piece of Rust code:

```

1 fn unbox<'a>(b: &'a mut Box<Point>) -> &'a mut u32 {
2   let bx = &mut *b;
3   &mut (*bx).x
4 }

```

Rust boxes translate to owned pointers. In Rust, the only semantic difference between the type T and $Box<T>$ is that the former lives on the stack, while the latter lives on the heap. λ_{Rust} does not distinguish between stack and heap, so the two concepts unify: Box is just **own**.

$$unbox := \forall \alpha. \mathbf{fn}(\alpha \text{ alive}; \mathbf{own} \ \&_{mut}^{\alpha} \mathbf{own} \ Point) \rightarrow \mathbf{own} \ \&_{mut}^{\alpha} \mathbf{int}$$

```

funrec unbox( $b$ ) ret  $ret :=$ 
  let  $b' = *b$  in let  $bx = *b'$  in
  letalloc  $r : \&_{mut}^{\alpha} int := bx.0$  in
  delete( $\&_{mut}^{\alpha} Point, b$ ); delete( $Point, bx$ ); jump  $ret(r)$ 

```


Context: α alive

```
{ret  $\triangleleft$  cont( $r$ .  $r$   $\triangleleft$  own  $\&_{\text{mut}}^{\alpha}$  int);  $b$   $\triangleleft$  own  $\&_{\text{mut}}^{\alpha}$  own Point}
{ $b$ } let  $b' = *b$  in { $b' \triangleleft \&_{\text{mut}}^{\alpha}$  own Point,  $b \triangleleft$  own  $\downarrow$ }
{ $b'$ } let  $bx = *b'$  in { $bx \triangleleft$  own Point}
{ $bx$ } Split { $bx.0 \triangleleft$  ownPoint int,  $bx.1 \triangleleft$  ownPoint int}
{ $bx.0$ } letalloc  $r : \&_{\text{mut}}^{\alpha}$  int :=  $bx.0$  in { $r \triangleleft$  own  $\&_{\text{mut}}^{\alpha}$  int,  $bx.0$ }
{ $bx.0, bx.1$ } Merge { $bx \triangleleft$  own Point}
{ $b, bx, r, \text{retval}$ } delete( $\&_{\text{mut}}^{\alpha}$  Point,  $b$ ); delete(Point,  $bx$ ); jump ret( $r$ )
```

Example 4: Struct initialization.

```
1 fn point(x: i32, y: i32) -> Point {
2   Point { x: x, y: y}
3 }
```

$point := \text{fn}(\text{own int}, \text{own int}) \rightarrow \text{own Point}$

```
funrec point(x, y) ret ret :=
  let  $x' = *x$  in let  $y' = *y$  in
  let  $r = \text{new}(\text{Point})$  in
   $r.0 := x'$ ;  $r.1 := y'$ ;
  delete(int,  $x$ ); delete(int,  $y$ ); jump ret( $r$ )
```

The part about having the lines match up does not really work out here any more...

```
{ret  $\triangleleft$  cont( $r$ .  $r$   $\triangleleft$  own Point);  $x \triangleleft$  own int,  $y \triangleleft$  own int}
{ $x, y$ } let  $x' = *x$  in let  $y' = *y$  in { $x \triangleleft$  own  $\downarrow$ ,  $x' \triangleleft$  int,  $y \triangleleft$  own  $\downarrow$ ,  $y' \triangleleft$  int}
{} let  $r = \text{new}(\text{Point})$  in { $r \triangleleft$  own  $\downarrow$  Point}
{ $r$ } Split { $r.0 \triangleleft$  ownPoint  $\downarrow$ ,  $r.1 \triangleleft$  ownPoint  $\downarrow$ }
{ $r.0, x'$ }  $r.0 := x'$ ; { $r.0 \triangleleft$  ownPoint int,  $x'$ }
{ $r.1, y'$ }  $r.1 := y'$ ; { $r.1 \triangleleft$  ownPoint int,  $y'$ }
{ $r.0, r.1$ } Merge { $r \triangleleft$  own Point}
{ $x, y, \text{ret}, r$ } delete(int,  $x$ ); delete(int,  $y$ ); jump ret( $r$ )
```

Example 5: Enum matching and initialization.

We assume the following *meta-level* type definition:

$Option := \lambda\tau. \text{unit} + \tau$

We will write $Option\langle\tau\rangle$ for τ applied to $Option$.

```
1 fn option_as_mut<T>(o: &mut Option<T>) -> Option<&mut T> {
2   match *o {
3     None => None,
4     Some(ref mut t) => Some(t)
5   }
6 }
```

option_as_mut is parametric over some type τ on the meta-level.

$$option_as_mut\langle\tau\rangle := \forall\alpha. \mathbf{fn}(\alpha \text{ alive}; \mathbf{own} \&_{\mathbf{mut}}^{\alpha} Option\langle\tau\rangle) \rightarrow \mathbf{own} Option\langle\&_{\mathbf{mut}}^{\alpha} \tau\rangle$$

```

funrec option_as_mut(o) ret ret :=
  let r = new(Option<&_{mut}^{\alpha} \tau>) in
havecont k in
  let o' = *o in case *o' of
    - r :=  $\stackrel{\text{inj } 0}{=}$  ();
      jump k()
    - r :=  $\stackrel{\text{inj } 1}{=}$  o'.1;
      jump k()
  wherecont k() :=
    delete(own &_{mut}^{\alpha} Option\langle\tau\rangle, o); jump ret(r)

```

Context: α alive

```

{ret <= cont(r. r <= own Option<&_{mut}^{\alpha} \tau>); o <= own &_{mut}^{\alpha} Option\langle\tau\rangle}
{ } let r = new(Option<&_{mut}^{\alpha} \tau>) in {r <= own &_{mut}^{\alpha} Option\langle\tau\rangle}
havecont k in
{k <= cont(r <= own Option<&_{mut}^{\alpha} \tau>, o <= own &_{mut}^{\alpha} Option\langle\tau\rangle, r <= own &_{mut}^{\alpha} Option\langle\tau\rangle)}
{o } let o' = *o in {o' <= &_{mut}^{\alpha} Option\langle\tau\rangle, o <= own &_{mut}^{\alpha} Option\langle\tau\rangle}
{o' } case *o' of
- {o'.1 <= &_{mut}^{\alpha} unit}
  {r } r :=  $\stackrel{\text{inj } 0}{=}$  () in {r <= own Option<&_{mut}^{\alpha} \tau>}
  {k, o, r } jump k()
- {o'.1 <= &_{mut}^{\alpha} \tau}
  { } r :=  $\stackrel{\text{inj } 1}{=}$  o'.1 in {r <= own Option<&_{mut}^{\alpha} \tau>}
  {k, o, r } jump k()
wherecont k() :=
{k <= cont(r <= own Option<&_{mut}^{\alpha} \tau>, o <= own &_{mut}^{\alpha} Option\langle\tau\rangle)}
{o, ret, r } delete(own &_{mut}^{\alpha} Option\langle\tau\rangle, o); jump ret(r)

```

Example 6: Moving out of an enum.

```

1 fn unwrap_or<T>(o: Option<T>, def: T) -> T {
2   match o {
3     None => def,
4     Some(t) => t
5   }
6 }

```

unwrap_or is parametric over some type τ on the meta-level.

$$unwrap_or := \forall\alpha. \mathbf{fn}\alpha \text{ alive}; \mathbf{own} Option\langle\tau\rangle, \mathbf{own} \tau \rightarrow \tau$$

```

funrec unwrap_or(o, def) ret ret :=
  case *o of
  - delete(Option(τ), o); jump ret(def)
  - letalloc r : τ := o.1 in
    delete(Option(τ), o); delete(τ, def); jump ret(r)

{ret <= cont(r. r <= own τ); o <= own Option(τ), def <= own τ}
{o} case *o of
- {o <= own Option(τ)}
  {o, ret, def} delete(Option(τ), o); jump ret(def)
- {o.0 <= own Option(τ) ↯, o.1 <= own Option(τ) τ}
  {o.1} letalloc r : τ := o.1 in {r <= own τ, o.1 <= own Option(τ) ↯ τ}
  {o.1, o.1} Merge {o <= own ↯ Option(τ)}
  {o, def, ret, r} delete(Option(τ), o); delete(τ, def); jump ret(r)

```

Example 7: Lazy lifetime initialization.

```

1 struct Two<'a> {
2   f: &'a i32,
3   g: &'a i32,
4 }
5
6 fn lazy_lft() {
7   let (mut t, f, g) : (Two, i32, i32);
8   f = 42;
9   t = Two { f: &f, g: &f };
10  *t.f; // The lifetime definitely is already active here
11  g = 23; // And g is definitely not yet borrowed.
12  t.g = &g; // But now we can borrow g at the *old* lifetime.
13 }

```

Let $Two(\kappa) := \&_{\text{shr}}^{\kappa} \text{int} \times \&_{\text{shr}}^{\kappa} \text{int}$.

$lazy_lft := \text{fn}() \rightarrow \text{own unit}$

```

funrec lazy_lft() ret ret :=
  newlft;
  let t = new(2) in let f = new(1) in let g = new(1) in
  f := 42;
  t.0 := f; t.1 := f;
  let t0' = *t.0 in *t0';
  g := 23;
  t.1 := g;
  let r = new(0) in
  endlft; delete(2, t); delete(1, f); delete(1, g); jump ret(r)

```

```

{ret  $\triangleleft$  cont( $r$ .  $r \triangleleft$  own unit}
  newlft;
Local lifetimes:  $\alpha \sqsubseteq []$ 
{ } let  $t = \text{new}(2)$  in let  $f = \text{new}(1)$  in let  $g = \text{new}(1)$  in
{  $t \triangleleft$  own  $\downarrow_2$ ,  $f \triangleleft$  own  $\downarrow_1$ ,  $g \triangleleft$  own  $\downarrow_1$  }
{  $f$  }  $f := 42$ ; {  $f \triangleleft$  own int }
{  $f$  } Borrow at  $\alpha$  {  $f \triangleleft \&_{\text{shr}}^\alpha \text{int}$ ,  $f \triangleleft^{\dagger\alpha} \text{own int}$  }
{  $t$  }  $t.0 := f$ ;  $t.1 := f$ ; {  $t \triangleleft$  own  $\text{Two}(\alpha)$  }
{  $t$  } let  $t0' = *t.0$  in {  $t0' \triangleleft \&_{\text{shr}}^\alpha \text{int}$ ,  $t$  }
{  $t0'$  }  $*t0'$ ; {  $t0'$  }
{  $g$  }  $g := 23$ ; {  $g \triangleleft$  own int }
{  $g$  } Borrow at  $\alpha$  {  $g \triangleleft \&_{\text{shr}}^\alpha \text{int}$ ,  $g \triangleleft^{\dagger\alpha} \text{own int}$  }
{  $g, t$  }  $t.1 := g$ ; {  $g, t$  }
{ } let  $r = \text{new}(0)$  in {  $r \triangleleft$  own unit }
{  $f, g$  } endlft; {  $g \triangleleft$  own int,  $f \triangleleft$  own int }
{  $t, f, g, \text{ret}$  } delete( $2, t$ ); delete( $1, f$ ); delete( $1, g$ ); jump  $\text{ret}(r)$ 

```

3 λ_{Rust} in Iris

Before we get started with the interesting parts, we do some preparatory work on the Iris level to enable reasoning about lambdaRust work. Mostly this is a straight-forward lifting of the operational semantics; the part we need to describe in more details is how we manage the heap.

Physical state. To manage the heap, we use two monoids: Finite partial functions from locations to pairs of lock states and (exclusive) values to talk about ownership of locations (with the instance named γ_{PhVal}), and their contents; and finite partial functions from blocks to (exclusive) tuples for start index and length of the block (instance name γ_{PhFree}). For that to work out, we give a monoid structure to lock states with some unit ε and **reading** $n \cdot \text{reading } m = \text{reading } n + m$.

$$\begin{aligned}
\text{PHVAL} &:= \text{Loc} \xrightarrow{\text{fin}} \text{FRAC}(\text{LockSt} \times \text{AG}(\text{Val})) \\
\text{PHFREE} &:= \mathbb{N} \xrightarrow{\text{fin}} \text{FRAC}(\text{POWFIN}(\mathbb{N})) \\
\ell \xrightarrow{q} v &:= [\circ [\ell \leftarrow q(\text{reading } 0, v)] : \text{AUTH}(\text{PHVAL})] \gamma_{\text{PhVal}} \\
\ell \mapsto v &:= \ell \xrightarrow{1} v \\
\ell \xrightarrow{q} \bar{v} &:= \bigstar_i \ell + i \xrightarrow{q} \bar{v}_i \\
\ell \mapsto \bar{v} &:= \ell \xrightarrow{1} \bar{v} \\
\uparrow_q^m \ell &:= \exists i, n. \ell = (i, n) \wedge [\circ [i \leftarrow q([\geq n, < n + m])]] \gamma_{\text{PhFree}} \\
\text{InvPhys} &:= \exists h, V, F. \text{Phy}(h) * [\bullet V] \gamma_{\text{PhVal}} * [\bullet F] \gamma_{\text{PhFree}} * \\
&\quad h = V * (\forall i. \text{dom}(h) \cap \{i\} \times \mathbb{N} = F(i))
\end{aligned}$$

where $\uparrow_1^m \ell$ is the permission to deallocate ℓ as block of length m . If the fraction is less than 1, this means that we don't own the entire block yet, and have to obtain more permissions before we can deallocate anything. We assume a global invariant namespace \mathcal{N}_{Ph} , and we assume $\boxed{\text{InvPhys}}^{\mathcal{N}_{\text{Ph}}}$ to be in the global context.

We obtain the usual triples for both atomic and non-atomic memory loads and stores, and (de)allocation, and some useful separations:

$$\begin{aligned}
\ell \xrightarrow{q} v * \ell \xrightarrow{q'} v' &\Leftrightarrow \ell \xrightarrow{q+q'} v * v = v' & \uparrow_q^m \ell * \uparrow_{q'}^{m'} \ell + m &\Leftrightarrow \uparrow_{q+q'}^{m+m'} \ell \\
\{\text{True}\} \text{alloc}(n) \{ \ell. \exists \bar{v}. \ell \mapsto \bar{v} * |\bar{v}| = n * \uparrow_1^n \ell \}_{\mathcal{N}_{\text{Ph}}} & & \{ \ell \mapsto \bar{v} * \uparrow_1^{|\bar{v}|} \ell \} \text{free}(|\bar{v}|, v) \{ \text{True} \}_{\mathcal{N}_{\text{Ph}}} \\
\{ \ell \xrightarrow{q} v \} *_{\text{sc}} \ell \{ v'. v' = v * \ell \xrightarrow{q} v \}_{\mathcal{N}_{\text{Ph}}} & & \{ \ell \xrightarrow{q} v \} * \ell \{ v'. v' = v * \ell \xrightarrow{q} v \}_{\mathcal{N}_{\text{Ph}}} \\
\{ \ell \mapsto v \} \ell :=_{\text{sc}} w \{ v'. v' = () * \ell \mapsto w \}_{\mathcal{N}_{\text{Ph}}} & & \{ \ell \mapsto v \} \ell := w \{ v'. v' = () * \ell \mapsto w \}_{\mathcal{N}_{\text{Ph}}} \\
\frac{|\bar{v}_1| = |\bar{v}_2| = n}{\{ \ell_1 \mapsto \bar{v}_1 * \ell_2 \xrightarrow{q} \bar{v}_2 \} \ell_1 :=_n * \ell_2 \{ \ell_1 \mapsto \bar{v}_2 * \ell_2 \xrightarrow{q} \bar{v}_2 \}_{\mathcal{N}_{\text{Ph}}}}
\end{aligned}$$

4 Lifetime logic

The core principle of lifetimes and borrows has applications beyond type systems. In the following, we develop a logic that includes primitives dealing with lifetimes, using the λ_{Rust} type system as a sample application.

4.1 Proof rules

Intuitively speaking, why ought a type system like the one in §1.4 be sound? What justifies doing a borrow, using that borrow, and obtaining ownership of the original permission again when the borrow ends? The purpose of this section is to develop an intuition for the proof rules of Figure 3, which are used in the soundness proof of the type system. These rules describe the *lifetime logic*.

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-CONS**. The rule is quite a mouthful, so it is worth looking at the following simpler (derived) version:

$$\langle \&_{\text{full}}^{\kappa} P * [\kappa]_q \rangle \Leftrightarrow \triangleright P \rangle_{\mathcal{N}_{\text{ift}}} \quad (1)$$

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-CONS** 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 3: 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
$[\text{Bor} : i]_q$	Ownership of fraction q of the indexed borrow i	Yes	No

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

$$\kappa \sqsubseteq \kappa' := (\forall q. \langle [\kappa]_q \Leftrightarrow q' \cdot [\kappa']_{q'} \rangle_{\mathcal{N}_{\text{lft}}} * ([\dagger\kappa'] \Rightarrow_{\mathcal{N}_{\text{lft}}} [\dagger\kappa]))$$

Lifetime creation and end.

LFTL-BEGIN $\text{True} \Rightarrow_{\mathcal{N}_{\text{lft}}} \exists \kappa. [\kappa]_1 * \square([\kappa]_1 \Rightarrow_{\emptyset}^{\mathcal{N}_{\text{lft}}} [\dagger\kappa])$	LFTL-TOK-FRACT $[\kappa]_{q+q'} \Leftrightarrow [\kappa]_q * [\kappa]_{q'}$	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 $\triangleright P \Rightarrow_{\mathcal{N}_{\text{lft}}} \&_{\text{full}}^\kappa P * ([\dagger\kappa] \Rightarrow_{\mathcal{N}_{\text{lft}}}^* \triangleright P)$	LFTL-BOR-SEP $\&_{\text{full}}^\kappa (P * Q) \Leftrightarrow_{\mathcal{N}_{\text{lft}}} \&_{\text{full}}^\kappa P * \&_{\text{full}}^\kappa Q$
LFTL-BOR-FAKE $[\dagger\kappa] \Rightarrow_{\mathcal{N}_{\text{lft}}} \&_{\text{full}}^\kappa P$	
LFTL-BOR-ACC-STRONG $\&_{\text{full}}^\kappa P * [\kappa]_q \Rightarrow_{\mathcal{N}_{\text{lft}}} \exists \kappa'. \kappa \sqsubseteq \kappa' * \triangleright P * (\forall Q. \triangleright (\triangleright Q * [\dagger\kappa'] \Rightarrow_{\emptyset}^* \triangleright P) * \triangleright Q \Rightarrow_{\mathcal{N}_{\text{lft}}}^* \&_{\text{full}}^{\kappa'} Q * [\kappa]_q)$	
LFTL-BOR-ACC-ATOMIC-STRONG $\&_{\text{full}}^\kappa P \xRightarrow{\mathcal{N}_{\text{lft}}} \emptyset \left(\exists \kappa'. \kappa \sqsubseteq \kappa' * \triangleright P * (\forall Q. \triangleright (\triangleright Q * [\dagger\kappa'] \Rightarrow_{\emptyset}^* \triangleright P) * \triangleright Q \Rightarrow_{\mathcal{N}_{\text{lft}}}^* \&_{\text{full}}^{\kappa'} Q) \right) \vee ([\dagger\kappa] * \emptyset \Rightarrow_{\mathcal{N}_{\text{lft}}} \text{True})$	

Indexed borrows.

LFTL-BOR-IDX $\&_{\text{full}}^\kappa P \Leftrightarrow \exists i. \&_i^\kappa P * [\text{Bor} : i]_1$	LFTL-BOR-FRACT $[\text{Bor} : i]_{q+q'} \Leftrightarrow [\text{Bor} : i]_q * [\text{Bor} : i]_{q'}$	LFTL-IDX-SHORTEN $\frac{\kappa' \sqsubseteq \kappa}{\&_i^\kappa P \Rightarrow \&_i^{\kappa'} P}$
LFTL-IDX-ACC $\&_i^\kappa P \vdash \langle [\text{Bor} : i]_1 * [\kappa]_q \Leftrightarrow \triangleright P \rangle_{\mathcal{N}_{\text{lft}}}$	LFTL-IDX-ACC-ATOMIC $\&_i^\kappa P \vdash \langle [\text{Bor} : i]_q \Leftrightarrow b. \text{if } b \text{ then } \triangleright P \text{ else } [\dagger\kappa] \rangle_{\mathcal{N}_{\text{lft}}}^\emptyset$	
23		
LFTL-IDX-BOR-UNNEST $\&_i^\kappa P * \&_{\text{full}}^{\kappa'}([\text{Bor} : i]_1) \Rightarrow_{\mathcal{N}_{\text{lft}}} \&_{\text{full}}^{\kappa \sqcap \kappa'} P$		

Finally, the rule **LFTL-BOR-ACC-ATOMIC-CONS** provides a way to access a full borrow *without* having a proof that the lifetime is still ongoing. The rule is again fairly involved due to incorporating a rule of consequence, so let us consider the following simpler (derived) version:

$$\langle \&_{\text{full}}^{\kappa} P \Leftrightarrow b. \text{if } b \text{ then } \triangleright P \text{ else } [\dagger\kappa] \rangle_{\mathcal{N}_{\text{ift}}}^{\emptyset} \quad (2)$$

The key differences to is that we do *not* need to provide $[\kappa]_q$. As a consequence, (a) the accessor is mask-changing, so it can only be used atomically, and (b) we may get *either* $\triangleright P$ (the content of the full borrow) *or* a proof that κ has, in fact, already ended.

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:

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

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' . Furthermore, tokens showing that κ' has ended must be convertible to tokens showing that κ has ended. 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. They can also be accessed atomically without a lifetime token **LFTL-IDX-ACC-ATOMIC**. The latter rule further shows that we actually only need *any fraction* of the borrow token to perform an atomic access, thus providing a way of sharing borrows (by distributing their fractions). 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.

4.2 Derived forms of borrowing

Figure 4 shows some rules that can be derived from the basic rules discussed in the previous subsection. **LFTL-BOR-FREEZE** is a very interesting derived rule, and it also demonstrates the full power of **LFTL-BOR-ACC-ATOMIC-CONS**, so we will briefly discuss it. After applying **LFTL-BOR-ACC-ATOMIC-CONS**, we distinguish two cases. Either we get $\triangleright \exists x \in \tau. P$. We then move the \triangleright down into the existential and destruct the latter to obtain an x and $\triangleright P$. We pick $Q := P$ and trivially show that $P * \exists x \in \tau. P$. Then we run the closing view shift to finish the proof. In the other case, all we

$$\begin{array}{c}
\text{LFTL-INCL-ISECT} \\
\kappa \sqcap \kappa' \sqsubseteq \kappa
\end{array}
\quad
\begin{array}{c}
\text{LFTL-INCL-GLB} \\
\frac{\kappa \sqsubseteq \kappa' \quad \kappa \sqsubseteq \kappa''}{\kappa \sqsubseteq \kappa' \sqcap \kappa''}
\end{array}
\quad
\begin{array}{c}
\text{LFTL-FRACT-LINCL} \\
\frac{\&_{\text{frac}}^{\kappa} q' \cdot [\kappa']_{q \cdot q'}}{\kappa \sqsubseteq \kappa'}
\end{array}
\quad
\begin{array}{c}
\text{LFTL-BOR-SHORTEN} \\
\frac{\kappa' \sqsubseteq \kappa}{\&_{\text{full}}^{\kappa} P \Rightarrow \&_{\text{full}}^{\kappa'} P}
\end{array}$$

$$\begin{array}{c}
\text{LFTL-REBORROW} \\
\kappa' \sqsubseteq \kappa \vdash \&_{\text{full}}^{\kappa} P \Rightarrow_{\mathcal{N}_{\text{ift}}} \&_{\text{full}}^{\kappa'} P * ([\dagger \kappa'] \Rightarrow_{\mathcal{N}_{\text{ift}}}^* \&_{\text{full}}^{\kappa} P)
\end{array}
\quad
\begin{array}{c}
\text{LFTL-BOR-UNNEST} \\
\&_{\text{full}}^{\kappa'} (\&_{\text{full}}^{\kappa} P) \Rightarrow_{\mathcal{N}_{\text{ift}}}^* \&_{\text{full}}^{\kappa \sqcap \kappa'} P
\end{array}$$

$$\begin{array}{c}
\text{LFTL-BOR-ACC-CONS} \\
\&_{\text{full}}^{\kappa} P * [\kappa]_q \Rightarrow_{\mathcal{N}_{\text{ift}}} \triangleright P * \forall Q. \triangleright (\triangleright Q \Rightarrow_{\emptyset}^* \triangleright P) * \triangleright Q \Rightarrow_{\mathcal{N}_{\text{ift}}}^* \&_{\text{full}}^{\kappa} Q * [\kappa]_q
\end{array}$$

$$\begin{array}{c}
\text{LFTL-BOR-ACC} \\
\langle [\kappa]_q * \&_{\text{full}}^{\kappa} P \Leftrightarrow \triangleright P \rangle_{\mathcal{N}_{\text{ift}}}
\end{array}
\quad
\begin{array}{c}
\text{LFTL-BOR-ACC-ATOMIC} \\
\langle \&_{\text{full}}^{\kappa} P \Leftrightarrow b. \text{ if } b \text{ then } \triangleright P \text{ else } [\dagger \kappa] \rangle_{\mathcal{N}_{\text{ift}}}^{\emptyset}
\end{array}$$

$$\begin{array}{c}
\text{LFTL-BOR-ACC-ATOMIC-CONS} \\
\&_{\text{full}}^{\kappa} P \xrightarrow{\mathcal{N}_{\text{ift}}}^{\emptyset} \left(\triangleright P * \forall Q. \triangleright (\triangleright Q \Rightarrow_{\emptyset}^* \triangleright P) * \triangleright Q \xrightarrow{\mathcal{N}_{\text{ift}}}^{\emptyset} \&_{\text{full}}^{\kappa} Q \right) \vee [\dagger \kappa] * \emptyset \xrightarrow{\mathcal{N}_{\text{ift}}} \text{True}
\end{array}$$

$$\begin{array}{c}
\text{LFTL-BOR-FREEZE} \\
\frac{\tau \text{ inhabited}}{(\&_{\text{full}}^{\kappa} \exists x : \tau. P) \Rightarrow_{\mathcal{N}_{\text{ift}}} \exists x : \tau. \&_{\text{full}}^{\kappa} P}
\end{array}$$

Figure 4: Lifetime logic derived rules

get is $[\dagger \kappa]$. We run the closing view shift, and now we use the fact that τ is inhabited to obtain some x . We pick that x as the witness for our goal, and then use **LFTL-BOR-FAKE** to finish the proof.

Furthermore, we introduce in **Figure 5** 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 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

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

$$\begin{array}{c}
\text{LFTL-BOR-AT} \\
\mathcal{N} \# \mathcal{N}_{\text{ift}} \vdash \&_{\mathbf{full}}^{\kappa} P \Rightarrow_{\mathcal{N}_{\text{ift}}} \&_{\mathbf{at}}^{\kappa/\mathcal{N}} P \\
\\
\text{LFTL-AT-ACC-ATOMIC} \\
\&_{\mathbf{at}}^{\kappa/\mathcal{N}} P \vdash \langle \text{True} \Leftrightarrow b. \text{if } b \text{ then } \triangleright P \text{ else } [\dagger\kappa] \rangle_{\mathcal{N}_{\text{ift}}, \mathcal{N}}^{\emptyset} \\
\\
\text{LFTL-AT-SHORTEN} \\
\frac{\kappa' \sqsubseteq \kappa}{\&_{\mathbf{at}}^{\kappa/\mathcal{N}} P \Rightarrow \&_{\mathbf{at}}^{\kappa'/\mathcal{N}} P} \\
\\
\text{LFTL-BOR-LFTNAMESP} \\
\&_{\mathbf{full}}^{\kappa} P \Rightarrow_{\mathcal{N}_{\text{ift}}} \&_{\mathbf{at}}^{\kappa/\mathcal{N}_{\text{ift}}} P \\
\\
\text{LFTL-AT-ACC} \\
\&_{\mathbf{at}}^{\kappa/\mathcal{N}} P \vdash \langle [\kappa]_q \Leftrightarrow \triangleright P \rangle_{\mathcal{N}_{\text{ift}}, \mathcal{N}}^{\mathcal{N}_{\text{ift}} - \mathcal{N}}
\end{array}$$

Non-atomic persistent borrows

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

Fractured borrows

$$\begin{array}{c}
\text{LFTL-BOR-FRACTURE} \\
\&_{\mathbf{full}}^{\kappa} \Phi(1) \Rightarrow_{\mathcal{N}_{\text{ift}}} \&_{\mathbf{frac}}^{\kappa} \Phi \\
\\
\text{LFTL-FRACT-ACC} \\
\frac{\forall q_1, q_2. \Phi(q_1 + q_2) \Leftrightarrow \Phi(q_1) * \Phi(q_2)}{\&_{\mathbf{frac}}^{\kappa} \Phi \vdash \langle [\kappa]_q \Leftrightarrow q'. \triangleright \Phi(q') \rangle_{\mathcal{N}_{\text{ift}}}} \\
\\
\text{LFTL-FRACT-ACC-ATOMIC} \\
\&_{\mathbf{frac}}^{\kappa} \Phi \vdash \langle \text{True} \Leftrightarrow (b, q). \text{if } b \text{ then } \triangleright \Phi(q) \text{ else } [\dagger\kappa] \rangle_{\mathcal{N}_{\text{ift}}}^{\emptyset} \\
\\
\text{LFTL-FRACT-SHORTEN} \\
\frac{\kappa' \sqsubseteq \kappa}{\&_{\mathbf{frac}}^{\kappa} \Phi \Rightarrow \&_{\mathbf{frac}}^{\kappa'} \Phi}
\end{array}$$

Figure 5: Lifetime logic derived forms

be defined as follows:

$$\&_{\text{at}}^{\kappa/\mathcal{N}} P := \exists i. \&_i^{\kappa} P * (\mathcal{N} \# \mathcal{N}_{\text{lft}} * [\text{Bor} : i]_1)^{\mathcal{N}} \vee \mathcal{N} = \mathcal{N}_{\text{lft}} * [\exists q. [\text{Bor} : i]_q]^{\mathcal{N}_{\text{lft}}}$$

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. The persistent borrows of this other flavor are called *non-atomic persistent borrows*. They can be defined by:

$$\&_{\text{na}}^{\kappa/p.\mathcal{N}} P := \exists i. \&_i^{\kappa} P * \text{Nalnv}^{p.\mathcal{N}}([\text{Bor} : i]_1)$$

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.

The way this works is that we have an atomic persistent borrow which contains some fraction of Φ , and some fraction of the lifetime token, such that the two fractions add up to 1. When the lifetime is ended, the full token of one of the intersected lifetimes is used up, so there cannot be any piece of the lifetime token within the fractured borrow – so the full $\Phi(1)$ is available. The rule of consequence **LFTL-BOR-ACC-STRONG** witnesses this fact by providing $[\dagger\kappa]$ to the view shift that is applied.

Fractured borrows can be defined as follows:

$$\&_{\text{frac}}^{\kappa} \Phi := \exists \kappa', \gamma. \kappa \sqsubseteq \kappa' * \&_{\text{at}}^{\kappa'/\mathcal{N}_{\text{lft}}} \exists q. \Phi(q) * [\dagger\kappa']^{\gamma} * (q = 1 \vee [\kappa']_{1-q})$$

Here, are using the FRAC RA.

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

4.3 Model

We will model lifetimes κ as multisets of *atomic lifetimes* $\Lambda \in \mathbb{N}$, which are just identifiers. This forms a cancellable positive commutative monoid with union for composition and \emptyset as the unit.

We will need the following datatypes and CMRAs:

$$\begin{aligned} \text{BorSt} &:= \text{in} + \text{open}(q) + \text{rebor}(\kappa) \\ \text{LftSt} &:= \text{LftStAlive} + \text{LftStDead} \\ \text{ALFT} &:= \text{AUTH}(\mathbb{N} \xrightarrow{\text{fin}} \text{FRAC} +_{\dagger} ()) \\ \text{ILFT} &:= \text{AUTH}(\wp^{\text{fin},+}(\mathbb{N}) \xrightarrow{\text{fin}} \text{AG}(\mathcal{G} \times \mathcal{G} \times \mathcal{G})) \\ \text{BORBOX} &:= \text{AUTH}(\mathbb{N} \xrightarrow{\text{fin}} \text{AG}(\text{BorSt}) \times \text{FRAC}) \\ \text{CNT} &:= \text{AUTH}(\mathbb{N}) \\ \text{INHBOX} &:= \text{AUTH}(\wp^{\text{fin}}(\mathbb{N})) \end{aligned}$$

We assume some globally known indices γ_a and γ_i for managing atomic and intersected lifetimes. The two tokens of the lifetime logic are easily modelled:

$$\begin{aligned} [\kappa]_q &:= \bigstar_{\Lambda \in \kappa} [\circ [\Lambda \leftarrow \text{inl}(q)]]^{\gamma_a} \\ [\dagger \kappa] &:= \exists \Lambda \in \kappa. [\circ [\Lambda \leftarrow \text{inr}()]]^{\gamma_a} \end{aligned}$$

We will use the following notation for a view shift that frames assertion P_F :

$$P \equiv \bigstar_{[P_F]_{\mathcal{E}}} Q := P * P_F \equiv \bigstar_{\mathcal{E}} Q * P_F$$

ILFT manages the intersected lifetimes; it just records the ghost names of the state managing those lifetimes. To simplify working with this indirection, we define:

$$\begin{aligned} \text{OwnBor}(\kappa, x) &:= \exists \gamma_{\text{bor}}. [\circ [\kappa \leftarrow \gamma_{\text{bor}}, _]]^{\gamma_i} * [x]^{\gamma_{\text{bor}}} \\ \text{OwnCnt}(\kappa, x) &:= \exists \gamma_{\text{cnt}}. [\circ [\kappa \leftarrow _ , \gamma_{\text{cnt}}, _]]^{\gamma_i} * [x]^{\gamma_{\text{cnt}}} \\ \text{OwnInh}(\kappa, x) &:= \exists \gamma_{\text{inh}}. [\circ [\kappa \leftarrow _ , _ , \gamma_{\text{inh}}]]^{\gamma_i} * [x]^{\gamma_{\text{inh}}} \end{aligned}$$

Now we can define the core of the model: the protocols for alive and dead lifetimes. We split the namespace \mathcal{N}_{ft} into three disjoint sub-namespaces $\mathcal{N}_{\text{mgmt}}$, \mathcal{N}_{bor} , \mathcal{N}_{inh} . Here, $A : \mathbb{N} \xrightarrow{\text{fin}} \mathbf{LftSt}$ is a map indicating the state of the atomic lifetimes, and $I : \wp^{\text{fin},+}(\mathbb{N}) \xrightarrow{\text{fin}} \mathcal{G} \times \mathcal{G} \times \mathcal{G}$ indicates which intersected lifetimes exist.

$$\text{bor_to_box}(s) := \begin{cases} \text{full} & \text{if } s = \text{in} \\ \text{empty} & \text{otherwise} \end{cases}$$

$$\text{LftBorAlive}(\kappa, P_B) := \exists B : \mathbb{N} \xrightarrow{\text{fin}} \text{BorSt}. \text{OwnBor}(\kappa, \bullet [i \leftarrow (B(i), 1) \mid i \in \text{dom}(B)]) * \text{Box}(\mathcal{N}_{\text{bor}}, P_B, [i \leftarrow \text{bor_to_box}(B(i)) \mid i \in \text{dom}(B)]) *$$

$$\bigstar_{i \in \text{dom}(B)} \begin{cases} \text{True} & \text{if } B(i) = \text{in} \\ [\kappa]_q & \text{if } B(i) = \text{open}(q) \\ \text{OwnCnt}(\kappa', \circ 1) * \kappa \subset \kappa' & \text{if } B(i) = \text{rebor}(\kappa') \end{cases}$$

$$\text{LftBorDead}(\kappa) := \exists B : \wp^{\text{fin}}(\mathbb{N}), P_B : \text{Prop}. \text{OwnBor}(\kappa, \bullet [i \leftarrow (\text{in}, 1) \mid i \in B]) * \text{Box}(\mathcal{N}_{\text{bor}}, P_B, [i \leftarrow \text{empty} \mid i \in \text{dom}(B)])$$

$$\text{LftInh}(\kappa, P_I, s) := \exists E : \wp^{\text{fin}}(\mathbb{N}). \text{OwnInh}(\kappa, \bullet E) * \text{Box}(\mathcal{N}_{\text{inh}}, P_I, [i \leftarrow s \mid i \in E])$$

$$\text{LftVs}(\kappa, P_B, P_I) := \exists n : \mathbb{N}. \text{OwnCnt}(\kappa, \bullet n) * \forall I : \wp^{\text{fin},+}(\mathbb{N}) \xrightarrow{\text{fin}} \mathcal{G} \times \mathcal{G} \times \mathcal{G}.$$

$$\triangleright P_B * [\dagger \kappa] \equiv \bigstar \left[\text{LftBorDead}(\kappa) * \left[\bullet I \right]^{\gamma_i} * \bigstar_{\substack{\kappa' \in \text{dom}(I) \\ \kappa' \subset \kappa}} \text{LftAlive}(\kappa') \right]_{\mathcal{N}_{\text{bor}}} \triangleright P_I * \text{OwnCnt}(\kappa \circ n)$$

$$\text{LftAlive}(\kappa) := \exists P_B, P_I. \text{LftBorAlive}(\kappa, P_B) * \text{LftVs}(\kappa, P_B, P_I) * \text{LftInh}(\kappa, P_I, \text{empty})$$

$$\text{LftDead}(\kappa) := \exists P_I. \text{LftBorDead}(\kappa) * \text{OwnCnt}(\kappa, \bullet 0) * \text{LftInh}(\kappa, P_I, \text{full})$$

$$\text{LftAliveIn}(A, \kappa) := \forall \Lambda \in \kappa. A(\Lambda) = \text{LftStAlive}$$

$$\text{LftDeadIn}(A, \kappa) := \exists \Lambda \in \kappa. A(\Lambda) = \text{LftStDead}$$

$$\text{LftInv}(A, \kappa) := \text{LftAlive}(\kappa) * \text{LftAliveIn}(A, \kappa) \vee \text{LftDead}(\kappa) * \text{LftDeadIn}(A, \kappa)$$

Notice that LftAlive and LftVs are defined mutually recursively, which is well-defined because the size of the lifetime κ gets strictly smaller.

The rough idea behind this setup is as follows: For every lifetime κ , we have two boxes: one tracking the borrows, and one tracking the inheritances. The latter are used to obtain resources from dead lifetimes, which is necessary to show the view shift obtained via **LFTL-BORROW** and **LFTL-REBORROW**. The borrow box contains assertion P_B , and we have an authoritative map B tracking which slices of the box are full and which are not. There are two ways a slice can be empty: either the borrow currently open, and some fraction of the lifetime token was put in here as a deposit. Alternatively, the borrow can be reborrowed to a strictly shorter lifetime (*i.e.*, a lifetime represented by a strictly larger multiset). Ownership of the fragments of B permit changing the state of the slice or removing it (*e.g.*, for splitting, where one slice is removed and two new ones are added). On the inheritance side, the box overall contains assertion P_I . LftInh ensures that the slices of the box are all in the same state, but there is still an authoritative set that manages ownership of slices for the purpose of removing them from the box. Finally, P_B and P_I are connected through LftVs , which roughly speaking says that one can view shift from P_B to P_I . This view shift is executed when a lifetime ends, after which P_I can be used to fill all the inheritance boxes. All that extra machinery in LftVs is needed to support reborrowing.

Based on this, we define LftLCtx , which we always assume to be in the context for the lifetime

logic rules.

$$\begin{aligned} \text{LftLInv} &:= \exists A, I. [\bullet A]^{\gamma_a} * [\bullet I]^{\gamma_i} * \bigstar_{\kappa \in \text{dom}(I)} \text{LftInv}(A, \kappa) \\ \text{LftLCtx} &:= \boxed{\text{LftLInv}}^{\mathcal{N}_{\text{mgmt}}} \end{aligned}$$

Now we can model the remaining assertions of the logic, and an assertion **RawBor** (“raw borrows”) that will be useful later in the proofs. The indices i in the lifetime logic rules are actually pairs of a lifetime κ' and a box index i .

$$\begin{aligned} \kappa \sqsubseteq \kappa' &:= (\forall q. \langle [\kappa]_q \Leftrightarrow q'. [\kappa']_{q'} \rangle_{\mathcal{N}_{\text{ift}}}) * ([\dagger \kappa'] \Rightarrow_{\mathcal{N}_{\text{ift}}} [\dagger \kappa]) \\ [\text{Bor} : \kappa', i]_q &:= \text{OwnBor}(\kappa, \circ i \leftarrow (\text{in}, q)) \\ \&_{\kappa', i}^{\kappa} P &:= \kappa \sqsubseteq \kappa' * \text{BoxSlice}(\mathcal{N}_{\text{bor}}, P, i) \\ \text{RawBor}(\kappa, P) &:= \exists i. \text{BoxSlice}(\mathcal{N}_{\text{bor}}, P, i) * [\text{Bor} : \kappa, i]_1 \\ \&_{\text{full}}^{\kappa} P &:= \exists \kappa'. \kappa \sqsubseteq \kappa' * \text{RawBor}(\kappa', P) \end{aligned}$$

Some proof rules are trivially justified: **LFTL-TOK-FRACT**, **LFTL-TOK-COMP**, **LFTL-TOK-UNIT**, **LFTL-NOT-OWN-END**, **LFTL-END-COMP**, **LFTL-END-UNIT**, **LFTL-BOR-IDX**, **LFTL-BOR-FRACT**, **LFTL-IDX-SHORTEN** are all simple implications. We briefly discuss the most important steps of most the remaining rules.¹ When we have a full borrow $\&_{\text{full}}^{\kappa} P$, we will call the actual lifetime of the borrow (the one in the existential quantifier) κ_0 . In particular, $\kappa \sqsubseteq \kappa_0$.

Proof sketch of LFTL-BORROW. First we have to check whether κ is already allocated in I ; if not, that’s easy to do:

$$\text{True} \Rightarrow_{\mathcal{N}_{\text{ift}}} [\circ [\kappa \leftarrow \text{---} \text{---} \text{---}]]^{\gamma_i} \quad (3)$$

The proof of this also extends A with new atomic lifetimes as necessary. In the following, we assume to have this in the context.

Next we take a look at $\text{LftInv}(A, \kappa)$. If the lifetime is dead, we use the following to create a “fake” full borrow:

$$\triangleright^b \text{LftBorDead}(\kappa) \Rightarrow_{\mathcal{N}_{\text{bor}}} \triangleright^b \text{LftBorDead}(\kappa) * \text{RawBor}(\kappa, P) \quad (4)$$

$$\triangleright \text{LftInh}(\kappa, P_I, \text{full}) * \triangleright P \Rightarrow_{\mathcal{N}_{\text{inh}}} \triangleright \text{LftInh}(\kappa, P_I * P, \text{full}) * ([\dagger \kappa] \Rightarrow_{\mathcal{N}_{\text{ift}}} \bigstar \triangleright P) \quad (5)$$

(4) just extends the B in LftBorDead with a fresh element and creates an empty slice in the box. (5) allocates a fresh element in E .

If the lifetime is alive, we use:

$$\triangleright \text{LftInh}(\kappa, P_I, \text{empty}) \Rightarrow_{\mathcal{N}_{\text{inh}}} \triangleright \text{LftInh}(\kappa, P_I * P, \text{empty}) * \exists j. \text{OwnInh}(\kappa, \circ \{j\}) * \text{BoxSlice}(\mathcal{N}_{\text{inh}}, P, j) \quad (6)$$

$$\triangleright \text{LftBorAlive}(\kappa, P_B) * \triangleright P \Rightarrow_{\mathcal{N}_{\text{bor}}} \triangleright \text{LftBorAlive}(\kappa, P_B * P) * \&_{\text{full}}^{\kappa} P \quad (7)$$

$$\triangleright \text{LftVs}(\kappa, P_B, P_I) \Rightarrow \triangleright \text{LftVs}(\kappa, P_B * P, P_I * P) \quad (8)$$

¹The proof of reborrowing is not covered here; it can (like the others) be found in the Coq development.

These are all easy consequences of boxing lemmas and allocating in the authoritative B and E . All that is left to do is show the view shift $[\dagger\kappa] \Rightarrow^* \triangleright P$. From $[\dagger\kappa]$ we learn that κ has ended, so we can get access to $\text{LftDead}(\kappa)$. We can now apply the following lemma to the resources we got from (6):

$$\triangleright \text{LftInh}(\kappa, P_I, \text{full}) * \text{OwnInh}(\kappa, \circ \{j\}) * \text{BoxSlice}(\mathcal{N}_{\text{inh}}, P, j) \Rightarrow_{\mathcal{N}_{\text{inh}}} \triangleright P * \exists P'_I. \triangleright \text{LftInh}(\kappa, P'_I, \text{full}) \quad (9)$$

This is shown easily by removing the slice from the box.

Proof sketch of LFTL-BOR-FAKE . This is a trivial consequence of (4).

Proof sketch of LFTL-BOR-ACC-CONS , LFTL-IDX-ACC . These two work fairly similarly. First we run the accessor in $\kappa \sqsubseteq \kappa_0$ to obtain a token for the actual lifetime κ_0 of the borrow. We have a witness of κ_0 being in I , so we can get $\text{LftInv}(A, \kappa_0)$. Since we have a token $[\kappa_0]_{q'}$, we can get $\text{LftAlive}(\kappa_0)$. All we care about is LftBorAlive , using the following lemma:

$$\begin{aligned} \text{BoxSlice}(\mathcal{N}_{\text{bor}}, P, i) \vdash \triangleright \text{LftBorAlive}(\kappa, P_B) * \text{OwnBor}(\kappa, \circ i \leftarrow \text{in}, \text{inl}(\text{ex}())) * [\kappa]_q &\Leftrightarrow_{\mathcal{N}_{\text{bor}}} \quad (10) \\ \triangleright \text{LftBorAlive}(\kappa, P_B) * \text{OwnBor}(\kappa, \circ i \leftarrow \text{open}(q), \text{inl}(\text{ex}())) * \triangleright P \end{aligned}$$

From $B(i) = \text{in}$ we know that the box slice is full. We empty the slice, obtaining $\triangleright P$. Next, we change $B(i)$ to $\text{open}(q)$. This allows us to re-establish LftBorAlive . Similar, for the right-to-left direction, we fill the empty slice and change $B(i)$ back to in .

This already pretty much completes the proof of LFTL-IDX-ACC . For the closing view shift, we follow the same path, using the right-to-left direction of (10).

To finish LFTL-BOR-ACC-CONS , we have to handle the rule of consequence embedded in the closing view shift of the accessor. We start out like above, until we get LftBorAlive . We use remove the empty slice from the borrow box to learn that P_B later decomposes into $P'_B * P$. Since LftBorAlive and LftVs are contractive in P_B , we can rewrite with this decomposition. After adding a new empty slice to the box, we can obtain $\triangleright \text{LftBorAlive}(\kappa_0, P'_B * Q)$ (this also requires removing the old slice from the authoritative map and adding a new one, but we own the fragment, so that's possible). Then we can finish the proof by applying (10) and

$$\triangleright \text{LftVs}(\kappa, P'_B * P) * \triangleright \left(\triangleright Q * [\dagger\kappa] \Rightarrow^*_{-\mathcal{N}_{\text{ft}}} \triangleright P \right) \Rightarrow \triangleright \text{LftVs}(\kappa, P'_B * Q) \quad (11)$$

This last lemma follows by composing the view shift in LftVs with the one we get (from Q to P).

Proof sketch of $\text{LFTL-BOR-ACC-ATOMIC-CONS}$, $\text{LFTL-IDX-ACC-ATOMIC}$. We start by inspecting whether κ_0 is alive. If it is not, we obtain $[\dagger\kappa_0]$, close the invariant again, and use $\kappa \sqsubseteq \kappa_0$ to obtain $[\dagger\kappa]$. The closing view shift is trivial.

If κ' is alive, we take apart $\triangleright \text{LftBorAlive}(\kappa, P_B)$. From the borrow token we got, we have $B(i) = \text{in}$, and thus the slice with P is full and we can empty it.

We are done with $\text{LFTL-IDX-ACC-ATOMIC}$ now: for the closing view shift, we get $\triangleright P$ again, so we fill the slice and are done.

For $\text{LFTL-BOR-ACC-ATOMIC-CONS}$, we instead continue like we did above with LFTL-BOR-ACC-CONS : we remove the empty slice and create a new one with Q in it, and then we update LftVs using (11).

Proof sketch of LFTL-BOR-SEP. First, we look at the left-to-right direction (splitting), which is simpler. We start by looking at $\text{LftInv}(\kappa')$. If it is dead, then we use (4) to just “fake” $\&_{\text{full}}^{\kappa'} P$ and $\&_{\text{full}}^{\kappa'} Q$, which can both be changed to borrows at κ using LFTL-IDX-SHORTEN.

So we can assume we have $\triangleright \text{LftAlive}(\kappa')$ and, in particular, $\triangleright \text{LftBorAlive}$. We split the slice containing $P * Q$ into two slices containing P and Q , respectively. We also have to fix up the authoritative map B , which is easy because we own the fragment corresponding to the slice we removed. This gives us fragments for the new new slices, so we can finish up the proof by putting these fragments into the proofs of $\&_{\text{full}}^{\kappa'} P$ and $\&_{\text{full}}^{\kappa'} Q$.

Next, we look at the right-to-left direction (merging). The trouble here is that we obtain *two* borrows, at two potentially different lifetimes κ_0 and κ_1 . We do have $\kappa \sqsubseteq \kappa_0$ and $\kappa \sqsubseteq \kappa_1$. We use reborrowing to obtain both borrows at lifetime $\kappa' := \kappa \sqcap \kappa_0 \sqcap \kappa_1$. Notice that we have $\kappa \sqsubseteq \kappa'$. Now we can start to do the actual merging. We check whether κ' is dead; if yes, so is κ and we can “fake” the result. So we obtain $\triangleright \text{LftAlive}(\kappa')$ and, in particular, $\triangleright \text{LftBorAlive}$. We merge the two slices containing P and Q , respectively, and turn them into a slice containing $P * Q$. We have all the fragments we need to fix up the authoritative B , so we can close everything up again and obtain $\&_{\text{full}}^{\kappa'} P * \&_{\text{full}}^{\kappa'} P$. By LFTL-BOR-SHORTEN, we are done.

Proof sketch of LFTL-BEGIN. The first step of this proof is easy, it just involves allocating a new atomic lifetime Λ in ALFT and returning a singleton $\kappa := \{\Lambda\}$.

This leaves us with proving the closing view shift. Before we come to the core proof, we show some helper lemmas. We start with a lemma to end a single intersected lifetime κ :

$$(\forall \kappa'. \kappa' \in \text{dom}(I) \wedge \kappa' \subset \kappa \Rightarrow \kappa' \in K) \wedge (\forall \kappa'. \kappa' \in \text{dom}(I) \wedge \kappa \subset \kappa' \Rightarrow \kappa' \in K') \vdash$$

$$\text{LftAlive}(\kappa) * [\uparrow \kappa] \Rightarrow \left[\left[\bullet I \right]^{\gamma_1} * \bigstar_{\kappa' \in K} \text{LftAlive}(\kappa') * \bigstar_{\kappa' \in K'} \text{LftDead}(\kappa') \right]_{-\mathcal{N}_{\text{mgmt}}} \text{LftDead}(\kappa) \quad (12)$$

The proof proceeds by taking a closer look at LftBorAlive (in $\text{LftAlive}(\kappa)$). The goal is to show that all $B(i)$ are in. We can rule out **open** because we have $[\uparrow \kappa]$. For **rebor**(κ'), we obtain $\text{OwnCnt}(\kappa', \circ 1)$ for κ' strictly larger than κ . However, this implies $\kappa' \in K'$ and thus we have $\text{LftDead}(\kappa')$, which says that the authoritative count is 0 (and the ghost names match) – a contradiction. This we know $B(i) = \text{in}$. After emptying the borrow box, we obtain $\triangleright P_B$ and $\text{LftBorDead}(\kappa)$. Next, we apply **LftVs**. We have all its preconditions: $\triangleright P_B$, $[\uparrow \kappa]$, the authoritative I and **LftAlive** for all strictly shorter lifetimes. We obtain $\triangleright P_I$ and $\text{OwnCnt}(\kappa, \circ n)$, which we use with the authoritative counter to set that to 0. We fill the inheritance box, obtaining $\text{LftInh}(\kappa, P_I, \text{full})$ which completes the proof.

Next, we show how to end a whole set K of lifetimes at once. To this end, the set K must be closed under smaller lifetimes, *i.e.*, larger sets. Furthermore, we need a proof that all the other lifetimes in I that have some sublifetime in K that’s alive according to A , are still alive (these lifetimes are collected in K').

$$(\forall \kappa \in K. \forall \kappa' \in \text{dom}(I). \kappa' \supseteq \kappa \Rightarrow \kappa' \in K) \wedge$$

$$(\forall \kappa \in K. \forall \kappa' \in \text{dom}(I). \text{LftAliveIn}(A, \kappa) \wedge \kappa' \notin K \wedge \kappa' \subset \kappa \Rightarrow \kappa' \in K') \vdash$$

$$\left(\bigstar_{\kappa \in K} \text{LftInv}(A, \kappa) * [\uparrow \kappa] \right) \Rightarrow \left[\left[\bullet I \right]^{\gamma_1} * \bigstar_{\kappa' \in K'} \text{LftAlive}(\kappa') \right]_{-\mathcal{N}_{\text{mgmt}}} \bigstar_{\kappa \in K} \text{LftDead}(\kappa) \quad (13)$$

To show (13), we perform induction over the metric $|K|$, *i.e.*, over the size of the kill-set K . We start by checking whether K contains any lifetimes κ such that $\text{LftAliveIn}(A, \kappa)$. If no, we have nothing to do. Otherwise, we select a *minimal* element $\kappa \in \{\kappa \in K \mid \text{LftAliveIn}(A, \kappa)\}$ according to the relation \subset . This relation is acyclic, so such an element has to exist. We let $K'' := K \setminus \{\kappa\}$ and $K''' := K' \cup \{\kappa\}$. Clearly, K'' is smaller than K , so we can invoke the induction hypothesis to kill K'' while framing K' . To this end, we have to show that K'' is up-closed. This is the case because κ is not only minimal in $\{\kappa \in K \mid \text{LftAliveIn}(A, \kappa)\}$, it is also minimal in K . Furthermore we have to show that K''' contains all lifetimes below something alive in K'' . This is the case, because such a lifetime κ' will either also be in the down-closure of K (and hence it is in K'), or it will be κ itself, which we added to K''' . After invoking the induction hypothesis, we have that all lifetimes in K'' are dead. To complete our goal, all that's left to do is end κ . To this end, we invoke (12). By our frame and by the fact that κ is a minimal alive lifetime in K , we know that all lifetimes strictly shorter than κ are in K' and hence alive. Furthermore, we know that all lifetimes strictly larger than κ are in K'' and hence dead. This completes the proof.

Now, we can come back to proving the closing view shift of **LFTL-BEGIN**. We start out assuming $[\{\Lambda\}]_1$. We open **LftLCtx** and take a step to obtain **LftLInv**. From the token we own, we know $A(\Lambda) = \text{LftStAlive}$. We update our token to obtain $[\dagger\{\Lambda\}]$. Now we want to apply (13) with $K := \{\kappa \in I \mid \Lambda \in \kappa\}$ and $K' := \{\kappa \in \text{dom}(I) \mid \kappa \notin K \wedge \exists \kappa' \in K. \text{LftAliveIn}(A, \kappa') \wedge \kappa \subset \kappa'\}$. From $[\dagger\{\Lambda\}]$ we can get the corresponding token for all $\kappa \in K$. We thus satisfy all requirements of (13), which finishes the proof.

5 λ_{Rust} model

Well-formed terms of the type system are interpreted as Iris terms. Valuable expressions are interpreted as the value they evaluate to. Variables in the type system are interpreted as Iris variables of appropriate sort.

5.1 Types

Types are complex beasts. After the preparations we made in §4, it should not come as a surprise that borrowing of λ_{Rust} types will be explained using the lifetime logic. The notation has already been suggestively chosen to match the one used in the syntactic type system.

The domain of semantic types is defined in Figure 6, and the interpretation of all primitive types is defined in Figure 7. In the following, we will develop this definition step-by-step, together with the semantic interpretation of the most important (and most complex) types: Owned pointers, as well as mutable and shared references. We will also talk about products (restricted to the representative case of pairs), which is probably the least surprising type.

The core of a semantic type is its notion of *ownership* $\llbracket \tau \rrbracket.\text{own} : TId \times \text{list}(Val) \rightarrow iProp$. This is a *thread-indexed predicate over a list of values*. Why *lists* of values? A type describes a *continuous region of memory*: Compound types like products and sums take up more than one memory location, because they need more than one value to be represented. For example, the interpretation of $\tau_1 \times \tau_2$ will demand that the given list is the *concatenation* of two lists accepted by τ_1 and τ_2 , respectively, while the uninitialized type just accepts any list of appropriate length:

$$\begin{aligned} \llbracket \tau_1 \times \tau_2 \rrbracket.\text{own}(t, \bar{v}) &:= \exists \bar{v}_1, \bar{v}_2. \bar{v} = \bar{v}_1 \mathbin{++} \bar{v}_2 * \llbracket \tau_1 \rrbracket.\text{own}(t, \bar{v}_1) * \llbracket \tau_2 \rrbracket.\text{own}(t, \bar{v}_2) \\ \llbracket \zeta_n \rrbracket.\text{own}(_, \bar{v}) &:= |\bar{v}| = n \end{aligned}$$

The thread-relative nature of these predicates is needed to model types which are not **Send** or **Sync**, types which crucially rely on being accessed from only a single thread.

It turns out that we need another bit of data to properly describe a type: We need to know its *size*. This is given as $\llbracket \tau \rrbracket.\text{size} : \mathbb{N}$, such that **TY-SIZE** holds. Note that we also need to have

Semantic Type A *semantic type* is a tuple $(\text{size} \in \mathbb{N}, \text{own} \in TId \times \text{list}(Val) \rightarrow iProp, \text{shr} \in Lft \times TId \times Loc \rightarrow iProp)$ such that

$$\begin{aligned} \forall t, \bar{v}. \text{own}(t, \bar{v}) &\Rightarrow |\bar{v}| = \text{size} && (\text{TY-SIZE}) \\ \forall \kappa, t, \ell. \text{shr}(\kappa, t, \ell) &\Rightarrow \Box \text{shr}(\kappa, t, \ell) && (\text{TY-SHR-PERSISTENT}) \\ \forall \kappa, t, \ell. \&\kappa_{\text{full}}^\kappa (\ell \mapsto \text{own}(t)) * [\kappa]_q &\Rightarrow_{\mathcal{N}_{\text{ift}}} \text{shr}(\kappa, t, \ell) * [\kappa]_q && (\text{TY-SHARE}) \\ \forall \kappa, \kappa', t, \ell. \kappa' &\sqsubseteq \kappa \Rightarrow \text{shr}(\kappa, t, \ell) \Rightarrow \text{shr}(\kappa', t, \ell) && (\text{TY-SHR-MONO}) \end{aligned}$$

Semantic type inclusion

$$\begin{aligned} \tau_1 \sqsubseteq^{\text{ty}} \tau_2 &:= \tau_1.\text{size} = \tau_2.\text{size} \wedge (\Box \forall t, \bar{v}. \tau_1.\text{own}(t, \bar{v}) \Rightarrow \tau_2.\text{own}(t, \bar{v})) \wedge \\ &\quad (\Box \forall \kappa, t, \ell. \tau_1.\text{shr}(\kappa, t, \ell) \Rightarrow \tau_2.\text{shr}(\kappa, t, \ell)) \end{aligned}$$

Figure 6: The semantic domain of types.

$\text{size}(\tau) = \llbracket \tau \rrbracket.\text{size}$.

5.1.1 Owned pointers and mutable references

With the foundations introduced above, we can now define

$$\begin{aligned} \llbracket \mathbf{own}_n \tau \rrbracket.\text{own}(t, \bar{v}) &:= \exists \ell. \bar{v} = [\ell] * \triangleright \ell \mapsto \llbracket \tau \rrbracket.\text{own}(t) * \left(\llbracket \tau \rrbracket.\text{size} = 0 \vee \llbracket n \rrbracket > 0 * \triangleright \dagger \frac{\llbracket \tau \rrbracket.\text{size}}{\llbracket \tau \rrbracket.\text{size}/\llbracket n \rrbracket} \ell \right) \\ \llbracket \&_{\mathbf{mut}}^\kappa \tau \rrbracket.\text{own}(t, \bar{v}) &:= \exists \ell. \bar{v} = [\ell] * \&_{\mathbf{full}}^\kappa \ell \mapsto \llbracket \tau \rrbracket.\text{own}(t) \end{aligned}$$

Here and in the remainder of this document, $\ell \mapsto \Phi$ is sugar for $\exists \bar{v}. \ell \mapsto \bar{v} * \Phi(\bar{v})$. Notice the close relationship between the two types: The mutable reference has a borrow of the content where the owned pointer, well, owns it. In particular, both definitions are *contractive*. The only remaining difference is that owned pointers have the permission to deallocate what they point to, which mutable references do not have.

With this setup, it is clear how borrowing can start (**C-BORROW**). It is also clear that owning a mutable reference *at an ongoing lifetime* justifies reading and writing through that pointer (**TREAD-BOR** and **TWRITE-BOR**).

Re-borrowing an already borrowed pointer (**C-REBORROW**) is justified by **LFTL-REBORROW**.

The most interesting rules are **S-DEREF-BOR-OWN** and **S-DEREF-BOR-BOR**. In both cases, after opening the borrow, we *freeze* the current value of the pointer: Since the original borrow is lost, we know it cannot be changed anymore. For **S-DEREF-BOR-BOR**, we also have to justify getting out the inner borrow: We start with a nested borrow, and after opening the outer one and taking a step, we use **LFTL-BOR-UNNEST** to both close the outer borrow and obtain the inner one at the outer lifetime.

5.1.2 Shared references

The story is unfortunately more complicated for shared references. This is because sharing is a rather non-uniform idea in Rust: For many types, sharing them (*i.e.*, having a shared reference to an element of such a type) implies that the contents of the variable are *frozen*, *i.e.*, all mutation is prohibited. This makes it trivially safe for everybody to perform read-only actions on this variable, without and risk of data races or invalid pointers. The obvious model for this in separation logic is to provide every holder of a shared reference with some fraction of the pointer.

Other types, however, provide *interior mutability*, which means that it *is* possible to perform mutation through a shared reference. This should be modeled by having *full* ownership of the pointer available, guarded through some protocol so that every holder of a shared reference can obtain the full permission if they follow the protocol.

The way we express this formally is that *every type decides itself what it means to be shared*. The interpretation of $\&_{\mathbf{shr}}^\kappa \tau$ is not given uniformly; instead, it is defined by the type itself:

$$\llbracket \&_{\mathbf{shr}}^\kappa \tau \rrbracket.\text{own}(t, \bar{v}) := \exists \ell. \bar{v} = [\ell] * \Box \llbracket \tau \rrbracket.\text{shr}(\kappa, t, \ell)$$

The sharing predicate has to satisfy two core properties: By **TY-SHR-PERSISTENT**, sharing is persistent. By **TY-SHARE**, sharing can be started from a mutable reference (compare the premise of this rule to the interpretation of mutable references above). Furthermore, sharing predicates have to satisfy some closure properties (**TY-SHR-MONO**).

$$\begin{aligned}
\llbracket \mathbf{!} \rrbracket_\gamma &:= \text{Type} \{ \text{size} := 0; \text{own} := \lambda _ . \mathbf{False}; \text{shr} := \lambda _ . _ . \mathbf{False} \} & (\text{TY-DEF-EMP}) \\
\llbracket \mathbf{unit} \rrbracket_\gamma &:= \text{Type} \{ \text{size} := 0; \text{own} := \lambda _ . \bar{v}. \bar{v} = []; \text{shr} := \lambda _ . _ . \mathbf{True} \} & (\text{TY-DEF-UNIT}) \\
\llbracket \mathbf{bool} \rrbracket_\gamma &:= \text{SimpleType}(\lambda _ . v. v = \mathbf{true} \vee v = \mathbf{false}) & (\text{TY-DEF-BOOL}) \\
\llbracket \mathbf{int} \rrbracket_\gamma &:= \text{SimpleType}(\lambda _ . \bar{v}. \exists z. v = z) & (\text{TY-DEF-NAT}) \\
\llbracket \mathbf{own}_n \tau \rrbracket_\gamma &:= \text{Type} \{ \text{size} := 1; \\
&\quad \text{own} := \lambda t, \bar{v}. \exists \ell. \bar{v} = [\ell] * \triangleright \ell \mapsto \llbracket \tau \rrbracket_\gamma . \text{own}(t) * \\
&\quad \left(\llbracket \tau \rrbracket_\gamma . \text{size} = 0 \vee \llbracket n \rrbracket_\gamma > 0 * \triangleright \dagger_{\llbracket \tau \rrbracket_\gamma . \text{size} / \llbracket n \rrbracket_\gamma}^{\llbracket \tau \rrbracket_\gamma . \text{size}} \ell \right); \\
&\quad \text{shr} := \lambda \kappa, t, \ell. \exists \ell'. \&\mathbf{frac}^\kappa(\lambda q'. \ell \xrightarrow{q'} \ell') * \\
&\quad \square(\forall q'. [\kappa]_{q'} \approx_{\mathcal{N}_{\text{shr}}, \mathcal{N}_{\text{ift}}}^{\mathcal{N}_{\text{ift}}} \llbracket \tau \rrbracket_\gamma . \text{shr}(\kappa, t, \ell') * [\kappa]_{q'}) \} \\
& & (\text{TY-DEF-OWN}) \\
\llbracket \&\mathbf{mut}^\kappa \tau \rrbracket_\gamma &:= \text{Type} \{ \text{size} := 1; \\
&\quad \text{own} := \lambda t, \bar{v}. \exists \ell. \bar{v} = [\ell] * \&\mathbf{full}^\kappa \ell \mapsto \llbracket \tau \rrbracket_\gamma . \text{own}(t); \\
&\quad \text{shr} := \lambda \kappa', t, \ell. \exists \ell'. \&\mathbf{frac}^{\kappa'}(\lambda q. \ell \xrightarrow{q} \ell') * \\
&\quad \square(\forall q. [\llbracket \kappa \rrbracket \sqcap \kappa']_q \approx_{\mathcal{E}, \mathcal{N}_{\text{ift}}}^{\mathcal{N}_{\text{ift}}} \llbracket \tau \rrbracket_\gamma . \text{shr}(\llbracket \kappa \rrbracket \sqcap \kappa', t, \ell') * [\llbracket \kappa \rrbracket \sqcap \kappa']_q) \} \\
& & (\text{TY-DEF-MUT}) \\
\llbracket \&\mathbf{shr}^\kappa \tau \rrbracket_\gamma &:= \text{SimpleType}(\lambda t, v. \exists \ell. v = \ell * \llbracket \tau \rrbracket_\gamma . \text{shr}(\llbracket \kappa \rrbracket, t, \ell)) & (\text{TY-DEF-SHR}) \\
\llbracket \Pi \bar{\tau} \rrbracket_\gamma &:= \text{Type} \{ \text{size} := \sum_i \llbracket \bar{\tau}_i \rrbracket_\gamma . \text{size}; \\
&\quad \text{own} := \lambda t, \bar{v}. \exists \bar{v}. \bar{v} = \sum_i \bar{v}_i * \bigstar_i \llbracket \bar{\tau}_i \rrbracket_\gamma . \text{own}(t, \bar{v}_i); \\
&\quad \text{shr} := \lambda \kappa, t, \ell. \bigstar_i \llbracket \bar{\tau}_i \rrbracket_\gamma . \text{shr}(\kappa, t, \ell + \sum_{j < i} \llbracket \bar{\tau}_j \rrbracket_\gamma . \text{size}) \} \\
& & (\text{TY-DEF-PROD}) \\
\llbracket \Sigma \bar{\tau} \rrbracket_\gamma &:= \text{Type} \{ \text{size} := 1 + \max_i \llbracket \bar{\tau}_i \rrbracket_\gamma . \text{size}; \\
&\quad \text{own} := \lambda t, \bar{v}. \exists i, \bar{v}', \bar{v}''. \bar{v} = [i] \uparrow \bar{v}' \uparrow \bar{v}'' * \llbracket \bar{\tau}_i \rrbracket_\gamma . \text{own}(t, \bar{v}') * \\
&\quad |\bar{v}''| = 1 + \max_j \llbracket \bar{\tau}_j \rrbracket_\gamma . \text{size}; \\
&\quad \text{shr} := \lambda \kappa, t, \ell. \exists i. \llbracket \bar{\tau}_i \rrbracket_\gamma . \text{shr}(\kappa, t, \ell + 1) * \&\mathbf{frac}^\kappa(\lambda q. \ell \xrightarrow{q} i * \\
&\quad (\ell + 1 + \llbracket \bar{\tau}_i \rrbracket_\gamma . \text{size} \mapsto \lambda \bar{v}. \llbracket \bar{\tau}_i \rrbracket_\gamma . \text{size} + |\bar{v}| = \max_j \llbracket \bar{\tau}_j \rrbracket_\gamma . \text{size})) \} \\
& & (\text{TY-DEF-SUM}) \\
\llbracket \downarrow_1 \rrbracket_\gamma &:= \text{SimpleType}(\lambda _ . \mathbf{True}) & (\text{TY-DEF-UNINIT1}) \\
\llbracket \downarrow_n \rrbracket_\gamma &:= \Pi[\downarrow_1, \dots, \downarrow_1] \quad \text{of length } \llbracket n \rrbracket & (\text{TY-DEF-UNINIT}) \\
\llbracket \forall \bar{\alpha}. \mathbf{fn}_{(\mathcal{F}} : \mathbf{E}; \bar{\tau}) \rightarrow \tau \rrbracket_\gamma &:= \text{SimpleType}(\lambda _ . v. \exists f, \bar{x}, k, F. v = \mathbf{funrec} f(\bar{x}) \mathbf{ret} k := F * \\
&\quad \triangleright \forall \bar{\kappa}, \kappa_{\mathcal{F}}, v_k, \bar{v}. \square[\mathbf{E};_{\mathcal{F}} \sqsubseteq_1 [] \mid k \triangleleft \mathbf{cont}_{(\mathcal{F}} \sqsubseteq_1 []; x. x \triangleleft \mathbf{own} \tau); \\
&\quad \bar{x} \triangleleft \mathbf{own} \bar{\tau} \vdash F[\mathbf{funrec} f(\bar{x}) \mathbf{ret} k := F/f, \bar{v}/\bar{x}, v_k/k] \rrbracket_{\gamma[\bar{\alpha} \leftarrow \bar{\kappa}][_{\mathcal{F}} \leftarrow \kappa_{\mathcal{F}}][\bar{x} \leftarrow \bar{v}][k \leftarrow v_k]}) \\
& & (\text{TY-DEF-FN}) \\
\llbracket \mu T. \tau \rrbracket &:= \text{fix}(\lambda T. \llbracket \tau \rrbracket_{\gamma[T \leftarrow T]})
\end{aligned}$$

Figure 7: The interpretations of all the primitive types

Simple sharing. The most simple form of sharing occurs when sharing “plain data”. For example, consider the type **int**. We define ownership and sharing of a number as follows:

$$\begin{aligned} \llbracket \mathbf{int} \rrbracket.\text{own}(t, \bar{v}) &:= \exists z. \bar{v} = [z] \\ \llbracket \mathbf{int} \rrbracket.\text{shr}(\kappa, -, \ell) &:= \exists \bar{v}. \&_{\mathbf{frac}}^{\kappa}(\lambda q. \ell \xrightarrow{q} \bar{v}) * \llbracket \mathbf{int} \rrbracket.\text{own}(t, \bar{v}) \end{aligned}$$

Actually, this kind of sharing is so common that we define a notion of *simple types* using the above sharing predicate:

$$\begin{aligned} \text{SimpleType}(\Phi) &:= \text{Type} \{ \text{size} := 1; \\ &\quad \text{own} := \lambda t, \bar{v}. \exists v. \bar{v} = [v] * \Phi(t, v); \\ &\quad \text{shr} := \lambda \kappa, t, \ell. \exists v. \&_{\mathbf{frac}}^{\kappa}(\lambda q. \ell \xrightarrow{q} v) * \triangleright \Phi(t, v) \} \end{aligned}$$

For a simple type, we only have to show that the ownership predicate Φ is persistent and that **TY-SIZE** holds. The remaining properties follow.

Sharing owned pointers and mutable references. It turns out that owned pointers and mutable references have interesting and similar sharing predicates. Here’s the full definition:

$$\begin{aligned} \llbracket \mathbf{own}_n \tau \rrbracket.\text{shr}(\kappa, t, \ell) &:= \exists \ell'. \&_{\mathbf{frac}}^{\kappa}(\lambda q. \ell \xrightarrow{q} \ell') * \square(\forall q. [\kappa]_q \approx_{\mathcal{N}_{\text{shr}}, \mathcal{N}_{\text{ift}}}^{\mathcal{N}_{\text{ift}}} \llbracket \tau \rrbracket.\text{shr}(\kappa, t, \ell') * [\kappa]_q) \\ \llbracket \&_{\mathbf{mut}}^{\kappa} \tau \rrbracket.\text{shr}(\kappa', t, \ell) &:= \exists \ell'. \&_{\mathbf{frac}}^{\kappa'}(\lambda q. \ell \xrightarrow{q} \ell') * \square(\forall q. [\kappa \sqcap \kappa']_q \approx_{\mathcal{E}, \mathcal{N}_{\text{ift}}}^{\mathcal{N}_{\text{ift}}} \llbracket \tau \rrbracket.\text{shr}(\kappa \sqcap \kappa', t, \ell') * [\kappa \sqcap \kappa']_q) \end{aligned}$$

using again the “magic update that takes a step” introduced in §4.

As you can see, a shared reference to an owned pointer or mutable reference consists of two parts: First of all, the outer pointer is shared. This is a straight-forward fractured borrow such that everybody gets a read-only permission to dereference the outer pointer. This is just like the simple sharing predicate defined above.

The more complicated, second part involves sharing the inner pointer. What we would like to have here (instead of the view shift that takes a step) is just $\llbracket \tau \rrbracket.\text{shr}(\kappa, t, \ell')$. Unfortunately, that would make it impossible to prove **TY-SHARE**: Since owned pointers are, well, pointers, they only own the data they point to *later* or under a borrow, respectively. This also means they cannot start sharing that data immediately, since execution view shifts under \triangleright is not possible. Instead, we remember that we can start sharing the data behind the inner pointer *once someone lets us take a step*. In other words, only when the outer pointer is dereferenced, the sharing of the inner data actually starts. We kind of lazily apply sharing down the pointer chain. To prove this view shift, there will be a small invariant managing this, since the shared references could actually be distributed accross different threads – there is a race for who gets to actually start the sharing.

5.1.3 Compound types: Sums and products

Both ownership and sharing of sums and products are fairly straight-forward, but there is some book-keeping and list manipulation going on that blurs the view.

$$\begin{aligned}
\llbracket \Pi \bar{\tau} \rrbracket.\text{own}(t, \bar{v}) &:= \exists \bar{v}. \bar{v} = \sum_i \bar{v}_i * \bigstar_i \llbracket \bar{\tau}_i \rrbracket.\text{own}(t, \bar{v}_i) \\
\llbracket \Pi \bar{\tau} \rrbracket.\text{shr}(\kappa, t, \ell) &:= \bigstar_i \llbracket \bar{\tau}_i \rrbracket.\text{shr}(\kappa, t, \ell + \sum_{j < i} \llbracket \bar{\tau}_j \rrbracket.\text{size}) \\
\llbracket \Sigma \bar{\tau} \rrbracket.\text{own}(t, \bar{v}) &:= \exists i, \bar{v}', \bar{v}''. \bar{v} = [i] \uparrow \bar{v}' \uparrow \bar{v}'' * |\bar{v}''| = 1 + \max_j \llbracket \bar{\tau}_j \rrbracket.\text{size} * \llbracket \bar{\tau}_i \rrbracket.\text{own}(t, \bar{v}') \\
\llbracket \Sigma \bar{\tau} \rrbracket.\text{shr}(\kappa, t, \ell) &:= \exists i. \&_{\text{frac}}^\kappa (\lambda q. \ell \xrightarrow{q} i * (\ell + 1 + \llbracket \bar{\tau}_i \rrbracket.\text{size} \mapsto \lambda \bar{v}. \llbracket \bar{\tau}_i \rrbracket.\text{size} + |\bar{v}| = \max_j \llbracket \bar{\tau}_j \rrbracket.\text{size})) * \\
&\quad \llbracket \bar{\tau}_i \rrbracket.\text{shr}(\kappa, t, \ell + 1)
\end{aligned}$$

5.1.4 Copy, Send, Sync

$$\begin{aligned}
\llbracket \tau \text{ copy} \rrbracket_\gamma &:= (\Box \forall t, \bar{v}. \llbracket \tau \rrbracket_\gamma.\text{own}(t, \bar{v}) \Rightarrow \Box \llbracket \tau \rrbracket_\gamma.\text{own}(t, \bar{v})) \wedge \\
&\quad \Box \forall \kappa, t, \ell, q. \llbracket \tau \rrbracket_\gamma.\text{shr}(\kappa, t, \ell) \multimap \langle [\text{Na} : t.\mathcal{N}_{\text{shr}}. [\geq \ell, < \ell + 1 + \llbracket \tau \rrbracket_\gamma.\text{size}]] * [\kappa]_q \Leftrightarrow \\
&\quad q'. [\text{Na} : t.\mathcal{N}_{\text{shr}}. (\ell + \llbracket \tau \rrbracket_\gamma.\text{size})] * \triangleright \ell \xrightarrow{q'} \llbracket \tau \rrbracket_\gamma.\text{own}(t) \rangle_{\mathcal{N}_{\text{shr}}, \mathcal{N}_{\text{ift}}} \\
\llbracket \tau \text{ send} \rrbracket_\gamma &:= \Box \forall t_1, t_2, \bar{v}. \llbracket \tau \rrbracket_\gamma.\text{own}(t_1, \bar{v}) \Rightarrow \llbracket \tau \rrbracket_\gamma.\text{own}(t_2, \bar{v}) \\
\llbracket \tau \text{ sync} \rrbracket_\gamma &:= \Box \forall \kappa, t_1, t_2, \ell. \llbracket \tau \rrbracket_\gamma.\text{shr}(\kappa, t_1, \ell) \Rightarrow \llbracket \tau \rrbracket_\gamma.\text{shr}(\kappa, t_2, \ell)
\end{aligned}$$

5.2 Type and continuation contexts

$$\begin{aligned}
\llbracket \mathbf{T} \rrbracket_\gamma &: TId \rightarrow iProp \\
\llbracket \emptyset \rrbracket_\gamma(t) &:= \text{True} \\
\llbracket \mathbf{T}, p \triangleleft \tau \rrbracket_\gamma(t) &:= \llbracket \tau \rrbracket_\gamma.\text{own}(t, [\llbracket p \rrbracket]) * \llbracket \mathbf{T} \rrbracket_\gamma(t) \\
\llbracket \mathbf{T}, p \triangleleft^{\dagger \kappa} \tau \rrbracket_\gamma(t) &:= ([\dagger \llbracket \kappa \rrbracket] \Rightarrow \bigstar_\top \llbracket \tau \rrbracket_\gamma.\text{own}(t, [\llbracket p \rrbracket])) * \llbracket \mathbf{T} \rrbracket_\gamma(t) \\
\llbracket \mathbf{K} \rrbracket_\gamma &: TId \rightarrow iProp \\
\llbracket \emptyset \rrbracket_\gamma(t) &:= \text{True} \\
\llbracket \mathbf{K}, k \triangleleft \text{cont}(\mathbf{L}; \bar{x}. \mathbf{T}) \rrbracket_\gamma(t) &:= (\forall \bar{v}. [\text{Na} : t] * \llbracket \mathbf{L} \rrbracket_\gamma(1) * \llbracket \mathbf{T} \rrbracket_{\gamma[\bar{x} \leftarrow \bar{v}]}(t) \multimap \text{wp } \llbracket k \rrbracket_\gamma(\bar{v}) \{ \text{True} \}) \wedge \llbracket \mathbf{K} \rrbracket_\gamma(t)
\end{aligned}$$

5.3 Lifetime contexts and judgments

The local and external lifetime contexts are interpreted as follows:

$$\begin{aligned}
\llbracket \mathbf{E} \rrbracket_\gamma &: iProp \\
\llbracket \emptyset \rrbracket_\gamma &:= \text{True} \\
\llbracket \mathbf{E}, \kappa \sqsubseteq_e \kappa' \rrbracket_\gamma &:= \llbracket \kappa \rrbracket_\gamma \sqsubseteq \llbracket \kappa' \rrbracket_\gamma * \llbracket \mathbf{E} \rrbracket_\gamma \\
\\
\llbracket \mathbf{L} \rrbracket_\gamma &: Fract \rightarrow iProp \\
\llbracket \emptyset \rrbracket_\gamma(q) &:= \text{True} \\
\llbracket \mathbf{L}, \kappa \sqsubseteq_1 \bar{\kappa} \rrbracket_\gamma(q) &:= \exists \kappa'. \llbracket \kappa \rrbracket_\gamma = \kappa' \sqcap (\sqcap \llbracket \bar{\kappa} \rrbracket_\gamma) * [\kappa']_q * \square([\kappa']_1 \multimap_{\emptyset}^{\mathcal{N}_{\text{ift}}} [\dagger \kappa']) * \llbracket \mathbf{L} \rrbracket_\gamma(q) \\
\\
\llbracket \mathbf{E}_1; \mathbf{L}_1 \vdash \mathbf{E}_2 \rrbracket &:= \forall q. \llbracket \mathbf{L} \rrbracket_\gamma(q) \multimap \square(\llbracket \mathbf{E}_1 \rrbracket \multimap \llbracket \mathbf{E}_2 \rrbracket) \\
\llbracket \mathbf{E}; \mathbf{L} \vdash \kappa_1 \sqsubseteq \kappa_2 \rrbracket &:= \forall q. \llbracket \mathbf{L} \rrbracket_\gamma(q) \multimap \square(\llbracket \mathbf{E} \rrbracket \multimap \llbracket \kappa_1 \rrbracket \sqsubseteq \llbracket \kappa_2 \rrbracket) \\
\llbracket \mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive} \rrbracket &:= \forall q. \llbracket \mathbf{E} \rrbracket \multimap \langle \llbracket \mathbf{L} \rrbracket_\gamma(q) \Leftrightarrow q'. \llbracket \kappa \rrbracket_{q'} \rangle
\end{aligned}$$

5.4 Judgments

$$\begin{aligned}
\llbracket \mathbf{E}; \mathbf{L} \vdash \tau_1 \Rightarrow \tau_2 \rrbracket_\gamma &:= \forall q. \llbracket \mathbf{L} \rrbracket_\gamma(q) \multimap \square(\llbracket \mathbf{E} \rrbracket_\gamma \multimap \tau_1 \sqsubseteq^{\text{by}} \tau_2) \\
\llbracket \mathbf{E}; \mathbf{L} \vdash \mathbf{T}_1 \Rightarrow \mathbf{T}_2 \rrbracket_\gamma &:= \forall t, q. \llbracket \mathbf{E} \rrbracket_\gamma * \llbracket \mathbf{L} \rrbracket_\gamma(q) * \llbracket \mathbf{T}_1 \rrbracket_\gamma(t) \multimap \llbracket \mathbf{T}_2 \rrbracket_\gamma(t) \\
\llbracket \mathbf{T}_1 \Rightarrow^{\dagger \kappa} \mathbf{T}_2 \rrbracket_\gamma &:= \forall t. [\dagger \kappa] * \llbracket \mathbf{T}_1 \rrbracket_\gamma(t) \multimap \llbracket \mathbf{T}_2 \rrbracket_\gamma(t) \\
\llbracket \mathbf{E} \vdash \mathbf{K}_1 \Rightarrow \mathbf{K}_2 \rrbracket_\gamma &:= \forall t. \llbracket \mathbf{E} \rrbracket_\gamma * \llbracket \mathbf{K}_1 \rrbracket_\gamma(t) \multimap \llbracket \mathbf{K}_2 \rrbracket_\gamma(t) \\
\llbracket \mathbf{E}; \mathbf{L} \vdash \tau_1 \multimap^\tau \tau_2 \rrbracket_\gamma &:= \forall \ell, t, q. \llbracket \mathbf{E} \rrbracket_\gamma * \llbracket \mathbf{L} \rrbracket_\gamma(q) * \llbracket \tau_1 \rrbracket_\gamma.\text{own}(t, [v]) \multimap_{\mathcal{N}_{\text{ift}}, \mathcal{N}_{\text{rust}}} \exists \bar{\ell}, \bar{v}. \ell = v * |\bar{v}| = \llbracket \tau \rrbracket_\gamma.\text{size} * \ell \mapsto \bar{v} * \\
&\quad \left(\triangleright \ell \mapsto \llbracket \tau \rrbracket_\gamma.\text{own}(t) \multimap_{\mathcal{N}_{\text{ift}}, \mathcal{N}_{\text{rust}}} \llbracket \mathbf{L} \rrbracket_\gamma(q) * \llbracket \tau_2 \rrbracket_\gamma.\text{own}(t, v) \right) \\
\llbracket \mathbf{E}; \mathbf{L} \vdash \tau_1 \multimap^\tau \tau_2 \rrbracket_\gamma &:= \forall \ell, t, q. \llbracket \mathbf{E} \rrbracket_\gamma * \llbracket \mathbf{L} \rrbracket_\gamma(q) * [\text{Na} : t] * \llbracket \tau_1 \rrbracket_\gamma.\text{own}(t, [v]) \multimap_{\mathcal{N}_{\text{ift}}, \mathcal{N}_{\text{rust}}} \\
&\quad \exists \bar{\ell}, \bar{v}, q'. \ell = v * \bar{\ell} \xrightarrow{q'} \bar{v} * \triangleright \llbracket \tau \rrbracket_\gamma.\text{own}(t, \bar{v}) * \\
&\quad \left(\bar{\ell} \xrightarrow{q'} \bar{v} \multimap_{\mathcal{N}_{\text{ift}}, \mathcal{N}_{\text{rust}}} \llbracket \mathbf{L} \rrbracket_\gamma(q) * [\text{Na} : t] * \llbracket \tau_2 \rrbracket_\gamma.\text{own}(t, v) \right) \\
\llbracket \mathbf{E}; \mathbf{L} \mid \mathbf{T}_1 \vdash I \dashv x. \mathbf{T}_2 \rrbracket_\gamma &:= \forall t. \llbracket \mathbf{E} \rrbracket_\gamma * \llbracket \mathbf{L} \rrbracket_\gamma(1) * [\text{Na} : t] * \llbracket \mathbf{T}_1 \rrbracket_\gamma(t) \multimap \text{wp } I \{ v. \llbracket \mathbf{L} \rrbracket_\gamma(1) * [\text{Na} : t] * \llbracket \mathbf{T}_2 \rrbracket_{\gamma[x \leftarrow v]}(t) \} \\
\llbracket \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \vdash F \rrbracket_\gamma &:= \forall t. \llbracket \mathbf{E} \rrbracket_\gamma * \llbracket \mathbf{L} \rrbracket_\gamma(1) * [\text{Na} : t] * \llbracket \mathbf{K} \rrbracket_\gamma(t) * \llbracket \mathbf{T} \rrbracket_\gamma(t) \multimap \text{wp } F \{ \text{True} \}
\end{aligned}$$

5.5 Theorems

Theorem 1 (Compatibility of the logical relation). *For any inference rule of the type system, if we wrap all judgments in semantic brackets $\llbracket - \rrbracket$, the resulting Iris theorem holds.*

Theorem 2 (Adequacy of the logical relation). *Let f be a λ_{Rust} function such that the Iris assertion $\llbracket \emptyset; \emptyset \mid \emptyset \vdash f \dashv x. x \triangleleft \mathbf{fn}() \rightarrow \Pi[] \rrbracket$ holds, then when we execute $f(\lambda x. x)$ (passing it the default continuation), no execution ends in a stuck state.*