

The Iris 3.0 Documentation

<http://plv.mpi-sws.org/iris/>

November 23, 2017

Abstract

This document describes formally the Iris program logic. Every result in this document has been fully verified in Coq. The latest versions of this document and the Coq formalization can be found in the git repository at <https://gitlab.mpi-sws.org/FP/iris-coq/>. For further information, visit the Iris project website at <http://plv.mpi-sws.org/iris/>.

Contents

1 Algebraic Structures	4
1.1 OFE	4
1.2 COFE	5
1.3 RA	5
1.4 CMRA	6
2 OFE and COFE constructions	9
2.1 Trivial pointwise lifting	9
2.2 Next (type-level later)	9
2.3 Uniform Predicates	9
3 RA and CMRA constructions	10
3.1 Product	10
3.2 Sum	10
3.3 Option	10
3.4 Finite partial function	11
3.5 Agreement	11
3.6 Exclusive CMRA	11
3.7 Authoritative	12
3.8 STS with tokens	12
4 Base Logic	15
4.1 Grammar	15
4.2 Types	15
4.3 Proof rules	16
4.4 Consistency	18
5 Model and semantics	19
6 Extensions of the Base Logic	22
6.1 Derived rules about base connectives	22
6.2 Persistent assertions	22
6.3 Timeless assertions and except-0	22
7 Language	24
7.1 Concurrent language	24
8 Program Logic	25
8.1 Dynamic Composeable Higher-Order Resources	25
8.2 World Satisfaction, Invariants, Fancy Updates	27
8.3 Weakest Precondition	28
8.4 Invariant Namespaces	31
8.5 Accessors	32

9	Derived constructions	33
9.1	Non-atomic (“thread-local”) invariants	33
9.2	Boxes	34
10	Logical paradoxes	36
10.1	Saved propositions without a later	36
10.2	Invariants without a later	37

1 Algebraic Structures

1.1 OFE

The model of Iris lives in the category of *Ordered Families of Equivalences* (OFEs). This definition varies slightly from the original one in [2].

Definition 1. An ordered family of equivalences (OFE) is a tuple $(T, (\overset{n}{=} \subseteq T \times T)_{n \in \mathbb{N}})$ satisfying

$$\forall n. (\overset{n}{=}) \text{ is an equivalence relation} \quad (\text{OFE-EQUIV})$$

$$\forall n, m. n \geq m \Rightarrow (\overset{n}{=}) \subseteq (\overset{m}{=}) \quad (\text{OFE-MONO})$$

$$\forall x, y. x = y \Leftrightarrow (\forall n. x \overset{n}{=} y) \quad (\text{OFE-LIMIT})$$

The key intuition behind OFEs is that elements x and y are n -equivalent, notation $x \overset{n}{=} y$, if they are *equivalent for n steps of computation*, i.e., if they cannot be distinguished by a program running for no more than n steps. In other words, as n increases, $\overset{n}{=}$ becomes more and more refined (OFE-MONO)—and in the limit, it agrees with plain equality (OFE-LIMIT).

Definition 2. An element $x \in T$ of an OFE is called *discrete* if

$$\forall y \in T. x \overset{0}{=} y \Rightarrow x = y$$

An OFE A is called *discrete* if all its elements are discrete. For a set X , we write ΔX for the discrete OFE with $x \overset{n}{=} x' \triangleq x = x'$

Definition 3. A function $f : T \rightarrow U$ between two OFEs is *non-expansive* (written $f : T \xrightarrow{ne} U$) if

$$\forall n, x \in T, y \in T. x \overset{n}{=} y \Rightarrow f(x) \overset{n}{=} f(y)$$

It is *contractive* if

$$\forall n, x \in T, y \in T. (\forall m < n. x \overset{m}{=} y) \Rightarrow f(x) \overset{n}{=} f(y)$$

Intuitively, applying a non-expansive function to some data will not suddenly introduce differences between seemingly equal data. Elements that cannot be distinguished by programs within n steps remain indistinguishable after applying f .

Definition 4. The category \mathcal{OFE} consists of OFEs as objects, and non-expansive functions as arrows.

Note that \mathcal{OFE} is cartesian closed. In particular:

Definition 5. Given two OFEs T and U , the set of non-expansive functions $\{f : T \xrightarrow{ne} U\}$ is itself an OFE with

$$f \overset{n}{=} g \triangleq \forall x \in T. f(x) \overset{n}{=} g(x)$$

Definition 6. A (bi)functor $F : \mathcal{OFE} \rightarrow \mathcal{OFE}$ is called *locally non-expansive* if its action F_1 on arrows is itself a non-expansive map. Similarly, F is called *locally contractive* if F_1 is a contractive map.

The function space $(-) \xrightarrow{ne} (-)$ is a locally non-expansive bifunctor. Note that the composition of non-expansive (bi)functors is non-expansive, and the composition of a non-expansive and a contractive (bi)functor is contractive.

1.2 COFE

COFEs are *complete OFEs*, which means that we can take limits of arbitrary chains.

Definition 7 (Chain). *Given some set T and an indexed family $(\stackrel{n}{\subseteq} \subseteq T \times T)_{n \in \mathbb{N}}$ of equivalence relations, a chain is a function $c : \mathbb{N} \rightarrow T$ such that $\forall n, m. n \leq m \Rightarrow c(m) \stackrel{n}{\subseteq} c(n)$.*

Definition 8. A complete ordered family of equivalences (COFE) is a tuple $(T : \mathcal{OFE}, \text{lim} : \text{chain}(T) \rightarrow T)$ satisfying

$$\forall n, c. \text{lim}(c) \stackrel{n}{\subseteq} c(n) \quad (\text{COFE-COMPL})$$

Definition 9. The category \mathcal{COFE} consists of COFEs as objects, and non-expansive functions as arrows.

The function space $T \xrightarrow{\text{ne}} U$ is a COFE if U is a COFE (i.e., the domain T can actually be just an OFE).

Completeness is necessary to take fixed-points. For once, every *contractive function* $f : T \rightarrow U$ where U is a COFE and inhabited has a *unique* fixed-point $\text{fix}(f)$ such that $\text{fix}(f) = f(\text{fix}(f))$. This also holds if f^k is contractive for an arbitrary k . Furthermore, by America and Rutten's theorem [1, 3], every contractive (bi)functor from \mathcal{COFE} to \mathcal{COFE} has a unique¹ fixed-point.

1.3 RA

Definition 10. A resource algebra (RA) is a tuple $(M, \mathcal{V} \subseteq M, |-| : M \rightarrow M^?, (\cdot) : M \times M \rightarrow M)$ satisfying:

$$\forall a, b, c. (a \cdot b) \cdot c = a \cdot (b \cdot c) \quad (\text{RA-ASSOC})$$

$$\forall a, b. a \cdot b = b \cdot a \quad (\text{RA-COMM})$$

$$\forall a. |a| \in M \Rightarrow |a| \cdot a = a \quad (\text{RA-CORE-ID})$$

$$\forall a. |a| \in M \Rightarrow ||a|| = |a| \quad (\text{RA-CORE-IDEM})$$

$$\forall a, b. |a| \in M \wedge a \preceq b \Rightarrow |b| \in M \wedge |a| \preceq |b| \quad (\text{RA-CORE-MONO})$$

$$\forall a, b. (a \cdot b) \in \mathcal{V} \Rightarrow a \in \mathcal{V} \quad (\text{RA-VALID-OP})$$

$$\text{where } M^? \triangleq M \uplus \{\perp\} \quad a^? \cdot \perp \triangleq \perp \cdot a^? \triangleq a^?$$

$$a \preceq b \triangleq \exists c \in M. b = a \cdot c \quad (\text{RA-INCL})$$

RAs are closely related to *Partial Commutative Monoids* (PCMs), with two key differences:

1. The composition operation on RAs is total (as opposed to the partial composition operation of a PCM), but there is a specific subset \mathcal{V} of *valid* elements that is compatible with the composition operation (RA-VALID-OP).

This take on partiality is necessary when defining the structure of *higher-order* ghost state, CMRAs, in the next subsection.

¹Uniqueness is not proven in Coq.

2. Instead of a single unit that is an identity to every element, we allow for an arbitrary number of units, via a function $|-|$ assigning to an element a its (*duplicable*) *core* $|a|$, as demanded by **RA-CORE-ID**. We further demand that $|-|$ is idempotent (**RA-CORE-IDEM**) and monotone (**RA-CORE-MONO**) with respect to the *extension order*, defined similarly to that for PCMs (**RA-INCL**).

Notice that the domain of the core is $M^?$, a set that adds a dummy element \perp to M . Thus, the core can be *partial*: not all elements need to have a unit. We use the metavariable $a^?$ to indicate elements of $M^?$. We also lift the composition (\cdot) to $M^?$. Partial cores help us to build interesting composite RAs from smaller primitives.

Notice also that the core of an RA is a strict generalization of the unit that any PCM must provide, since $|-|$ can always be picked as a constant function.

Definition 11. *It is possible to do a frame-preserving update from $a \in M$ to $B \subseteq M$, written $a \rightsquigarrow B$, if*

$$\forall a_f^? \in M^?. a \cdot a_f^? \in \mathcal{V} \Rightarrow \exists b \in B. b \cdot a_f^? \in \mathcal{V}$$

We further define $a \rightsquigarrow b \triangleq a \rightsquigarrow \{b\}$.

The assertion $a \rightsquigarrow B$ says that every element $a_f^?$ compatible with a (we also call such elements *frames*), must also be compatible with some $b \in B$. Notice that $a_f^?$ could be \perp , so the frame-preserving update can also be applied to elements that have *no* frame. Intuitively, this means that whatever assumptions the rest of the program is making about the state of γ , if these assumptions are compatible with a , then updating to b will not invalidate any of these assumptions. Since Iris ensures that the global ghost state is valid, this means that we can soundly update the ghost state from a to a non-deterministically picked $b \in B$.

1.4 CMRA

Definition 12. *A CMRA is a tuple $(M : \mathcal{OFE}, (\mathcal{V}_n \subseteq M)_{n \in \mathbb{N}}, |-| : M \xrightarrow{ne} M^?, (\cdot) : M \times M \xrightarrow{ne} M)$ satisfying:*

$$\forall n, a, b. a \stackrel{n}{=} b \wedge a \in \mathcal{V}_n \Rightarrow b \in \mathcal{V}_n \quad (\text{CMRA-VALID-NE})$$

$$\forall n, m. n \geq m \Rightarrow \mathcal{V}_n \subseteq \mathcal{V}_m \quad (\text{CMRA-VALID-MONO})$$

$$\forall a, b, c. (a \cdot b) \cdot c = a \cdot (b \cdot c) \quad (\text{CMRA-ASSOC})$$

$$\forall a, b. a \cdot b = b \cdot a \quad (\text{CMRA-COMM})$$

$$\forall a. |a| \in M \Rightarrow |a| \cdot a = a \quad (\text{CMRA-CORE-ID})$$

$$\forall a. |a| \in M \Rightarrow ||a|| = |a| \quad (\text{CMRA-CORE-IDEM})$$

$$\forall a, b. |a| \in M \wedge a \preceq b \Rightarrow |b| \in M \wedge |a| \preceq |b| \quad (\text{CMRA-CORE-MONO})$$

$$\forall n, a, b. (a \cdot b) \in \mathcal{V}_n \Rightarrow a \in \mathcal{V}_n \quad (\text{CMRA-VALID-OP})$$

$$\forall n, a, b_1, b_2. a \in \mathcal{V}_n \wedge a \stackrel{n}{=} b_1 \cdot b_2 \Rightarrow \exists c_1, c_2. a = c_1 \cdot c_2 \wedge c_1 \stackrel{n}{=} b_1 \wedge c_2 \stackrel{n}{=} b_2 \quad (\text{CMRA-EXTEND})$$

where

$$a \preceq b \triangleq \exists c. b = a \cdot c \quad (\text{CMRA-INCL})$$

$$a \preceq^n b \triangleq \exists c. b \stackrel{n}{=} a \cdot c \quad (\text{CMRA-INCLN})$$

This is a natural generalization of RAs over OFEs. All operations have to be non-expansive, and the validity predicate \mathcal{V} can now also depend on the step-index. We define the plain \mathcal{V} as the “limit” of the \mathcal{V}_n :

$$\mathcal{V} \triangleq \bigcap_{n \in \mathbb{N}} \mathcal{V}_n$$

The extension axiom (CMRA-EXTEND). Notice that the existential quantification in this axiom is *constructive*, *i.e.*, it is a sigma type in Coq. The purpose of this axiom is to compute a_1 , a_2 completing the following square:

$$\begin{array}{ccc} a & \stackrel{n}{=} & b \\ \parallel & & \parallel \\ a_1 \cdot a_2 & \stackrel{n}{=} & b_1 \cdot b_2 \end{array}$$

where the n -equivalence at the bottom is meant to apply to the pairs of elements, *i.e.*, we demand $a_1 \stackrel{n}{=} b_1$ and $a_2 \stackrel{n}{=} b_2$. In other words, extension carries the decomposition of b into b_1 and b_2 over the n -equivalence of a and b , and yields a corresponding decomposition of a into a_1 and a_2 . This operation is needed to prove that \triangleright commutes with separating conjunction:

$$\triangleright(P * Q) \Leftrightarrow \triangleright P * \triangleright Q$$

Definition 13. An element ε of a CMRA M is called the unit of M if it satisfies the following conditions:

1. ε is valid:
 $\forall n. \varepsilon \in \mathcal{V}_n$
2. ε is a left-identity of the operation:
 $\forall a \in M. \varepsilon \cdot a = a$
3. ε is its own core:
 $|\varepsilon| = \varepsilon$

Lemma 1. If M has a unit ε , then the core $|-|$ is total, *i.e.*, $\forall a. |a| \in M$.

Definition 14. It is possible to do a frame-preserving update from $a \in M$ to $B \subseteq M$, written $a \rightsquigarrow B$, if

$$\forall n, a_f^? . a \cdot a_f^? \in \mathcal{V}_n \Rightarrow \exists b \in B. b \cdot a_f^? \in \mathcal{V}_n$$

We further define $a \rightsquigarrow b \triangleq a \rightsquigarrow \{b\}$.

Note that for RAs, this and the RA-based definition of a frame-preserving update coincide.

Definition 15. A CMRA M is discrete if it satisfies the following conditions:

1. M is a discrete COFE
2. \mathcal{V} ignores the step-index:
 $\forall a \in M. a \in \mathcal{V}_0 \Rightarrow \forall n, a \in \mathcal{V}_n$

Note that every RA is a discrete CMRA, by picking the discrete COFE for the equivalence relation. Furthermore, discrete CMRAs can be turned into RAs by ignoring their COFE structure, as well as the step-index of \mathcal{V} .

Definition 16. A function $f : M_1 \rightarrow M_2$ between two CMRAs is monotone (written $f : M_1 \xrightarrow{\text{mon}} M_2$) if it satisfies the following conditions:

1. f is non-expansive
2. f preserves validity:
 $\forall n, a \in M_1. a \in \mathcal{V}_n \Rightarrow f(a) \in \mathcal{V}_n$
3. f preserves CMRA inclusion:
 $\forall a \in M_1, b \in M_1. a \preceq b \Rightarrow f(a) \preceq f(b)$

Definition 17. The category \mathcal{CMRA} consists of CMRAs as objects, and monotone functions as arrows.

Note that every object/arrow in \mathcal{CMRA} is also an object/arrow of \mathcal{CFE} . The notion of a locally non-expansive (or contractive) bifunctor naturally generalizes to bifunctors between these categories.

2 OFE and COFE constructions

2.1 Trivial pointwise lifting

The (C)OFE structure on many types can be easily obtained by pointwise lifting of the structure of the components. This is what we do for option $T^?$, product $(M_i)_{i \in I}$ (with I some finite index set), sum $T + T'$ and finite partial functions $K \xrightarrow{\text{fin}} M$ (with K infinite countable).

2.2 Next (type-level later)

Given a OFE T , we define $\blacktriangleright T$ as follows (using a datatype-like notation to define the type):

$$\begin{aligned} \blacktriangleright T &\triangleq \text{next}(x : T) \\ \text{next}(x) \stackrel{n}{=} \text{next}(y) &\triangleq n = 0 \vee x \stackrel{n-1}{=} y \end{aligned}$$

Note that in the definition of the carrier $\blacktriangleright T$, next is a constructor (like the constructors in Coq), *i.e.*, this is short for $\{\text{next}(x) \mid x \in T\}$.

$\blacktriangleright(-)$ is a locally *contractive* functor from \mathcal{OFE} to \mathcal{OFE} .

2.3 Uniform Predicates

Given a CMRA M , we define the COFE $UPred(M)$ of *uniform predicates* over M as follows:

$$UPred(M) \triangleq \left\{ \Phi : \mathbb{N} \times M \rightarrow Prop \mid \begin{array}{l} (\forall n, x, y. \Phi(n, x) \wedge x \stackrel{n}{=} y \Rightarrow \Phi(n, y)) \wedge \\ (\forall n, m, x, y. \Phi(n, x) \wedge x \preceq y \wedge m \leq n \wedge y \in \mathcal{V}_m \Rightarrow \Phi(m, y)) \end{array} \right\}$$

where $Prop$ is the set of meta-level propositions, *e.g.*, Coq 's Prop . $UPred(-)$ is a locally non-expansive functor from \mathcal{CMRA} to \mathcal{COFE} .

One way to understand this definition is to re-write it a little. We start by defining the COFE of *step-indexed propositions*: For every step-index, the proposition either holds or does not hold.

$$\begin{aligned} SProp &\triangleq \wp^\downarrow(\mathbb{N}) \\ &\triangleq \{X \in \wp(\mathbb{N}) \mid \forall n, m. n \geq m \Rightarrow n \in X \Rightarrow m \in X\} \\ X \stackrel{n}{=} Y &\triangleq \forall m \leq n. m \in X \Leftrightarrow m \in Y \end{aligned}$$

Notice that this notion of $SProp$ is already hidden in the validity predicate \mathcal{V}_n of a CMRA: We could equivalently require every CMRA to define $\mathcal{V}_-(-) : M \xrightarrow{\text{ne}} SProp$, replacing **CMRA-VALID-NE** and **CMRA-VALID-MONO**.

Now we can rewrite $UPred(M)$ as monotone step-indexed predicates over M , where the definition of a “monotone” function here is a little funny.

$$\begin{aligned} UPred(M) &\cong M \xrightarrow{\text{mon}} SProp \\ &\triangleq \left\{ \Phi : M \xrightarrow{\text{ne}} SProp \mid \forall n, m, x, y. n \in \Phi(x) \wedge x \preceq y \wedge m \leq n \wedge y \in \mathcal{V}_m \Rightarrow m \in \Phi(y) \right\} \end{aligned}$$

The reason we chose the first definition is that it is easier to work with in Coq .

3 RA and CMRA constructions

3.1 Product

Given a family $(M_i)_{i \in I}$ of CMRAs (I finite), we construct a CMRA for the product $\prod_{i \in I} M_i$ by lifting everything pointwise.

Frame-preserving updates on the M_i lift to the product:

$$\frac{\text{PROD-UPDATE} \quad a \rightsquigarrow_{M_i} B}{f [i \leftarrow a] \rightsquigarrow \{f [i \leftarrow b] \mid b \in B\}}$$

3.2 Sum

The *sum CMRA* $M_1 +_{\downarrow} M_2$ for any CMRAs M_1 and M_2 is defined as (again, we use a datatype-like notation):

$$\begin{aligned} M_1 +_{\downarrow} M_2 &\triangleq \text{inl}(a_1 : M_1) \mid \text{inr}(a_2 : M_2) \mid \downarrow \\ \mathcal{V}_n &\triangleq \{\text{inl}(a_1) \mid a_1 \in \mathcal{V}'_n\} \cup \{\text{inr}(a_2) \mid a_2 \in \mathcal{V}''_n\} \\ \text{inl}(a_1) \cdot \text{inl}(b_1) &\triangleq \text{inl}(a_1 \cdot b_1) \\ |\text{inl}(a_1)| &\triangleq \begin{cases} \perp & \text{if } |a_1| = \perp \\ \text{inl}(|a_1|) & \text{otherwise} \end{cases} \end{aligned}$$

The composition and core for inr are defined symmetrically. The remaining cases of the composition and core are all \downarrow . Above, \mathcal{V}' refers to the validity of M_1 , and \mathcal{V}'' to the validity of M_2 .

Notice that we added the artificial “invalid” (or “undefined”) element \downarrow to this CMRA just in order to make certain compositions of elements (in this case, inl and inr) invalid.

The step-indexed equivalence is inductively defined as follows:

$$\frac{x \stackrel{n}{=} y}{\text{inl}(x) \stackrel{n}{=} \text{inl}(y)} \quad \frac{x \stackrel{n}{=} y}{\text{inr}(x) \stackrel{n}{=} \text{inr}(y)} \quad \downarrow \stackrel{n}{=} \downarrow$$

We obtain the following frame-preserving updates, as well as their symmetric counterparts:

$$\frac{\text{SUM-UPDATE} \quad a \rightsquigarrow_{M_1} B}{\text{inl}(a) \rightsquigarrow \{\text{inl}(b) \mid b \in B\}} \quad \frac{\text{SUM-SWAP} \quad \forall a_{\dagger}, n. a \cdot a_{\dagger} \notin \mathcal{V}'_n \quad b \in \mathcal{V}''}{\text{inl}(a) \rightsquigarrow \text{inr}(b)}$$

Crucially, the second rule allows us to *swap* the “side” of the sum that the CMRA is on if \mathcal{V} has *no possible frame*.

3.3 Option

The definition of the (CM)RA axioms already lifted the composition operation on M to one on $M^?$. We can easily extend this to a full CMRA by defining a suitable core, namely

$$\begin{aligned} |\perp| &\triangleq \perp \\ |a^?| &\triangleq |a| \quad \text{If } a^? \neq \perp \end{aligned}$$

Notice that this core is total, as the result always lies in $M^?$ (rather than in $M^{?^?}$).

3.4 Finite partial function

Given some infinite countable K and some CMRA M , the set of finite partial functions $K \xrightarrow{\text{fin}} M$ is equipped with a CMRA structure by lifting everything pointwise.

We obtain the following frame-preserving updates:

$$\begin{array}{ccc} \text{FPFN-ALLOC-STRONG} & \text{FPFN-ALLOC} & \text{FPFN-UPDATE} \\ \frac{G \text{ infinite} \quad a \in \mathcal{V}}{\emptyset \rightsquigarrow \{[\gamma \leftarrow a] \mid \gamma \in G\}} & \frac{a \in \mathcal{V}}{\emptyset \rightsquigarrow \{[\gamma \leftarrow a] \mid \gamma \in K\}} & \frac{a \rightsquigarrow_M B}{f[i \leftarrow a] \rightsquigarrow \{f[i \leftarrow b] \mid b \in B\}} \end{array}$$

Above, \mathcal{V} refers to the validity of M .

$K \xrightarrow{\text{fin}} (-)$ is a locally non-expansive functor from \mathcal{CMRA} to \mathcal{CMRA} .

3.5 Agreement

Given some OFE T , we define the CMRA $\text{AG}(T)$ as follows:

$$\begin{aligned} \text{AG}(T) &\triangleq \{a \in \wp^{\text{fin}}(T) \mid a \neq \emptyset\} / \sim \\ a \stackrel{n}{=} b &\triangleq (\forall x \in a. \exists y \in b. x \stackrel{n}{=} y) \wedge (\forall y \in b. \exists x \in a. x \stackrel{n}{=} y) \\ \text{where } a \sim b &\triangleq \forall n. a \stackrel{n}{=} b \end{aligned}$$

$$\begin{aligned} \mathcal{V}_n &\triangleq \{a \in \text{AG}(T) \mid \forall x, y \in a. x \stackrel{n}{=} y\} \\ |a| &\triangleq a \\ a \cdot b &\triangleq a \cup b \end{aligned}$$

$\text{AG}(-)$ is a locally non-expansive functor from \mathcal{OFE} to \mathcal{CMRA} .

We define a non-expansive injection ag into $\text{AG}(T)$ as follows:

$$\text{ag}(x) \triangleq \{x\}$$

There are no interesting frame-preserving updates for $\text{AG}(T)$, but we can show the following:

$$\begin{array}{ccc} \text{AG-VAL} & \text{AG-DUP} & \text{AG-AGREE} \\ \text{ag}(x) \in \mathcal{V}_n & \text{ag}(x) = \text{ag}(x) \cdot \text{ag}(x) & \text{ag}(x) \cdot \text{ag}(y) \in \mathcal{V}_n \Leftrightarrow x \stackrel{n}{=} y \end{array}$$

3.6 Exclusive CMRA

Given an OFE T , we define a CMRA $\text{EX}(T)$ such that at most one $x \in T$ can be owned:

$$\begin{aligned} \text{EX}(T) &\triangleq \text{ex}(T) \mid \downarrow \\ \mathcal{V}_n &\triangleq \{a \in \text{EX}(T) \mid a \neq \downarrow\} \end{aligned}$$

All cases of composition go to \downarrow .

$$|\text{ex}(x)| \triangleq \perp \qquad |\downarrow| \triangleq \downarrow$$

Remember that \perp is the “dummy” element in $M^?$ indicating (in this case) that $\text{ex}(x)$ has no core.

The step-indexed equivalence is inductively defined as follows:

$$\frac{x \stackrel{n}{=} y}{\text{ex}(x) \stackrel{n}{=} \text{ex}(y)} \quad \downarrow \stackrel{n}{=} \downarrow$$

$\text{EX}(-)$ is a locally non-expansive functor from $\mathcal{OF}\mathcal{E}$ to \mathcal{CMRA} .

We obtain the following frame-preserving update:

$$\begin{array}{c} \text{EX-UPDATE} \\ \text{ex}(x) \rightsquigarrow \text{ex}(y) \end{array}$$

3.7 Authoritative

Given a CMRA M , we construct $\text{AUTH}(M)$ modeling someone owning an *authoritative* element a of M , and others potentially owning fragments $b \preceq a$ of a . We assume that M has a unit ε , and hence its core is total. (If M is an exclusive monoid, the construction is very similar to a half-ownership monoid with two asymmetric halves.)

$$\begin{aligned} \text{AUTH}(M) &\triangleq \text{EX}(M)^\uparrow \times M \\ \mathcal{V}_n &\triangleq \{(x, b) \in \text{AUTH}(M) \mid b \in \mathcal{V}_n \wedge (x = \perp \vee \exists a. x = \text{ex}(a) \wedge b \preceq_n a)\} \\ (x_1, b_1) \cdot (x_2, b_2) &\triangleq (x_1 \cdot x_2, b_2 \cdot b_1) \\ |(x, b)| &\triangleq (\perp, |b|) \\ (x_1, b_1) \stackrel{n}{=} (x_2, b_2) &\triangleq x_1 \stackrel{n}{=} x_2 \wedge b_1 \stackrel{n}{=} b_2 \end{aligned}$$

Note that (\perp, ε) is the unit and asserts no ownership whatsoever, but $(\text{ex}(\varepsilon), \varepsilon)$ asserts that the authoritative element is ε .

Let $a, b \in M$. We write $\bullet a$ for full ownership $(\text{ex}(a), \varepsilon)$ and $\circ b$ for fragmental ownership (\perp, b) and $\bullet a, \circ b$ for combined ownership $(\text{ex}(a), b)$.

The frame-preserving update involves the notion of a *local update*:

Definition 18. *It is possible to do a local update from a_1 and b_1 to a_2 and b_2 , written $(a_1, b_1) \rightsquigarrow^1 (a_2, b_2)$, if*

$$\forall n, a_f^\uparrow. a_1 \in \mathcal{V}_n \wedge a_1 \stackrel{n}{=} b_1 \cdot a_f^\uparrow \Rightarrow a_2 \in \mathcal{V}_n \wedge a_2 \stackrel{n}{=} b_2 \cdot a_f^\uparrow$$

In other words, the idea is that for every possible frame a_f^\uparrow completing b_1 to a_1 , the same frame also completes b_2 to a_2 .

We then obtain

$$\begin{array}{c} \text{AUTH-UPDATE} \\ \frac{(a_1, b_1) \rightsquigarrow^1 (a_2, b_2)}{\bullet a_1, \circ b_1 \rightsquigarrow \bullet a_2, \circ b_2} \end{array}$$

3.8 STS with tokens

Given a state-transition system (STS, *i.e.*, a directed graph) $(\mathcal{S}, \rightarrow \subseteq \mathcal{S} \times \mathcal{S})$, a set of tokens \mathcal{T} , and a labeling $\mathcal{L} : \mathcal{S} \rightarrow \wp(\mathcal{T})$ of *protocol-owned* tokens for each state, we construct an RA modeling

an authoritative current state and permitting transitions given a *bound* on the current state and a set of *locally-owned* tokens.

The construction follows the idea of STSs as described in CaReSL [6]. We first lift the transition relation to $\mathcal{S} \times \wp(\mathcal{T})$ (implementing a *law of token conservation*) and define a stepping relation for the *frame* of a given token set:

$$\begin{aligned} (s, T) \rightarrow (s', T') &\triangleq s \rightarrow s' \wedge \mathcal{L}(s) \uplus T = \mathcal{L}(s') \uplus T' \\ s \xrightarrow{T} s' &\triangleq \exists T_1, T_2. T_1 \# \mathcal{L}(s) \cup T \wedge (s, T_1) \rightarrow (s', T_2) \end{aligned}$$

We further define *closed* sets of states (given a particular set of tokens) as well as the *closure* of a set:

$$\begin{aligned} \text{closed}(S, T) &\triangleq \forall s \in S. \mathcal{L}(s) \# T \wedge (\forall s'. s \xrightarrow{T} s' \Rightarrow s' \in S) \\ \uparrow(S, T) &\triangleq \{s' \in \mathcal{S} \mid \exists s \in S. s \xrightarrow{T}^* s'\} \end{aligned}$$

The STS RA is defined as follows

$$\begin{aligned} M &\triangleq \text{auth}(s : \mathcal{S}, T : \wp(\mathcal{T}) \mid \mathcal{L}(s) \# T) \mid \\ &\quad \text{frag}(S : \wp(\mathcal{S}), T : \wp(\mathcal{T}) \mid \text{closed}(S, T) \wedge S \neq \emptyset) \mid \not\downarrow \\ \mathcal{V} &\triangleq \{a \in M \mid a \neq \not\downarrow\} \\ \text{frag}(S_1, T_1) \cdot \text{frag}(S_2, T_2) &\triangleq \text{frag}(S_1 \cap S_2, T_1 \cup T_2) \quad \text{if } T_1 \# T_2 \text{ and } S_1 \cap S_2 \neq \emptyset \\ \text{frag}(S, T) \cdot \text{auth}(s, T') &\triangleq \text{auth}(s, T') \cdot \text{frag}(S, T) \triangleq \text{auth}(s, T \cup T') \quad \text{if } T \# T' \text{ and } s \in S \\ |\text{frag}(S, T)| &\triangleq \text{frag}(\uparrow(S, \emptyset), \emptyset) \\ |\text{auth}(s, T)| &\triangleq \text{frag}(\uparrow(\{s\}, \emptyset), \emptyset) \end{aligned}$$

The remaining cases are all $\not\downarrow$.

We will need the following frame-preserving update:

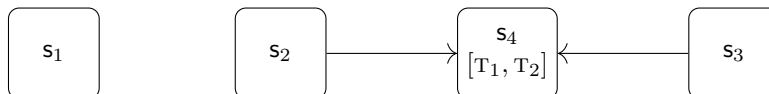
$$\frac{\text{STS-STEP} \quad (s, T) \rightarrow^* (s', T')}{\text{auth}(s, T) \rightsquigarrow \text{auth}(s', T')} \quad \frac{\text{STS-WEAKEN} \quad \text{closed}(S_2, T_2) \quad S_1 \subseteq S_2 \quad T_2 \subseteq T_1}{\text{frag}(S_1, T_1) \rightsquigarrow \text{frag}(S_2, T_2)}$$

The core is not a homomorphism. The core of the STS construction is only satisfying the RA axioms because we are *not* demanding the core to be a homomorphism—all we demand is for the core to be monotone with respect the **RA-INCL**.

In other words, the following does *not* hold for the STS core as defined above:

$$|a| \cdot |b| = |a \cdot b|$$

To see why, consider the following STS:



Now consider the following two elements of the STS RA:

$$a \triangleq \text{frag}(\{s_1, s_2\}, \{T_1\}) \quad b \triangleq \text{frag}(\{s_1, s_3\}, \{T_2\})$$

We have:

$$\begin{aligned} a \cdot b &= \text{frag}(\{s_1\}, \{T_1, T_2\}) & |a| &= \text{frag}(\{s_1, s_2, s_4\}, \emptyset) & |b| &= \text{frag}(\{s_1, s_3, s_4\}, \emptyset) \\ |a| \cdot |b| &= \text{frag}(\{s_1, s_4\}, \emptyset) \neq |a \cdot b| = \text{frag}(\{s_1\}, \emptyset) \end{aligned}$$

4 Base Logic

The base logic is parameterized by an arbitrary CMRA M having a unit ε . By [Lemma 1](#), this means that the core of M is a total function, so we will treat it as such in the following. This defines the structure of resources that can be owned.

As usual for higher-order logics, you can furthermore pick a *signature* $\mathcal{S} = (\mathcal{T}, \mathcal{F}, \mathcal{A})$ to add more types, symbols and axioms to the language. You have to make sure that \mathcal{T} includes the base types:

$$\mathcal{T} \supseteq \{M, \text{iProp}\}$$

Elements of \mathcal{T} are ranged over by T .

Each function symbol in \mathcal{F} has an associated *arity* comprising a natural number n and an ordered list of $n + 1$ types τ (the grammar of τ is defined below, and depends only on \mathcal{T}). We write

$$F : \tau_1, \dots, \tau_n \rightarrow \tau_{n+1} \in \mathcal{F}$$

to express that F is a function symbol with the indicated arity.

Furthermore, \mathcal{A} is a set of *axioms*, that is, terms t of type iProp . Again, the grammar of terms and their typing rules are defined below, and depends only on \mathcal{T} and \mathcal{F} , not on \mathcal{A} . Elements of \mathcal{A} are ranged over by A .

4.1 Grammar

Syntax. Iris syntax is built up from a signature \mathcal{S} and a countably infinite set Var of variables (ranged over by metavariables x, y, z). Below, a ranges over M and i ranges over $\{1, 2\}$.

$$\begin{aligned} \tau ::= & T \mid 1 \mid \tau \times \tau \mid \tau \rightarrow \tau \\ t, P, \Phi ::= & x \mid F(t_1, \dots, t_n) \mid () \mid (t, t) \mid \pi_i t \mid \lambda x : \tau. t \mid t(t) \mid a \mid |t| \mid t \cdot t \mid \\ & \text{False} \mid \text{True} \mid t =_\tau t \mid P \Rightarrow P \mid P \wedge P \mid P \vee P \mid P * P \mid P \multimap P \mid \\ & \mu x : \tau. t \mid \exists x : \tau. P \mid \forall x : \tau. P \mid \text{Own}(t) \mid \mathcal{V}(t) \mid \Box P \mid \triangleright P \mid \boxplus P \end{aligned}$$

Recursive predicates must be *guarded*: in $\mu x. t$, the variable x can only appear under the later \triangleright modality.

Note that the modalities \boxplus , \Box and \triangleright bind more tightly than $*$, \multimap , \wedge , \vee , and \Rightarrow .

Variable conventions. We assume that, if a term occurs multiple times in a rule, its free variables are exactly those binders which are available at every occurrence.

4.2 Types

Iris terms are simply-typed. The judgment $\Gamma \vdash t : \tau$ expresses that, in variable context Γ , the term t has type τ .

A variable context, $\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n$, declares a list of variables and their types. In writing $\Gamma, x : \tau$, we presuppose that x is not already declared in Γ .

Well-typed terms

 $\boxed{\Gamma \vdash_{\mathcal{S}} t : \tau}$

$$\begin{array}{c}
 x : \tau \vdash x : \tau \quad \frac{\Gamma \vdash t : \tau}{\Gamma, x : \tau' \vdash t : \tau} \quad \frac{\Gamma, x : \tau', y : \tau' \vdash t : \tau}{\Gamma, x : \tau' \vdash t[x/y] : \tau} \quad \frac{\Gamma_1, x : \tau', y : \tau'', \Gamma_2 \vdash t : \tau}{\Gamma_1, x : \tau'', y : \tau', \Gamma_2 \vdash t[y/x, x/y] : \tau} \\
 \\
 \frac{\Gamma \vdash t_1 : \tau_1 \quad \cdots \quad \Gamma \vdash t_n : \tau_n \quad F : \tau_1, \dots, \tau_n \rightarrow \tau_{n+1} \in \mathcal{F}}{\Gamma \vdash F(t_1, \dots, t_n) : \tau_{n+1}} \quad \Gamma \vdash () : 1 \\
 \\
 \frac{\Gamma \vdash t : \tau_1 \quad \Gamma \vdash u : \tau_2}{\Gamma \vdash (t, u) : \tau_1 \times \tau_2} \quad \frac{\Gamma \vdash t : \tau_1 \times \tau_2 \quad i \in \{1, 2\}}{\Gamma \vdash \pi_i t : \tau_i} \quad \frac{\Gamma, x : \tau \vdash t : \tau'}{\Gamma \vdash \lambda x. t : \tau \rightarrow \tau'} \\
 \\
 \frac{\Gamma \vdash t : \tau \rightarrow \tau' \quad u : \tau}{\Gamma \vdash t(u) : \tau'} \quad \Gamma \vdash \varepsilon : \mathbf{M} \quad \frac{\Gamma \vdash a : \mathbf{M}}{\Gamma \vdash |a| : \mathbf{M}} \quad \frac{\Gamma \vdash a : \mathbf{M} \quad \Gamma \vdash b : \mathbf{M}}{\Gamma \vdash a \cdot b : \mathbf{M}} \\
 \\
 \Gamma \vdash \text{False} : \text{iProp} \quad \Gamma \vdash \text{True} : \text{iProp} \quad \frac{\Gamma \vdash t : \tau \quad \Gamma \vdash u : \tau}{\Gamma \vdash t =_{\tau} u : \text{iProp}} \quad \frac{\Gamma \vdash P : \text{iProp} \quad \Gamma \vdash Q : \text{iProp}}{\Gamma \vdash P \Rightarrow Q : \text{iProp}} \\
 \\
 \frac{\Gamma \vdash P : \text{iProp} \quad \Gamma \vdash Q : \text{iProp}}{\Gamma \vdash P \wedge Q : \text{iProp}} \quad \frac{\Gamma \vdash P : \text{iProp} \quad \Gamma \vdash Q : \text{iProp}}{\Gamma \vdash P \vee Q : \text{iProp}} \\
 \\
 \frac{\Gamma \vdash P : \text{iProp} \quad \Gamma \vdash Q : \text{iProp}}{\Gamma \vdash P * Q : \text{iProp}} \quad \frac{\Gamma \vdash P : \text{iProp} \quad \Gamma \vdash Q : \text{iProp}}{\Gamma \vdash P \multimap Q : \text{iProp}} \\
 \\
 \frac{\Gamma, x : \tau \vdash t : \tau \quad x \text{ is guarded in } t}{\Gamma \vdash \mu x : \tau. t : \tau} \quad \frac{\Gamma, x : \tau \vdash P : \text{iProp}}{\Gamma \vdash \exists x : \tau. P : \text{iProp}} \quad \frac{\Gamma, x : \tau \vdash P : \text{iProp}}{\Gamma \vdash \forall x : \tau. P : \text{iProp}} \\
 \\
 \frac{\Gamma \vdash a : \mathbf{M}}{\Gamma \vdash \text{Own}(a) : \text{iProp}} \quad \frac{\Gamma \vdash a : \tau \quad \tau \text{ is a CMRA}}{\Gamma \vdash \mathcal{V}(a) : \text{iProp}} \quad \frac{\Gamma \vdash P : \text{iProp}}{\Gamma \vdash \square P : \text{iProp}} \quad \frac{\Gamma \vdash P : \text{iProp}}{\Gamma \vdash \triangleright P : \text{iProp}} \\
 \\
 \frac{\Gamma \vdash P : \text{iProp}}{\Gamma \vdash \boxplus P : \text{iProp}}
 \end{array}$$

4.3 Proof rules

The judgment $\Gamma \mid P \vdash Q$ says that with free variables Γ , proposition Q holds whenever assumption P holds. Most of the rules will entirely omit the variable contexts Γ . In this case, we assume the same arbitrary context is used for every constituent of the rules. Axioms $\Gamma \mid P \dashv\vdash Q$ indicate that both $\Gamma \mid P \vdash Q$ and $\Gamma \mid Q \vdash P$ are proof rules of the logic.

 $\boxed{\Gamma \mid P \vdash Q}$

Laws of intuitionistic higher-order logic with equality. This is entirely standard.

$$\begin{array}{c}
\text{ASM} \\
\frac{}{P \vdash P} \\
\\
\text{CUT} \\
\frac{P \vdash Q \quad Q \vdash R}{P \vdash R} \\
\\
\text{EQ} \\
\frac{\Gamma, x : \tau \vdash Q : \text{iProp} \quad \Gamma \mid P \vdash Q[t/x] \quad \Gamma \mid P \vdash t =_{\tau} t'}{\Gamma \mid P \vdash Q[t'/x]} \\
\\
\text{REFL} \\
\text{True} \vdash t =_{\tau} t \\
\\
\perp\text{E} \\
\text{False} \vdash P \\
\\
\top\text{I} \\
P \vdash \text{True} \\
\\
\wedge\text{I} \\
\frac{P \vdash Q \quad P \vdash R}{P \vdash Q \wedge R} \\
\\
\wedge\text{EL} \\
\frac{P \vdash Q \wedge R}{P \vdash Q} \\
\\
\wedge\text{ER} \\
\frac{P \vdash Q \wedge R}{P \vdash R} \\
\\
\vee\text{IL} \\
\frac{P \vdash Q}{P \vdash Q \vee R} \\
\\
\vee\text{IR} \\
\frac{P \vdash R}{P \vdash Q \vee R} \\
\\
\vee\text{E} \\
\frac{P \vdash R \quad Q \vdash R}{P \vee Q \vdash R} \\
\\
\Rightarrow\text{I} \\
\frac{P \wedge Q \vdash R}{P \vdash Q \Rightarrow R} \\
\\
\Rightarrow\text{E} \\
\frac{P \vdash Q \Rightarrow R \quad P \vdash Q}{P \vdash R} \\
\\
\forall\text{I} \\
\frac{\Gamma, x : \tau \mid P \vdash Q}{\Gamma \mid P \vdash \forall x : \tau. Q} \\
\\
\forall\text{E} \\
\frac{\Gamma \mid P \vdash \forall x : \tau. Q \quad \Gamma \vdash t : \tau}{\Gamma \mid P \vdash Q[t/x]} \\
\\
\exists\text{I} \\
\frac{\Gamma \mid P \vdash Q[t/x] \quad \Gamma \vdash t : \tau}{\Gamma \mid P \vdash \exists x : \tau. Q} \\
\\
\exists\text{E} \\
\frac{\Gamma, x : \tau \mid P \vdash Q}{\Gamma \mid \exists x : \tau. P \vdash Q}
\end{array}$$

Furthermore, we have the usual η and β laws for projections, λ and μ .

Laws of (affine) bunched implications.

$$\begin{array}{c}
\text{True} * P \dashv\vdash P \\
P * Q \vdash Q * P \\
(P * Q) * R \vdash P * (Q * R) \\
\\
\text{*MONO} \\
\frac{P_1 \vdash Q_1 \quad P_2 \vdash Q_2}{P_1 * P_2 \vdash Q_1 * Q_2} \\
\\
\text{*I-E} \\
\frac{P * Q \vdash R}{P \vdash Q * R}
\end{array}$$

Laws for the always modality.

$$\begin{array}{c}
\text{\(\square\)-MONO} \\
\frac{P \vdash Q}{\square P \vdash \square Q} \\
\\
\text{\(\square\)-E} \\
\square P \vdash P \\
\\
\text{True} \vdash \square \text{True} \\
\square (P \wedge Q) \vdash \square (P * Q) \\
\square P \wedge Q \vdash \square P * Q \\
\\
\square P \vdash \square \square P \\
\forall x. \square P \vdash \square \forall x. P \\
\square \exists x. P \vdash \exists x. \square P
\end{array}$$

Laws for the later modality.

$$\begin{array}{c}
\text{\(\triangleright\)-MONO} \\
\frac{P \vdash Q}{\triangleright P \vdash \triangleright Q} \\
\\
\text{LÖB} \\
(\triangleright P \Rightarrow P) \vdash P \\
\\
\forall x. \triangleright P \vdash \triangleright \forall x. P \\
\triangleright \exists x. P \vdash \triangleright \text{False} \vee \exists x. \triangleright P \\
\triangleright P \vdash \triangleright \text{False} \vee (\triangleright \text{False} \Rightarrow P) \\
\\
\triangleright (P * Q) \dashv\vdash \triangleright P * \triangleright Q \\
\square \triangleright P \dashv\vdash \triangleright \square P
\end{array}$$

Laws for resources and validity.

$$\begin{array}{l}
\text{Own}(a) * \text{Own}(b) \dashv\vdash \text{Own}(a \cdot b) \\
\text{Own}(a) \vdash \Box \text{Own}(|a|) \\
\text{True} \vdash \text{Own}(\varepsilon) \\
\triangleright \text{Own}(a) \vdash \exists b. \text{Own}(b) \wedge \triangleright(a = b)
\end{array}
\qquad
\begin{array}{l}
\text{Own}(a) \vdash \mathcal{V}(a) \\
\mathcal{V}(a \cdot b) \vdash \mathcal{V}(a) \\
\mathcal{V}(a) \vdash \Box \mathcal{V}(a)
\end{array}$$

Laws for the basic update modality.

$$\begin{array}{c}
\text{UPD-MONO} \\
\frac{P \vdash Q}{\Rightarrow P \vdash \Rightarrow Q}
\end{array}
\qquad
\begin{array}{c}
\text{UPD-INTRO} \\
P \vdash \Rightarrow P
\end{array}
\qquad
\begin{array}{c}
\text{UPD-TRANS} \\
\Rightarrow \Rightarrow P \vdash \Rightarrow P
\end{array}
\qquad
\begin{array}{c}
\text{UPD-FRAME} \\
Q * \Rightarrow P \vdash \Rightarrow(Q * P)
\end{array}$$

$$\begin{array}{c}
\text{UPD-UPDATE} \\
\frac{a \rightsquigarrow B}{\text{Own}(a) \vdash \Rightarrow \exists b \in B. \text{Own}(b)}
\end{array}$$

The premise in **UPD-UPDATE** is a *meta-level* side-condition that has to be proven about a and B .

4.4 Consistency

The consistency statement of the logic reads as follows: For any n , we have

$$\neg(\text{True} \vdash (\Rightarrow \triangleright)^n \text{False})$$

where $(\Rightarrow \triangleright)^n$ is short for $\Rightarrow \triangleright$ being nested n times.

The reason we want a stronger consistency than the usual $\neg(\text{True} \vdash \text{False})$ is our modalities: it should be impossible to derive a contradiction below the modalities. For \Box , this follows from the elimination rule, but the other two modalities do not have an elimination rule. Hence we declare that it is impossible to derive a contradiction below any combination of these two modalities.

5 Model and semantics

The semantics closely follows the ideas laid out in [2].

Semantic domains. The semantic domains are interpreted as follows:

$$\begin{array}{ll} \llbracket \mathbf{iProp} \rrbracket \triangleq UPred(M) & \llbracket \mathbf{1} \rrbracket \triangleq \Delta\{\} \\ \llbracket \mathbf{M} \rrbracket \triangleq M & \llbracket \tau \times \tau' \rrbracket \triangleq \llbracket \tau \rrbracket \times \llbracket \tau' \rrbracket \\ & \llbracket \tau \rightarrow \tau' \rrbracket \triangleq \llbracket \tau \rrbracket \xrightarrow{ne} \llbracket \tau' \rrbracket \end{array}$$

For the remaining base types τ defined by the signature \mathcal{S} , we pick an object X_τ in \mathcal{OFE} and define

$$\llbracket \tau \rrbracket \triangleq X_\tau$$

For each function symbol $F : \tau_1, \dots, \tau_n \rightarrow \tau_{n+1} \in \mathcal{F}$, we pick a function $\llbracket F \rrbracket : \llbracket \tau_1 \rrbracket \times \dots \times \llbracket \tau_n \rrbracket \xrightarrow{ne} \llbracket \tau_{n+1} \rrbracket$.

Interpretation of assertions.

$$\boxed{\llbracket \Gamma \vdash t : \mathbf{iProp} \rrbracket : \llbracket \Gamma \rrbracket \xrightarrow{ne} UPred(M)}$$

Remember that $UPred(M)$ is isomorphic to $M \xrightarrow{\text{mon}} SProp$. We are thus going to define the assertions as mapping CMRA elements to sets of step-indices.

$$\begin{array}{l} \llbracket \Gamma \vdash t =_\tau u : \mathbf{iProp} \rrbracket_\gamma \triangleq \lambda_. \left\{ n \mid \llbracket \Gamma \vdash t : \tau \rrbracket_\gamma \stackrel{n}{\approx} \llbracket \Gamma \vdash u : \tau \rrbracket_\gamma \right\} \\ \llbracket \Gamma \vdash \mathbf{False} : \mathbf{iProp} \rrbracket_\gamma \triangleq \lambda_. \emptyset \\ \llbracket \Gamma \vdash \mathbf{True} : \mathbf{iProp} \rrbracket_\gamma \triangleq \lambda_. \mathbb{N} \\ \llbracket \Gamma \vdash P \wedge Q : \mathbf{iProp} \rrbracket_\gamma \triangleq \lambda a. \llbracket \Gamma \vdash P : \mathbf{iProp} \rrbracket_\gamma(a) \cap \llbracket \Gamma \vdash Q : \mathbf{iProp} \rrbracket_\gamma(a) \\ \llbracket \Gamma \vdash P \vee Q : \mathbf{iProp} \rrbracket_\gamma \triangleq \lambda a. \llbracket \Gamma \vdash P : \mathbf{iProp} \rrbracket_\gamma(a) \cup \llbracket \Gamma \vdash Q : \mathbf{iProp} \rrbracket_\gamma(a) \\ \llbracket \Gamma \vdash P \Rightarrow Q : \mathbf{iProp} \rrbracket_\gamma \triangleq \lambda a. \left\{ n \mid \begin{array}{l} \forall m, b. m \leq n \wedge a \preceq b \wedge b \in \mathcal{V}_m \Rightarrow \\ m \in \llbracket \Gamma \vdash P : \mathbf{iProp} \rrbracket_\gamma(b) \Rightarrow \\ m \in \llbracket \Gamma \vdash Q : \mathbf{iProp} \rrbracket_\gamma(b) \end{array} \right\} \\ \llbracket \Gamma \vdash \forall x : \tau. P : \mathbf{iProp} \rrbracket_\gamma \triangleq \lambda a. \left\{ n \mid \forall v \in \llbracket \tau \rrbracket. n \in \llbracket \Gamma, x : \tau \vdash P : \mathbf{iProp} \rrbracket_{\gamma[x \leftarrow v]}(a) \right\} \\ \llbracket \Gamma \vdash \exists x : \tau. P : \mathbf{iProp} \rrbracket_\gamma \triangleq \lambda a. \left\{ n \mid \exists v \in \llbracket \tau \rrbracket. n \in \llbracket \Gamma, x : \tau \vdash P : \mathbf{iProp} \rrbracket_{\gamma[x \leftarrow v]}(a) \right\} \end{array}$$

$$\begin{aligned}
\llbracket \Gamma \vdash P * Q : \mathbf{iProp} \rrbracket_\gamma &\triangleq \lambda a. \left\{ n \mid \begin{array}{l} \exists b_1, b_2. a \stackrel{n}{=} b_1 \cdot b_2 \wedge \\ n \in \llbracket \Gamma \vdash P : \mathbf{iProp} \rrbracket_\gamma(b_1) \wedge n \in \llbracket \Gamma \vdash Q : \mathbf{iProp} \rrbracket_\gamma(b_2) \end{array} \right\} \\
\llbracket \Gamma \vdash P \multimap Q : \mathbf{iProp} \rrbracket_\gamma &\triangleq \lambda a. \left\{ n \mid \begin{array}{l} \forall m, b. m \leq n \wedge a \cdot b \in \mathcal{V}_m \Rightarrow \\ m \in \llbracket \Gamma \vdash P : \mathbf{iProp} \rrbracket_\gamma(b) \Rightarrow \\ m \in \llbracket \Gamma \vdash Q : \mathbf{iProp} \rrbracket_\gamma(a \cdot b) \end{array} \right\} \\
\llbracket \Gamma \vdash \Box P : \mathbf{iProp} \rrbracket_\gamma &\triangleq \lambda a. \llbracket \Gamma \vdash P : \mathbf{iProp} \rrbracket_\gamma(|a|) \\
\llbracket \Gamma \vdash \triangleright P : \mathbf{iProp} \rrbracket_\gamma &\triangleq \lambda a. \{n \mid n = 0 \vee n - 1 \in \llbracket \Gamma \vdash P : \mathbf{iProp} \rrbracket_\gamma(a)\} \\
\llbracket \Gamma \vdash \text{Own}(t) : \mathbf{iProp} \rrbracket_\gamma &\triangleq \lambda b. \{n \mid \llbracket \Gamma \vdash t : \mathbf{M} \rrbracket_\gamma \stackrel{n}{\approx} b\} \\
\llbracket \Gamma \vdash \mathcal{V}(t) : \mathbf{iProp} \rrbracket_\gamma &\triangleq \lambda _ . \{n \mid \llbracket \Gamma \vdash t : \mathbf{M} \rrbracket_\gamma \in \mathcal{V}_n\} \\
\llbracket \Gamma \vdash \multimap P : \mathbf{iProp} \rrbracket_\gamma &\triangleq \lambda a. \left\{ n \mid \begin{array}{l} \forall m, a'. m \leq n \wedge (a \cdot a') \in \mathcal{V}_m \Rightarrow \\ \exists b. (b \cdot a') \in \mathcal{V}_m \wedge m \in \llbracket \Gamma \vdash P : \mathbf{iProp} \rrbracket_\gamma(b) \end{array} \right\}
\end{aligned}$$

For every definition, we have to show all the side-conditions: The maps have to be non-expansive and monotone.

Interpretation of non-propositional terms

$$\boxed{\llbracket \Gamma \vdash t : \tau \rrbracket : \llbracket \Gamma \rrbracket \xrightarrow{\text{ne}} \llbracket \tau \rrbracket}$$

$$\begin{aligned}
\llbracket \Gamma \vdash x : \tau \rrbracket_\gamma &\triangleq \gamma(x) \\
\llbracket \Gamma \vdash F(t_1, \dots, t_n) : \tau_{n+1} \rrbracket_\gamma &\triangleq \llbracket F \rrbracket(\llbracket \Gamma \vdash t_1 : \tau_1 \rrbracket_\gamma, \dots, \llbracket \Gamma \vdash t_n : \tau_n \rrbracket_\gamma) \\
\llbracket \Gamma \vdash \lambda x : \tau. t : \tau \rightarrow \tau' \rrbracket_\gamma &\triangleq \lambda u : \llbracket \tau \rrbracket. \llbracket \Gamma, x : \tau \vdash t : \tau \rrbracket_{\gamma[x \leftarrow u]} \\
\llbracket \Gamma \vdash t(u) : \tau' \rrbracket_\gamma &\triangleq \llbracket \Gamma \vdash t : \tau \rightarrow \tau' \rrbracket_\gamma(\llbracket \Gamma \vdash u : \tau \rrbracket_\gamma) \\
\llbracket \Gamma \vdash \mu x : \tau. t : \tau \rrbracket_\gamma &\triangleq \text{fix}(\lambda u : \llbracket \tau \rrbracket. \llbracket \Gamma, x : \tau \vdash t : \tau \rrbracket_{\gamma[x \leftarrow u]}) \\
\llbracket \Gamma \vdash () : 1 \rrbracket_\gamma &\triangleq () \\
\llbracket \Gamma \vdash (t_1, t_2) : \tau_1 \times \tau_2 \rrbracket_\gamma &\triangleq (\llbracket \Gamma \vdash t_1 : \tau_1 \rrbracket_\gamma, \llbracket \Gamma \vdash t_2 : \tau_2 \rrbracket_\gamma) \\
\llbracket \Gamma \vdash \pi_i(t) : \tau_i \rrbracket_\gamma &\triangleq \pi_i(\llbracket \Gamma \vdash t : \tau_1 \times \tau_2 \rrbracket_\gamma) \\
\llbracket a : \mathbf{M} \rrbracket_\gamma &\triangleq a \\
\llbracket \Gamma \vdash |t| : \mathbf{M} \rrbracket_\gamma &\triangleq |\llbracket \Gamma \vdash t : \mathbf{M} \rrbracket_\gamma| \\
\llbracket \Gamma \vdash t \cdot u : \mathbf{M} \rrbracket_\gamma &\triangleq \llbracket \Gamma \vdash t : \mathbf{M} \rrbracket_\gamma \cdot \llbracket \Gamma \vdash u : \mathbf{M} \rrbracket_\gamma
\end{aligned}$$

An environment Γ is interpreted as the set of finite partial functions ρ , with $\text{dom}(\rho) = \text{dom}(\Gamma)$ and $\rho(x) \in \llbracket \Gamma(x) \rrbracket$.

Logical entailment. We can now define *semantic* logical entailment.

Interpretation of entailment

$$\boxed{[\Gamma \mid \Theta \vdash P] : Prop}$$

$$\begin{aligned} [\Gamma \mid P \vdash Q] &\triangleq \forall n \in \mathbb{N}. \forall r \in M. \forall \gamma \in [\Gamma], \\ &n \in [\Gamma \vdash P : \mathbf{iProp}]_\gamma(r) \Rightarrow n \in [\Gamma \vdash Q : \mathbf{iProp}]_\gamma(r) \end{aligned}$$

The following soundness theorem connects syntactic and semantic entailment. It is proven by showing that all the syntactic proof rules of §4 can be validated in the model.

$$\Gamma \mid P \vdash Q \Rightarrow [\Gamma \mid P \vdash Q]$$

It now becomes straight-forward to show consistency of the logic.

6 Extensions of the Base Logic

In this section we discuss some additional constructions that we define within and on top of the base logic. These are not “extensions” in the sense that they change the proof power of the logic, they just form useful derived principles.

6.1 Derived rules about base connectives

We collect here some important and frequently used derived proof rules.

$$\begin{array}{l}
 P \Rightarrow Q \vdash P * Q \quad P * \exists x. Q \dashv\vdash \exists x. P * Q \quad P * \forall x. Q \vdash \forall x. P * Q \quad \Box(P * Q) \dashv\vdash \Box P * \Box Q \\
 \Box(P \Rightarrow Q) \vdash \Box P \Rightarrow \Box Q \quad \Box(P * Q) \vdash \Box P * \Box Q \quad \Box(P * Q) \dashv\vdash \Box(P \Rightarrow Q) \\
 \triangleright(P \Rightarrow Q) \vdash \triangleright P \Rightarrow \triangleright Q \quad \triangleright(P * Q) \vdash \triangleright P * \triangleright Q \quad P \vdash \triangleright P
 \end{array}$$

6.2 Persistent assertions

We call an assertion P *persistent* if $P \vdash \Box P$. These are assertions that “don’t own anything”, so we can (and will) treat them like “normal” intuitionistic assertions.

Of course, $\Box P$ is persistent for any P . Furthermore, by the proof rules given in §4.3, **True**, **False**, $t = t'$ as well as $\overline{[a]}$ and $\mathcal{V}(a)$ are persistent. Persistence is preserved by conjunction, disjunction, separating conjunction as well as universal and existential quantification and \triangleright .

6.3 Timeless assertions and except-0

One of the troubles of working in a step-indexed logic is the “later” modality \triangleright . It turns out that we can somewhat mitigate this trouble by working below the following *except-0* modality:

$$\diamond P \triangleq \triangleright \text{False} \vee P$$

This modality is useful because there is a class of assertions which we call *timeless* assertions, for which we have

$$\text{timeless}(P) \triangleq \triangleright P \vdash \diamond P$$

In other words, when working below the except-0 modality, we can *strip away* the later from timeless assertions.

The following rules can be derived about except-0:

$$\begin{array}{l}
 \text{EX0-MONO} \quad \frac{P \vdash Q}{\diamond P \vdash \diamond Q} \quad \text{EX0-INTRO} \quad P \vdash \diamond P \quad \text{EX0-IDEM} \quad \diamond \diamond P \vdash \diamond P \\
 \diamond(P * Q) \dashv\vdash \diamond P * \diamond Q \quad \diamond(P \wedge Q) \dashv\vdash \diamond P \wedge \diamond Q \quad \diamond(P \vee Q) \dashv\vdash \diamond P \vee \diamond Q \\
 \diamond \forall x. P \dashv\vdash \forall x. \diamond P \quad \diamond \exists x. P \dashv\vdash \exists x. \diamond P \\
 \diamond \Box P \dashv\vdash \Box \diamond P \quad \diamond \triangleright P \vdash \triangleright \diamond P
 \end{array}$$

The following rules identify the class of timeless assertions:

$$\begin{array}{c}
\frac{\Gamma \vdash \text{timeless}(P) \quad \Gamma \vdash \text{timeless}(Q)}{\Gamma \vdash \text{timeless}(P \wedge Q)} \qquad \frac{\Gamma \vdash \text{timeless}(P) \quad \Gamma \vdash \text{timeless}(Q)}{\Gamma \vdash \text{timeless}(P \vee Q)} \\
\\
\frac{\Gamma \vdash \text{timeless}(P) \quad \Gamma \vdash \text{timeless}(Q)}{\Gamma \vdash \text{timeless}(P * Q)} \qquad \frac{\Gamma \vdash \text{timeless}(P)}{\Gamma \vdash \text{timeless}(\Box P)} \qquad \frac{\Gamma \vdash \text{timeless}(Q)}{\Gamma \vdash \text{timeless}(P \Rightarrow Q)} \\
\\
\frac{\Gamma \vdash \text{timeless}(Q)}{\Gamma \vdash \text{timeless}(P \neg * Q)} \qquad \frac{\Gamma, x : \tau \vdash \text{timeless}(P)}{\Gamma \vdash \text{timeless}(\forall x : \tau. P)} \qquad \frac{\Gamma, x : \tau \vdash \text{timeless}(P)}{\Gamma \vdash \text{timeless}(\exists x : \tau. P)} \qquad \text{timeless}(\text{True}) \\
\\
\text{timeless}(\text{False}) \qquad \frac{t \text{ or } t' \text{ is a discrete OFE element}}{\text{timeless}(t =_{\tau} t')} \qquad \frac{a \text{ is a discrete OFE element}}{\text{timeless}(\text{Own}(a))} \\
\\
\frac{a \text{ is an element of a discrete CMRA}}{\text{timeless}(\mathcal{V}(a))}
\end{array}$$

7 Language

A language Λ consists of a set $Expr$ of *expressions* (metavariable e), a set Val of *values* (metavariable v), and a nonempty set $State$ of *states* (metavariable σ) such that

- There exist functions $\text{val_to_expr} : Val \rightarrow Expr$ and $\text{expr_to_val} : Expr \rightarrow Val$ (notice the latter is partial), such that

$$\forall e, v. \text{expr_to_val}(e) = v \Rightarrow \text{val_to_expr}(v) = e \quad \forall v. \text{expr_to_val}(\text{val_to_expr}(v)) = v$$

- There exists a *primitive reduction relation*

$$(-, - \rightarrow_{\mathfrak{t}} -, -, -) \subseteq Expr \times State \times Expr \times State \times List(Expr)$$

A reduction $e_1, \sigma_1 \rightarrow_{\mathfrak{t}} e_2, \sigma_2, \bar{e}$ indicates that, when e_1 reduces to e_2 , the new threads in the list \bar{e} is forked off. We will write $e_1, \sigma_1 \rightarrow_{\mathfrak{t}} e_2, \sigma_2$ for $e_1, \sigma_1 \rightarrow_{\mathfrak{t}} e_2, \sigma_2, ()$, *i.e.*, when no threads are forked off.

- All values are stuck:

$$e, _ \rightarrow_{\mathfrak{t}} _, _, _ \Rightarrow \text{expr_to_val}(e) = \perp$$

Definition 19. An expression e and state σ are *reducible* (written $\text{red}(e, \sigma)$) if

$$\exists e_2, \sigma_2, \bar{e}. e, \sigma \rightarrow_{\mathfrak{t}} e_2, \sigma_2, \bar{e}$$

Definition 20. An expression e is *atomic* if it reduces in one step to something irreducible:

$$\forall \sigma_1, e_2, \sigma_2, \bar{e}. e, \sigma_1 \rightarrow_{\mathfrak{t}} e_2, \sigma_2, \bar{e} \Rightarrow \neg \text{red}(e_2, \sigma_2)$$

Definition 21 (Context). A function $K : Expr \rightarrow Expr$ is a *context* if the following conditions are satisfied:

1. K does not turn non-values into values:

$$\forall e. \text{expr_to_val}(e) = \perp \Rightarrow \text{expr_to_val}(K(e)) = \perp$$

2. One can perform reductions below K :

$$\forall e_1, \sigma_1, e_2, \sigma_2, \bar{e}. e_1, \sigma_1 \rightarrow_{\mathfrak{t}} e_2, \sigma_2, \bar{e} \Rightarrow K(e_1), \sigma_1 \rightarrow_{\mathfrak{t}} K(e_2), \sigma_2, \bar{e}$$

3. Reductions stay below K until there is a value in the hole:

$$\forall e'_1, \sigma_1, e_2, \sigma_2, \bar{e}. \text{expr_to_val}(e'_1) = \perp \wedge K(e'_1), \sigma_1 \rightarrow_{\mathfrak{t}} e_2, \sigma_2, \bar{e} \Rightarrow \exists e'_2. e_2 = K(e'_2) \wedge e'_1, \sigma_1 \rightarrow_{\mathfrak{t}} e'_2, \sigma_2, \bar{e}$$

7.1 Concurrent language

For any language Λ , we define the corresponding thread-pool semantics.

Machine syntax

$$T \in \text{ThreadPool} \triangleq List(Expr)$$

Machine reduction

$$\boxed{T; \sigma \rightarrow_{\mathfrak{t}} T'; \sigma'}$$

$$\frac{e_1, \sigma_1 \rightarrow_{\mathfrak{t}} e_2, \sigma_2, \bar{e}}{T \uparrow [e_1] \uparrow T'; \sigma_1 \rightarrow_{\mathfrak{t}} T \uparrow [e_2] \uparrow T' \uparrow \bar{e}; \sigma_2}$$

8 Program Logic

This section describes how to build a program logic for an arbitrary language (*c.f.* §7) on top of the base logic. So in the following, we assume that some language Λ was fixed.

8.1 Dynamic Composeable Higher-Order Resources

The base logic described in §4 works over an arbitrary CMRA M defining the structure of the resources. It turns out that we can generalize this further and permit picking CMRAs “ $\Sigma(\text{iProp})$ ” that depend on the structure of assertions themselves. Of course, iProp is just the syntactic type of assertions; for this to make sense we have to look at the semantics.

Furthermore, there is a composability problem with the given logic: if we have one proof performed with CMRA M_1 , and another proof carried out with a *different* CMRA M_2 , then the two proofs are actually carried out in two *entirely separate logics* and hence cannot be combined.

Finally, in many cases just having a single “instance” of a CMRA available for reasoning is not enough. For example, when reasoning about a dynamically allocated data structure, every time a new instance of that data structure is created, we will want a fresh resource governing the state of this particular instance. While it would be possible to handle this problem whenever it comes up, it turns out to be useful to provide a general solution.

The purpose of this section is to describe how we solve these issues.

Picking the resources. The key ingredient that we will employ on top of the base logic is to give some more fixed structure to the resources. To instantiate the program logic, the user picks a family of locally contractive bifunctors $(\Sigma_i : \mathcal{OFE} \rightarrow \mathcal{CMRA})_{i \in \mathcal{I}}$. (This is in contrast to the base logic, where the user picks a single, fixed CMRA that has a unit.)

From this, we construct the bifunctor defining the overall resources as follows:

$$\begin{aligned} GName &\triangleq \mathbb{N} \\ ResF(T^{\text{op}}, T) &\triangleq \prod_{i \in \mathcal{I}} GName \xrightarrow{\text{fin}} \Sigma_i(T^{\text{op}}, T) \end{aligned}$$

We will motivate both the use of a product and the finite partial function below. $ResF(T^{\text{op}}, T)$ is a CMRA by lifting the individual CMRAs pointwise, and it has a unit (using the empty finite partial functions). Furthermore, since the Σ_i are locally contractive, so is $ResF$.

Now we can write down the recursive domain equation:

$$iPreProp \cong UPred(ResF(iPreProp, iPreProp))$$

$iPreProp$ is a COFE defined as the fixed-point of a locally contractive bifunctor. This fixed-point exists and is unique² by America and Rutten’s theorem [1, 3]. We do not need to consider how the object is constructed. We only need the isomorphism, given by

$$\begin{aligned} Res &\triangleq ResF(iPreProp, iPreProp) \\ iProp &\triangleq UPred(Res) \\ \xi &: iProp \xrightarrow{\text{ne}} iPreProp \\ \xi^{-1} &: iPreProp \xrightarrow{\text{ne}} iProp \end{aligned}$$

²We have not proven uniqueness in Coq.

Notice that $iProp$ is the semantic model of assertions for the base logic described in §4 with Res :

$$\llbracket \text{iProp} \rrbracket \triangleq iProp = UPred(Res)$$

Effectively, we just defined a way to instantiate the base logic with Res as the CMRA of resources, while providing a way for Res to depend on $iPreProp$, which is isomorphic to $\llbracket \text{iProp} \rrbracket$.

We thus obtain all the rules of §4, and furthermore, we can use the maps ξ and ξ^{-1} in the logic to convert between logical assertions $\llbracket \text{iProp} \rrbracket$ and the domain $iPreProp$ which is used in the construction of Res – so from elements of $iPreProp$, we can construct elements of $\llbracket \text{M} \rrbracket$, which are the elements that can be owned in our logic.

Proof composability. To make our proofs composable, we *generalize* our proofs over the family of functors. This is possible because we made Res a *product* of all the CMRA’s picked by the user, and because we can actually work with that product “pointwise”. So instead of picking a *concrete* family, proofs will assume to be given an *arbitrary* family of functors, plus a proof that this family *contains the functors they need*. Composing two proofs is then merely a matter of conjoining the assumptions they make about the functors. Since the logic is entirely parametric in the choice of functors, there is no trouble reasoning without full knowledge of the family of functors.

Only when the top-level proof is completed we will “close” the proof by picking a concrete family that contains exactly those functors the proof needs.

Dynamic resources. Finally, the use of finite partial functions lets us have as many instances of any CMRA as we could wish for: Because there can only ever be finitely many instances already allocated, it is always possible to create a fresh instance with any desired (valid) starting state. This is best demonstrated by giving some proof rules.

So let us first define the notion of ghost ownership that we use in this logic. Assuming that the family of functors contains the functor Σ_i at index i , and furthermore assuming that $M_i = \Sigma_i(iPreProp, iPreProp)$, given some $a \in M_i$ we define:

$$\{\!\!\{ a : M_i \}\!\!\}^\gamma \triangleq \text{Own}((\dots, \emptyset, i : [\gamma \leftarrow a], \emptyset, \dots))$$

This is ownership of the pair (element of the product over all the functors) that has the empty finite partial function in all components *except for* the component corresponding to index i , where we own the element a at index γ in the finite partial function.

We can show the following properties for this form of ownership:

$$\begin{array}{c} \text{RES-ALLOC} \\ \frac{G \text{ infinite} \quad a \in \mathcal{V}_{M_i}}{\text{True} \vdash \exists \gamma \in G. \{\!\!\{ a : M_i \}\!\!\}^\gamma} \end{array} \quad \begin{array}{c} \text{RES-UPDATE} \\ \frac{a \rightsquigarrow_{M_i} B}{\{\!\!\{ a : M_i \}\!\!\}^\gamma \vdash \exists b \in B. \{\!\!\{ b : M_i \}\!\!\}^\gamma} \end{array} \quad \begin{array}{c} \text{RES-EMPTY} \\ \frac{\varepsilon \text{ is a unit of } M_i}{\text{True} \vdash \exists \varepsilon. \{\!\!\{ \varepsilon \}\!\!\}^\gamma} \end{array}$$

$$\begin{array}{c} \text{RES-OP} \\ \{\!\!\{ a : M_i \}\!\!\}^\gamma * \{\!\!\{ b : M_i \}\!\!\}^\gamma \dashv\vdash \{\!\!\{ a \cdot b : M_i \}\!\!\}^\gamma \end{array} \quad \begin{array}{c} \text{RES-VALID} \\ \{\!\!\{ a : M_i \}\!\!\}^\gamma \Rightarrow \mathcal{V}_{M_i}(a) \end{array} \quad \begin{array}{c} \text{RES-TIMELESS} \\ \frac{a \text{ is a discrete OFE element}}{\text{timeless}(\{\!\!\{ a : M_i \}\!\!\}^\gamma)} \end{array}$$

Below, we will always work within (an instance of) the logic as described here. Whenever a CMRA is used in a proof, we implicitly assume it to be available in the global family of functors. We will typically leave the M_i implicit when asserting ghost ownership, as the type of a will be clear from the context.

8.2 World Satisfaction, Invariants, Fancy Updates

To introduce invariants into our logic, we will define weakest precondition to explicitly thread through the proof that all the invariants are maintained throughout program execution. However, in order to be able to access invariants, we will also have to provide a way to *temporarily disable* (or “open”) them. To this end, we use tokens that manage which invariants are currently enabled.

We assume to have the following four CMRAs available:

$$\begin{aligned}
\text{InvName} &\triangleq \mathbb{N} \\
\text{INV} &\triangleq \text{AUTH}(\text{InvName} \stackrel{\text{fin}}{\multimap} \text{AG}(\blacktriangleright i\text{PreProp})) \\
\text{EN} &\triangleq \wp(\text{InvName}) \\
\text{DIS} &\triangleq \wp^{\text{fin}}(\text{InvName})
\end{aligned}$$

The last two are the tokens used for managing invariants, INV is the monoid used to manage the invariants themselves.

We assume that at the beginning of the verification, instances named γ_{STATE} , γ_{INV} , γ_{EN} and γ_{DIS} of these CMRAs have been created, such that these names are globally known.

World Satisfaction. We can now define the assertion W (*world satisfaction*) which ensures that the enabled invariants are actually maintained:

$$W \triangleq \exists I : \text{InvName} \stackrel{\text{fin}}{\multimap} \text{iProp}. \left[\bullet \left[\iota \leftarrow \text{ag}(\text{next}(\xi(I(\iota)))) \mid \iota \in \text{dom}(I) \right] \right]^{\gamma_{\text{INV}}} * \left[\blacktriangleright_{\iota \in \text{dom}(I)} \left(I(\iota) * \left[\left[\iota \right] \right]^{\gamma_{\text{DIS}}} \vee \left[\left[\iota \right] \right]^{\gamma_{\text{EN}}} \right) \right]$$

Invariants. The following assertion states that an invariant with name ι exists and maintains assertion P :

$$\boxed{P}^\iota \triangleq \left[\circ \left[\iota \leftarrow \text{ag}(\text{next}(\xi(P))) \right] \right]^{\gamma_{\text{INV}}}$$

Fancy Updates and View Shifts. Next, we define *fancy updates*, which are essentially the same as the basic updates of the base logic (§4), except that they also have access to world satisfaction and can enable and disable invariants:

$$\mathcal{E}_1 \stackrel{\mathcal{E}_2}{\Rightarrow} P \triangleq W * \left[\left[\mathcal{E}_1 \right] \right]^{\gamma_{\text{EN}}} \multimap \Rightarrow \diamond (W * \left[\left[\mathcal{E}_2 \right] \right]^{\gamma_{\text{EN}}} * P)$$

Here, \mathcal{E}_1 and \mathcal{E}_2 are the *masks* of the view update, defining which invariants have to be (at least!) available before and after the update. We use \top as symbol for the largest possible mask, \mathbb{N} , and \perp for the smallest possible mask \emptyset . We will write $\stackrel{\mathcal{E}}{\Rightarrow} P$ for $\top \stackrel{\mathcal{E}}{\Rightarrow} P$. Fancy updates satisfy the following basic proof rules:

$$\begin{array}{c}
\text{FUP-MONO} \\
\frac{P \vdash Q}{\mathcal{E}_1 \stackrel{\mathcal{E}_2}{\Rightarrow} P \vdash \mathcal{E}_1 \stackrel{\mathcal{E}_2}{\Rightarrow} Q}
\end{array}
\quad
\begin{array}{c}
\text{FUP-INTRO-MASK} \\
\frac{\mathcal{E}_2 \subseteq \mathcal{E}_1}{P \vdash \mathcal{E}_1 \stackrel{\mathcal{E}_2}{\Rightarrow} \mathcal{E}_2 \stackrel{\mathcal{E}_1}{\Rightarrow} P}
\end{array}
\quad
\begin{array}{c}
\text{FUP-TRANS} \\
\frac{\mathcal{E}_1 \stackrel{\mathcal{E}_2}{\Rightarrow} \mathcal{E}_2 \stackrel{\mathcal{E}_3}{\Rightarrow} P \vdash \mathcal{E}_1 \stackrel{\mathcal{E}_3}{\Rightarrow} P}{\mathcal{E}_1 \stackrel{\mathcal{E}_2}{\Rightarrow} \mathcal{E}_2 \stackrel{\mathcal{E}_3}{\Rightarrow} P \vdash \mathcal{E}_1 \stackrel{\mathcal{E}_3}{\Rightarrow} P}
\end{array}
\quad
\begin{array}{c}
\text{FUP-UPD} \\
\frac{}{\Rightarrow P \vdash \Rightarrow_{\mathcal{E}} P}
\end{array}$$

$$\begin{array}{c}
\text{FUP-FRAME} \\
\frac{}{Q * \mathcal{E}_1 \stackrel{\mathcal{E}_2}{\Rightarrow} P \vdash \mathcal{E}_1 \uplus \mathcal{E}_f \stackrel{\mathcal{E}_2 \uplus \mathcal{E}_f}{\Rightarrow} Q * P}
\end{array}
\quad
\begin{array}{c}
\text{FUP-UPDATE} \\
\frac{a \rightsquigarrow B}{\text{Own}(a) \vdash \Rightarrow_{\mathcal{E}} \exists b \in B. \text{Own}(b)}
\end{array}
\quad
\begin{array}{c}
\text{FUP-TIMELESS} \\
\frac{\text{timeless}(P)}{\triangleright P \vdash \Rightarrow_{\mathcal{E}} P}
\end{array}$$

(There are no rules related to invariants here. Those rules will be discussed later, in §8.4.)

We can further define the notions of *view shifts* and *linear view shifts*:

$$\begin{aligned} P \varepsilon_1 \Rightarrow^* \varepsilon_2 Q &\triangleq P * \varepsilon_1 \Rightarrow^* \varepsilon_2 Q \\ P \varepsilon_1 \Rightarrow \varepsilon_2 Q &\triangleq \Box(P * \varepsilon_1 \Rightarrow^* \varepsilon_2 Q) \\ P \Rightarrow_{\varepsilon} Q &\triangleq P \varepsilon \Rightarrow^{\varepsilon} Q \end{aligned}$$

These two are useful when writing down specifications and for comparing with previous versions of Iris, but for reasoning, it is typically easier to just work directly with fancy updates. Still, just to give an idea of what view shifts “are”, here are some proof rules for them:

$$\begin{array}{c} \text{VS-UPDATE} \\ \frac{a \rightsquigarrow B}{\llbracket a \rrbracket^{\gamma} \Rightarrow_{\emptyset} \exists b \in B. \llbracket b \rrbracket^{\gamma}} \end{array} \quad \begin{array}{c} \text{VS-TRANS} \\ \frac{P \varepsilon_1 \Rightarrow^{\varepsilon_2} Q \quad Q \varepsilon_2 \Rightarrow^{\varepsilon_3} R}{P \varepsilon_1 \Rightarrow^{\varepsilon_3} R} \end{array} \quad \begin{array}{c} \text{VS-IMP} \\ \frac{\Box(P \Rightarrow Q)}{P \Rightarrow_{\emptyset} Q} \end{array} \quad \begin{array}{c} \text{VS-MASK-FRAME} \\ \frac{P \varepsilon_1 \Rightarrow^{\varepsilon_2} Q}{P \varepsilon_1 \uplus \varepsilon' \Rightarrow^{\varepsilon_2 \uplus \varepsilon'} Q} \end{array}$$

$$\begin{array}{c} \text{VS-FRAME} \\ \frac{P \varepsilon_1 \Rightarrow^{\varepsilon_2} Q}{P * R \varepsilon_1 \Rightarrow^{\varepsilon_2} Q * R} \end{array} \quad \begin{array}{c} \text{VS-TIMELESS} \\ \frac{\text{timeless}(P)}{\triangleright P \Rightarrow_{\emptyset} P} \end{array} \quad \begin{array}{c} \text{VS-DISJ} \\ \frac{P \varepsilon_1 \Rightarrow^{\varepsilon_2} R \quad Q \varepsilon_1 \Rightarrow^{\varepsilon_2} R}{P \vee Q \varepsilon_1 \Rightarrow^{\varepsilon_2} R} \end{array} \quad \begin{array}{c} \text{VS-EXIST} \\ \frac{\forall x. (P \varepsilon_1 \Rightarrow^{\varepsilon_2} Q)}{(\exists x. P) \varepsilon_1 \Rightarrow^{\varepsilon_2} Q} \end{array}$$

$$\begin{array}{c} \text{VS-ALWAYS} \\ \frac{\Box Q \vdash P \varepsilon_1 \Rightarrow^{\varepsilon_2} R}{P \wedge \Box Q \varepsilon_1 \Rightarrow^{\varepsilon_2} R} \end{array} \quad \begin{array}{c} \text{VS-FALSE} \\ \text{False} \varepsilon_1 \Rightarrow^{\varepsilon_2} P \end{array}$$

8.3 Weakest Precondition

Finally, we can define the core piece of the program logic, the assertion that reasons about program behavior: Weakest precondition, from which Hoare triples will be derived.

Defining weakest precondition. We assume that everything making up the definition of the language, *i.e.*, values, expressions, states, the conversion functions, reduction relation and all their properties, are suitably reflected into the logic (*i.e.*, they are part of the signature \mathcal{S}). We further assume (as a parameter) a predicate $I : \text{State} \rightarrow iProp$ that interprets the physical state as an Iris assertion. This can be instantiated, for example, with ownership of an authoritative RA to tie the physical state to fragments that are used for user-level proofs.

$$\begin{aligned} wp &\triangleq \mu wp. \lambda \mathcal{E}, e, \Phi. \\ &(\exists v. \text{expr_to_val}(e) = v \wedge \varepsilon \Rightarrow^* \Phi(v)) \vee \\ &\left(\text{expr_to_val}(e) = \perp \wedge \forall \sigma. I(\sigma) \varepsilon \Rightarrow^* \emptyset \right. \\ &\quad \left. \text{red}(e, \sigma) * \triangleright \forall e', \sigma', \vec{e}. (e, \sigma \rightarrow_{\tau} e', \sigma', \vec{e}) \emptyset \Rightarrow^* \varepsilon \right. \\ &\quad \left. I(\sigma') * wp(\mathcal{E}, e', \Phi) * \bigstar_{e'' \in \vec{e}} wp(\top, e'', \lambda _ . \text{True}) \right) \\ wp_{\varepsilon} e \{v. P\} &\triangleq wp(\mathcal{E}, e, \lambda v. P) \end{aligned}$$

If we leave away the mask, we assume it to default to \top .

Laws of weakest precondition. The following rules can all be derived:

$$\begin{array}{c}
\text{WP-VALUE} \\
\frac{}{P[v/x] \vdash \text{wp}_{\mathcal{E}} v \{x. P\}} \\
\\
\text{WP-MONO} \\
\frac{\mathcal{E}_1 \subseteq \mathcal{E}_2 \quad \Gamma, x : \text{val} \mid P \vdash Q}{\Gamma \mid \text{wp}_{\mathcal{E}_1} e \{x. P\} \vdash \text{wp}_{\mathcal{E}_2} e \{x. Q\}} \\
\\
\text{FUP-WP} \\
\frac{}{\models_{\mathcal{E}} \text{wp}_{\mathcal{E}} e \{x. P\} \vdash \text{wp}_{\mathcal{E}} e \{x. P\}} \\
\\
\text{WP-FUP} \\
\frac{}{\text{wp}_{\mathcal{E}} e \{x. \models_{\mathcal{E}} P\} \vdash \text{wp}_{\mathcal{E}} e \{x. P\}} \\
\\
\text{WP-ATOMIC} \\
\frac{\text{atomic}(e)}{\mathcal{E}_1 \models_{\mathcal{E}_2} \text{wp}_{\mathcal{E}_2} e \{x. \mathcal{E}_2 \models_{\mathcal{E}_1} P\} \vdash \text{wp}_{\mathcal{E}_1} e \{x. P\}} \\
\\
\text{WP-FRAME} \\
\frac{}{Q * \text{wp}_{\mathcal{E}} e \{x. P\} \vdash \text{wp}_{\mathcal{E}} e \{x. Q * P\}} \\
\\
\text{WP-FRAME-STEP} \\
\frac{\text{expr_to_val}(e) = \perp \quad \mathcal{E}_2 \subseteq \mathcal{E}_1}{\text{wp}_{\mathcal{E}_2} e \{x. P\} * \mathcal{E}_1 \models_{\mathcal{E}_2} \triangleright \mathcal{E}_2 \models_{\mathcal{E}_1} Q \vdash \text{wp}_{\mathcal{E}_1} e \{x. Q * P\}} \\
\\
\text{WP-BIND} \\
\frac{K \text{ is a context}}{\text{wp}_{\mathcal{E}} e \{x. \text{wp}_{\mathcal{E}} K(\text{val_to_expr}(x)) \{y. P\}\} \vdash \text{wp}_{\mathcal{E}} K(e) \{y. P\}}
\end{array}$$

We will also want a rule that connect weakest preconditions to the operational semantics of the language.

$$\begin{array}{c}
\text{WP-LIFT-STEP} \\
\frac{\text{expr_to_val}(e_1) = \perp}{\forall \sigma_1. I(\sigma_1) \overset{\mathcal{E}}{\equiv}^{\emptyset} \text{red}(e_1, \sigma_1) * \triangleright \forall e_2, \sigma_2, \vec{e}. (e_1, \sigma_1 \rightarrow_t e_2, \sigma_2, \vec{e}) \overset{\emptyset}{\equiv}^{\mathcal{E}} \left(I(\sigma_2) * \text{wp}_{\mathcal{E}} e_2 \{x. P\} * \bigstar_{e_f \in \vec{e}} \text{wp}_{\top} e_f \{ _ . \text{True} \} \right) \vdash \text{wp}_{\mathcal{E}} e_1 \{x. P\}}
\end{array}$$

Adequacy of weakest precondition. The purpose of the adequacy statement is to show that our notion of weakest preconditions is *realistic* in the sense that it actually has anything to do with the actual behavior of the program. There are two properties we are looking for: First of all, the postcondition should reflect actual properties of the values the program can terminate with. Second, a proof of a weakest precondition with any postcondition should imply that the program is *safe*, *i.e.*, that it does not get stuck.

Definition 22 (Adequacy). *A program e in some initial state σ is adequate for a set $V \subseteq \text{Val}$ of legal return values ($e, \sigma \vDash V$) if for all T', σ' such that $([e], \sigma) \rightarrow_{\text{tp}}^* (T', \sigma')$ we have*

1. *Safety: For any $e' \in T'$ we have that either e' is a value, or $\text{red}(e', \sigma')$:*

$$\forall e' \in T'. \text{expr_to_val}(e') \neq \perp \vee \text{red}(e', \sigma')$$

Notice that this is stronger than saying that the thread pool can reduce; we actually assert that every non-finished thread can take a step.

2. *Legal return value: If T'_1 (the main thread) is a value v' , then $v' \in V$:*

$$\forall v', T''. T' = [v'] \uparrow T'' \Rightarrow v' \in V$$

To express the adequacy statement for functional correctness, we assume that the signature \mathcal{S} adds a predicate Φ to the logic:

$$\Phi : \text{Val} \rightarrow \text{iProp} \in \mathcal{F}$$

Furthermore, we assume that the *interpretation* $\llbracket \Phi \rrbracket$ of Φ reflects some set V of legal return values into the logic (also see §5):

$$\begin{aligned} \llbracket \Phi \rrbracket & : \llbracket \text{Val} \rrbracket \xrightarrow{\text{ne}} \llbracket \text{iProp} \rrbracket \\ \llbracket \Phi \rrbracket & \triangleq \lambda v. \lambda _ . \{n \mid v \in V\} \end{aligned}$$

The signature can of course state arbitrary additional properties of Φ , as long as they are proven sound. The adequacy statement now reads as follows:

$$\begin{aligned} & \forall \mathcal{E}, e, v, \sigma. \\ & (\text{True} \vdash \models_{\mathcal{E}} \exists I. I(\sigma) * \text{wp}_{\mathcal{E}} e \{x. \Phi(x)\}) \Rightarrow \\ & e, \sigma \models V \end{aligned}$$

Notice that the state invariant S used by the weakest precondition is chosen *after* doing a fancy update, which allows it to depend on the names of ghost variables that are picked in that initial fancy update.

Hoare triples. It turns out that weakest precondition is actually quite convenient to work with, in particular when performing these proofs in Coq. Still, for a more traditional presentation, we can easily derive the notion of a Hoare triple:

$$\{P\} e \{v. Q\}_{\mathcal{E}} \triangleq \square (P \multimap \text{wp}_{\mathcal{E}} e \{v. Q\})$$

We only give some of the proof rules for Hoare triples here, since we usually do all our reasoning directly with weakest preconditions and use Hoare triples only to write specifications.

$$\begin{array}{c} \text{HT-BIND} \\ \frac{K \text{ is a context} \quad \{P\} e \{v. Q\}_{\mathcal{E}} \quad \forall v. \{Q\} K(v) \{w. R\}_{\mathcal{E}}}{\{P\} K(e) \{w. R\}_{\mathcal{E}}} \\ \text{HT-RET} \\ \frac{}{\{\text{True}\} w \{v. v = w\}_{\mathcal{E}}} \\ \text{HT-CSQ} \\ \frac{P \Rightarrow P' \quad \{P'\} e \{v. Q'\}_{\mathcal{E}} \quad \forall v. Q' \Rightarrow Q}{\{P\} e \{v. Q\}_{\mathcal{E}}} \\ \text{HT-FRAME} \\ \frac{\{P\} e \{v. Q\}_{\mathcal{E}}}{\{P * R\} e \{v. Q * R\}_{\mathcal{E}}} \\ \text{HT-ATOMIC} \\ \frac{P \xrightarrow{\mathcal{E} \circ \mathcal{E}'} \Rightarrow^{\mathcal{E}} P' \quad \{P'\} e \{v. Q'\}_{\mathcal{E}} \quad \forall v. Q' \xrightarrow{\mathcal{E} \circ \mathcal{E}'} \Rightarrow^{\mathcal{E} \circ \mathcal{E}'} Q \quad \text{atomic}(e)}{\{P\} e \{v. Q\}_{\mathcal{E} \circ \mathcal{E}'}} \\ \text{HT-FALSE} \\ \frac{}{\{\text{False}\} e \{v. P\}_{\mathcal{E}}} \\ \text{HT-DISJ} \\ \frac{\{P\} e \{v. R\}_{\mathcal{E}} \quad \{Q\} e \{v. R\}_{\mathcal{E}}}{\{P \vee Q\} e \{v. R\}_{\mathcal{E}}} \\ \text{HT-EXIST} \\ \frac{\forall x. \{P\} e \{v. Q\}_{\mathcal{E}}}{\{\exists x. P\} e \{v. Q\}_{\mathcal{E}}} \\ \text{HT-BOX} \\ \frac{\square Q \vdash \{P\} e \{v. R\}_{\mathcal{E}}}{\{P \wedge \square Q\} e \{v. R\}_{\mathcal{E}}} \end{array}$$

8.4 Invariant Namespaces

In §8.2, we defined an assertion \boxed{P}^ι expressing knowledge (*i.e.*, the assertion is persistent) that P is maintained as invariant with name ι . The concrete name ι is picked when the invariant is allocated, so it cannot possibly be statically known – it will always be a variable that’s threaded through everything. However, we hardly care about the actual, concrete name. All we need to know is that this name is *different* from the names of other invariants that we want to open at the same time. Keeping track of the n^2 mutual inequalities that arise with n invariants quickly gets in the way of the actual proof.

To solve this issue, instead of remembering the exact name picked for an invariant, we will keep track of the *namespace* the invariant was allocated in. Namespaces are sets of invariants, following a tree-like structure: Think of the name of an invariant as a sequence of identifiers, much like a fully qualified Java class name. A *namespace* \mathcal{N} then is like a Java package: it is a sequence of identifiers that we think of as *containing* all invariant names that begin with this sequence. For example, `org.mpi-sws.iris` is a namespace containing the invariant name `org.mpi-sws.iris.heap`.

The crux is that all namespaces contain infinitely many invariants, and hence we can *freely pick* the namespace an invariant is allocated in – no further, unpredictable choice has to be made. Furthermore, we will often know that namespaces are *disjoint* just by looking at them. The namespaces $\mathcal{N}.\text{iris}$ and $\mathcal{N}.\text{gps}$ are disjoint no matter the choice of \mathcal{N} . As a result, there is often no need to track disjointness of namespaces, we just have to pick the namespaces that we allocate our invariants in accordingly.

Formally speaking, let $\mathcal{N} \in \text{InvNamesp} \triangleq \text{List}(\mathbb{N})$ be the type of *invariant namespaces*. We use the notation $\mathcal{N}.\iota$ for the namespace $[\iota] \uplus \mathcal{N}$. (In other words, the list is “backwards”. This is because cons-ing to the list, like the dot does above, is easier to deal with in Coq than appending at the end.)

The elements of a namespaces are *structured invariant names* (think: Java fully qualified class name). They, too, are lists of \mathbb{N} , the same type as namespaces. In order to connect this up to the definitions of §8.2, we need a way to map structured invariant names to *InvName*, the type of “plain” invariant names. Any injective mapping `namesp_inj` will do; and such a mapping has to exist because $\text{List}(\mathbb{N})$ is countable and *InvName* is infinite. Whenever needed, we (usually implicitly) coerce \mathcal{N} to its encoded suffix-closure, *i.e.*, to the set of encoded structured invariant names contained in the namespace:

$$\mathcal{N}^\uparrow \triangleq \{\iota \mid \exists \mathcal{N}'. \iota = \text{namesp_inj}(\mathcal{N}' \uplus \mathcal{N})\}$$

We will overload the notation for invariant assertions for using namespaces instead of names:

$$\boxed{P}^\mathcal{N} \triangleq \exists \iota \in \mathcal{N}^\uparrow. \boxed{P}^\iota$$

We can now derive the following rules (this involves unfolding the definition of fancy updates):

$$\begin{array}{c}
\text{INV-PERSIST} \\
\boxed{P}^\mathcal{N} \vdash \square \boxed{P}^\mathcal{N} \\
\\
\text{INV-ALLOC} \\
\triangleright P \vdash \Rightarrow_{\emptyset} \boxed{P}^\mathcal{N} \\
\\
\text{INV-OPEN} \\
\frac{\mathcal{N} \subseteq \mathcal{E}}{\boxed{P}^\mathcal{N} \xrightarrow{\mathcal{E} \Rightarrow \mathcal{E} \setminus \mathcal{N}} \triangleright P * (\triangleright P \xrightarrow{\mathcal{E} \setminus \mathcal{N}} \star^\mathcal{E} \text{True})} \\
\\
\text{INV-OPEN-TIMELESS} \\
\frac{\mathcal{N} \subseteq \mathcal{E} \quad \text{timeless}(P)}{\boxed{P}^\mathcal{N} \xrightarrow{\mathcal{E} \Rightarrow \mathcal{E} \setminus \mathcal{N}} P * (P \xrightarrow{\mathcal{E} \setminus \mathcal{N}} \star^\mathcal{E} \text{True})}
\end{array}$$

8.5 Accessors

The two rules **INV-OPEN** and **INV-OPEN-TIMELESS** above may look a little surprising, in the sense that it is not clear on first sight how they would be applied. The rules are the first *accessors* that show up in this document. Accessors are assertions of the form

$$P \xrightarrow{\mathcal{E}_1 \Rightarrow \mathcal{E}_2} \exists x. Q * (\forall y. Q' \xrightarrow{\mathcal{E}_2} \star^{\mathcal{E}_1} R)$$

One way to think about such assertions is as follows: Given some accessor, if during our verification we have the assertion P and the mask \mathcal{E}_1 available, we can use the accessor to *access* Q and obtain the witness x . We call this *opening* the accessor, and it changes the mask to \mathcal{E}_2 . Additionally, opening the accessor provides us with $\forall y. Q' \xrightarrow{\mathcal{E}_2} \star^{\mathcal{E}_1} R$, a *linear view shift* (i.e., a view shift that can only be used once). This linear view shift tells us that in order to *close* the accessor again and go back to mask \mathcal{E}_1 , we have to pick some y and establish the corresponding Q' . After closing, we will obtain R .

Using **VS-TRANS** and **HT-ATOMIC** (or the corresponding proof rules for fancy updates and weakest preconditions), we can show that it is possible to open an accessor around any view shift and any *atomic* expression:

$$\frac{\text{ACC-VS} \quad P \xrightarrow{\mathcal{E}_1 \Rightarrow \mathcal{E}_2} \exists x. Q * (\forall y. Q' \xrightarrow{\mathcal{E}_2} \star^{\mathcal{E}_1} R) \quad \forall x. Q * P_F \Rightarrow_{\mathcal{E}_2} \exists y. Q' * P_F}{P * P_F \Rightarrow_{\mathcal{E}_1} R * P_F}$$

$$\frac{\text{ACC-HT} \quad P \xrightarrow{\mathcal{E}_1 \Rightarrow \mathcal{E}_2} \exists x. Q * (\forall y. Q' \xrightarrow{\mathcal{E}_2} \star^{\mathcal{E}_1} R) \quad \forall x. \{Q * P_F\} e \{\exists y. Q' * P_F\}_{\mathcal{E}_2} \quad \text{atomic}(e)}{\{P * P_F\} e \{R * P_F\}_{\mathcal{E}_1}}$$

Furthermore, in the special case that $\mathcal{E}_1 = \mathcal{E}_2$, the accessor can be opened around *any* expression. For this reason, we also call such accessors *non-atomic*.

The reasons accessors are useful is that they let us talk about “opening X” (e.g., “opening invariants”) without having to care what X is opened around. Furthermore, as we construct more sophisticated and more interesting things that can be opened (e.g., invariants that can be “cancelled”, or STSs), accessors become a useful interface that allows us to mix and match different abstractions in arbitrary ways.

For the special case that $P = R$ and $Q = Q'$, we use the following notation that avoids repetition:

$$\langle P \Leftrightarrow x.Q \rangle_{\mathcal{E}_1}^{\mathcal{E}_2} \triangleq P \xrightarrow{\mathcal{E}_1 \Rightarrow \mathcal{E}_2} \exists x. Q * (Q \xrightarrow{\mathcal{E}_2} \star^{\mathcal{E}_1} P)$$

This accessor is “idempotent” in the sense that it doesn’t actually change the state. After applying it, we get our P back so we end up where we started.

9 Derived constructions

9.1 Non-atomic (“thread-local”) invariants

Sometimes it is necessary to maintain invariants that we need to open non-atomically. Clearly, for this mechanism to be sound we need something that prevents us from opening the same invariant twice, something like the masks that avoid reentrancy on the “normal”, atomic invariants. The idea is to use tokens³ that guard access to non-atomic invariants. Having the token $[\text{NaInv} : p.\mathcal{E}]$ indicates that we can open all invariants in \mathcal{E} . The p here is the name of the *invariant pool*. This mechanism allows us to have multiple, independent pools of invariants that all have their own namespaces.

One way to think about non-atomic invariants is as “thread-local invariants”, where every pool is a thread. Every thread thus has its own, independent set of invariants. Every thread threads through all the tokens for its own pool, so that each invariant can only be opened in the thread it belongs to. As a consequence, they can be kept open around any sequence of expressions (*i.e.*, there is no restriction to atomic expressions) – after all, there cannot be any races with other threads.

Concretely, this is the monoid structure we need:

$$\begin{aligned} PId &\triangleq GName \\ \text{NATOK} &\triangleq \wp^{\text{fin}}(\text{InvName}) \times \wp(\text{InvName}) \end{aligned}$$

For every pool, there is a set of tokens designating which invariants are *enabled* (closed). This corresponds to the mask of “normal” invariants. We re-use the structure given by namespaces for non-atomic invariants. Furthermore, there is a *finite* set of invariants that is *disabled* (open).

Owning tokens is defined as follows:

$$\begin{aligned} [\text{NaInv} : p.\mathcal{E}] &\triangleq \{ \overline{\overline{\overline{\overline{\emptyset, \mathcal{E}}}} \}^p \\ [\text{NaInv} : p] &\triangleq [\text{NaInv} : p.\top] \end{aligned}$$

Next, we define non-atomic invariants. To simplify this construction, we piggy-back into “normal” invariants.

$$\text{NaInv}^{p.\mathcal{N}}(P) \triangleq \exists \iota \in \mathcal{N}. \boxed{P * \{ \overline{\overline{\overline{\overline{\{\iota\}, \emptyset}}}} \}^p \vee [\text{NaInv} : p.\{\iota\}]}^{\mathcal{N}}$$

We easily obtain:

$$\begin{aligned} \text{True} &\Rightarrow_{\perp} \exists p. [\text{NaInv} : p] & [\text{NaInv} : p.\mathcal{E}_1 \uplus \mathcal{E}_2] &\Leftrightarrow [\text{NaInv} : p.\mathcal{E}_1] * [\text{NaInv} : p.\mathcal{E}_2] \\ \triangleright P &\Rightarrow_{\mathcal{N}} \square \text{NaInv}^{p.\mathcal{N}}(P) & \text{NaInv}^{p.\mathcal{N}}(P) \vdash \langle [\text{NaInv} : p.\mathcal{N}] \Leftrightarrow \triangleright P \rangle_{\mathcal{N}} \end{aligned}$$

from which we can derive

$$\frac{\mathcal{N} \subseteq \mathcal{E}}{\text{NaInv}^{p.\mathcal{N}}(P) \vdash \langle [\text{NaInv} : p.\mathcal{E}] \Leftrightarrow \triangleright P * [\text{NaInv} : p.\mathcal{E} \setminus \mathcal{N}] \rangle_{\mathcal{N}}}$$

³Very much like the tokens that are used to encode “normal”, atomic invariants

9.2 Boxes

The idea behind the *boxes* is to have an assertion P that is actually split into a number of pieces, each of which can be taken out and back in separately. In some sense, this is a replacement for having an “authoritative PCM of Iris assertions itself”. It is similar to the pattern involving saved propositions that was used for the barrier [5], but more complicated because there are some operations that we want to perform without a later.

Roughly, the idea is that a *box* is a container for an assertion P . A box consists of a bunch of *slices* which decompose P into a separating conjunction of the assertions Q_i governed by the individual slices. Each slice is either *full* (it right now contains Q_i), or *empty* (it does not contain anything currently). The assertion governing the box keeps track of the state of all the slices that make up the box. The crux is that opening and closing of a slice can be done even if we only have ownership of the boxes “later” (\triangleright).

The interface for boxes is as follows: The two core assertions are: $\text{BoxSlice}(\mathcal{N}, P, i)$, saying that there is a slice in namespace \mathcal{N} with name i and content P ; and $\text{Box}(\mathcal{N}, P, f)$, saying that f describes the slices of a box in namespace \mathcal{N} , such that all the slices together contain P . Here, f is of type $\mathbb{N} \xrightarrow{\text{fin}} \text{BoxState}$ mapping names to states, where $\text{BoxState} \triangleq \{\text{full}, \text{empty}\}$.

$$\begin{array}{c} \text{BOX-CREATE} \\ \text{True} \Rightarrow_{\mathcal{N}} \text{Box}(\mathcal{N}, \text{True}, \emptyset) \end{array}$$

$$\begin{array}{c} \text{SLICE-INSERT-EMPTY} \\ \triangleright^b \text{Box}(\mathcal{N}, P, f) \Rightarrow_{\mathcal{N}} \exists i \notin \text{dom}(f). \square \text{BoxSlice}(\mathcal{N}, Q, i) * \triangleright^b \text{Box}(\mathcal{N}, P * Q, f [i \leftarrow \text{empty}]) \end{array}$$

$$\begin{array}{c} \text{SLICE-DELETE-EMPTY} \\ \frac{f(i) = \text{empty}}{\text{BoxSlice}(\mathcal{N}, Q, i) \vdash \triangleright^b \text{Box}(\mathcal{N}, P, f) \Rightarrow_{\mathcal{N}} \exists P'. \triangleright^b (\triangleright(P = P' * Q) * \text{Box}(\mathcal{N}, P', f [i \leftarrow \perp]))} \end{array}$$

$$\begin{array}{c} \text{SLICE-FILL} \\ \frac{f(i) = \text{empty}}{\text{BoxSlice}(\mathcal{N}, Q, i) \vdash \triangleright^b Q * \triangleright \text{Box}(\mathcal{N}, P, f) \Rightarrow_{\mathcal{N}} \triangleright^b \text{Box}(\mathcal{N}, P, f [i \leftarrow \text{full}])} \end{array}$$

$$\begin{array}{c} \text{SLICE-EMPTY} \\ \frac{f(i) = \text{full}}{\text{BoxSlice}(\mathcal{N}, Q, i) \vdash \triangleright^b \text{Box}(\mathcal{N}, P, f) \Rightarrow_{\mathcal{N}} \triangleright Q * \triangleright^b \text{Box}(\mathcal{N}, P, f [i \leftarrow \text{empty}])} \end{array}$$

$$\begin{array}{c} \text{BOX-FILL} \\ \frac{\forall i \in \text{dom}(f). f(i) = \text{empty}}{\triangleright P * \text{Box}(\mathcal{N}, P, f) \Rightarrow_{\mathcal{N}} \text{Box}(\mathcal{N}, P, f [i \leftarrow \text{full} \mid i \in \text{dom}(f)])} \end{array}$$

$$\begin{array}{c} \text{BOX-EMPTY} \\ \frac{\forall i \in \text{dom}(f). f(i) = \text{full}}{\text{Box}(\mathcal{N}, P, f) \Rightarrow_{\mathcal{N}} \triangleright P * \text{Box}(\mathcal{N}, P, f [i \leftarrow \text{empty} \mid i \in \text{dom}(f)])} \end{array}$$

Above, $\triangleright^b P$ is syntactic sugar for $\triangleright P$ (if b is 1) or P (if b is 0). This is essentially an *optional later*, indicating that the lemmas can be applied with Box being owned now or later, and that ownership is returned the same way.

Model. The above rules are validated by the following model. We need a CMRA as follows:

$$\begin{aligned} \text{BoxState} &\triangleq \text{full} + \text{empty} \\ \text{Box} &\triangleq \text{AUTH}(\text{EX}(\text{BoxState})^?) \times \text{AG}(\blacktriangleright i\text{Prop})^? \end{aligned}$$

Now we can define the assertions:

$$\begin{aligned} \text{SliceInv}(i, P) &\triangleq \exists b. \{ \overline{\overline{\overline{\overline{\bullet b, \varepsilon}}}} \}^i * ((b = \text{full}) \Rightarrow P) \\ \text{BoxSlice}(\mathcal{N}, P, i) &\triangleq \{ \overline{\overline{\overline{\overline{\varepsilon, P}}}} \}^i * \boxed{\text{SliceInv}(i, P)}^{\mathcal{N}} \\ \text{Box}(\mathcal{N}, P, f) &\triangleq \exists Q : \mathbb{N} \rightarrow i\text{Prop}. \triangleright (P = \bigstar_{i \in \text{dom}(f)} Q(i)) * \\ &\quad \bigstar_{i \in \text{dom}(f)} \{ \overline{\overline{\overline{\overline{\circ f(i), Q(i)}}}} \}^i * \boxed{\text{SliceInv}(i, Q(i))}^{\mathcal{N}} \end{aligned}$$

Derived rules. Here are some derived rules:

$$\begin{array}{c} \text{SLICE-INSERT-FULL} \\ \triangleright Q * \triangleright^b \text{Box}(\mathcal{N}, P, f) \Rightarrow_{\mathcal{N}} \exists i \notin \text{dom}(f). \square \text{BoxSlice}(\mathcal{N}, Q, i) * \triangleright^b \text{Box}(\mathcal{N}, P * Q, f [i \leftarrow \text{full}]) \end{array}$$

$$\begin{array}{c} \text{SLICE-DELETE-FULL} \\ \text{BoxSlice}(\mathcal{N}, Q, i) \vdash \triangleright^b \text{Box}(\mathcal{N}, P, f) \Rightarrow_{\mathcal{N}} \triangleright Q * \exists P'. \triangleright^b (\triangleright (P = P' * Q) * \text{Box}(\mathcal{N}, P', f [i \leftarrow \perp])) \end{array}$$

$$\begin{array}{c} \text{SLICE-SPLIT} \\ \text{BoxSlice}(\mathcal{N}, Q_1 * Q_2, i) \vdash \triangleright^b \text{Box}(\mathcal{N}, P, f) \Rightarrow_{\mathcal{N}} \exists i_1 \notin \text{dom}(f), i_2 \notin \text{dom}(f). i_1 \neq i_2 \wedge \\ \square \text{BoxSlice}(\mathcal{N}, Q_1, i_1) * \square \text{BoxSlice}(\mathcal{N}, Q_2, i_2) * \triangleright^b \text{Box}(\mathcal{N}, P, f [i \leftarrow \perp] [i_1 \leftarrow s] [i_2 \leftarrow s]) \end{array}$$

$$\begin{array}{c} \text{SLICE-MERGE} \\ \text{BoxSlice}(\mathcal{N}, Q_1, i_1), \text{BoxSlice}(\mathcal{N}, Q_2, i_2) \vdash \triangleright^b \text{Box}(\mathcal{N}, P, f) \Rightarrow_{\mathcal{N}} \exists i \notin \text{dom}(f) \setminus \{i_1, i_2\}. \\ \square \text{BoxSlice}(\mathcal{N}, Q_1 * Q_2, i) * \triangleright^b \text{Box}(\mathcal{N}, P, f [i_1 \leftarrow \perp] [i_2 \leftarrow \perp] [i \leftarrow s]) \end{array}$$

10 Logical paradoxes

In this section we provide proofs of some logical inconsistencies that arise when slight changes are made to the Iris logic.

10.1 Saved propositions without a later

As a preparation for the proof about invariants in §10.2, we show that omitting the later modality from a variant of *saved propositions* leads to a contradiction. Saved propositions have been introduced in prior work [4, 5] to prove correctness of synchronization primitives; we will explain all that is necessary here. The counterexample assumes a higher-order logic with separating conjunction, magic wand and the modalities \Box and \Vdash satisfying the rules in §4.

Theorem 1. *If there exists a type $GName$ and an assertion $_ \Vdash _ : GName \rightarrow iProp \rightarrow iProp$ associating names $\gamma : GName$ to propositions and satisfying:*

$$\begin{aligned} \vdash \Vdash \exists \gamma : GName. \gamma \Vdash P(\gamma) & \quad (\text{SPROP-ALLOC}) \\ \gamma \Vdash P \vdash \Box(\gamma \Vdash P) & \quad (\text{SPROP-PERSIST}) \\ \gamma \Vdash P * \gamma \Vdash Q \vdash P \Leftrightarrow Q & \quad (\text{SPROP-AGREE}) \end{aligned}$$

then $\vdash \Vdash \text{False}$.

The type $GName$ should be thought of as the type of “locations” and $\gamma \Vdash P$ should be read as stating that location γ “stores” proposition P . Notice that these are immutable locations, so the maps-to assertion is persistent. The rule **SPROP-ALLOC** is then thought of as allocation, and the rule **SPROP-AGREE** states that a given location γ can only store *one* proposition, so multiple witnesses covering the same location must agree.

The conclusion of **SPROP-AGREE** usually is guarded by a \triangleright . The point of this theorem is to show that said later is *essential*, as removing it introduces inconsistency. The key to proving **Theorem 1** is the following assertion:

Definition 23. $A(\gamma) \triangleq \exists P : iProp. \Box \neg P \wedge \gamma \Vdash P$.

Intuitively, $A(\gamma)$ says that the saved proposition named γ does *not* hold, *i.e.*, we can disprove it. Using **SPROP-PERSIST**, it is immediate that $A(\gamma)$ is persistent.

Now, by applying **SPROP-ALLOC** with A , we obtain a proof of $P \triangleq \gamma \Vdash A(\gamma)$: this says that the proposition named γ is the assertion saying that it, itself, doesn’t hold. In other words, P says that the assertion named γ expresses its own negation. Unsurprisingly, that leads to a contradiction, as is shown in the following lemma:

Lemma 2. *We have $\gamma \Vdash A(\gamma) \vdash \Box \neg A(\gamma)$ and $\gamma \Vdash A(\gamma) \vdash A(\gamma)$.*

Proof.

- First we show $\gamma \Vdash A(\gamma) \vdash \Box \neg A(\gamma)$. Since $\gamma \Vdash A(\gamma)$ is persistent it suffices to show $\gamma \Vdash A(\gamma) \vdash \neg A(\gamma)$. Suppose $\gamma \Vdash A(\gamma)$ and $A(\gamma)$. Then by definition of A there is a P such that $\Box \neg P$ and $\gamma \Vdash P$. By **SPROP-AGREE** we have $P \Leftrightarrow A(\gamma)$ and so from $\neg P$ we get $\neg A(\gamma)$, which leads to a contradiction with $A(\gamma)$.

- Using the first item we can now prove $\gamma \Rightarrow A(\gamma) \vdash A(\gamma)$. We need to prove

$$\exists P : iProp. \Box \neg P \wedge \gamma \Rightarrow P.$$

We do so by picking P to be $A(\gamma)$, which leaves us to prove $\Box \neg A(\gamma) \wedge \gamma \Rightarrow A(\gamma)$. The last conjunct holds by assumption, and the first conjunct follows from the previous item of this lemma. □

With this lemma in hand, the proof of [Theorem 1](#) is simple.

Theorem 1. Using the previous lemmas we have

$$\vdash \forall \gamma. \neg(\gamma \Rightarrow A(\gamma)).$$

Together with the rule [SPROP-ALLOC](#) we thus derive $\Rightarrow \text{False}$. □

10.2 Invariants without a later

Now we come to the main paradox: if we remove the \triangleright from [INV-OPEN](#), the logic becomes inconsistent. The theorem is stated as general as possible so taht it also applies to previous, less powerful versions of Iris.

Theorem 2. *Assume a higher-order separation logic with \Box and an update modality with a binary mask $\Rightarrow_{\{0,1\}}$ (think: empty mask and full mask) satisfying strong monad rules with respect to separating conjunction and such that:*

$$\begin{array}{c} \text{WEAKEN-MASK} \\ \Rightarrow_0 P \vdash \Rightarrow_1 P \end{array}$$

Assume a type $InvName$ and an assertion $\Box[\cdot] : InvName \rightarrow iProp \rightarrow iProp$ satisfying:

$$\begin{array}{c} \text{INV-ALLOC} \\ P \vdash \Rightarrow_1 \exists \iota. \Box[P]^\iota \end{array} \quad \begin{array}{c} \text{INV-PERSIST} \\ \Box[P]^\iota \vdash \Box \Box[P]^\iota \end{array} \quad \begin{array}{c} \text{INV-OPEN-NOLATER} \\ \frac{P * Q \vdash \Rightarrow_0 (P * R)}{\Box[P]^\iota * Q \vdash \Rightarrow_1 R} \end{array}$$

Finally, assume the existence of a type $GName$ and two tokens $\Box[S]^\gamma : GName \rightarrow iProp$ and $\Box[F]^\gamma : GName \rightarrow iProp$ parameterized by $GName$ and satisfying the following properties:

$$\begin{array}{cccc} \text{START-ALLOC} & \text{START-FINISH} & \text{START-NOT-FINISHED} & \text{FINISHED-DUP} \\ \vdash \Rightarrow_0 \exists \gamma. \Box[S]^\gamma & \Box[S]^\gamma \vdash \Rightarrow_0 \Box[F]^\gamma & \Box[S]^\gamma * \Box[F]^\gamma \vdash \text{False} & \Box[F]^\gamma \vdash \Box[F]^\gamma * \Box[F]^\gamma \end{array}$$

Then $\text{True} \vdash \Rightarrow_1 \text{False}$.

The core of the proof is defining the \Rightarrow from the previous counterexample using invariants. Then, using the standard proof rules for invariants, we show that it satisfies [SPROP-ALLOC](#) and [SPROP-PERSIST](#). Furthermore, assuming the rule for opening invariants without a \triangleright , we can prove a slightly weaker version of [SPROP-AGREE](#), which is sufficient for deriving a contradiction.

We start by defining \Rightarrow satisfying (almost) the assumptions of [Lemma 4](#).

Definition 24. We define $_ \Rightarrow _ : GName \rightarrow iProp \rightarrow iProp$ as:

$$\gamma \Rightarrow P \triangleq \exists l. \boxed{[\overline{S}]^\gamma \vee [\overline{F}]^\gamma * \square P}^l.$$

Note that using **INV-PERSIST**, it is immediate that $\gamma \Rightarrow P$ is persistent.

We use the tokens $[\overline{S}]^\gamma$ and $[\overline{F}]^\gamma$ to model invariants that can be initialized “lazily”: $[\overline{S}]^\gamma$ indicates that the invariant is still not initialized, whereas the duplicable $[\overline{F}]^\gamma$ indicates it has been initialized with a resource satisfying P .

We can show variants of **SPROP-AGREE** and **SPROP-ALLOC** for the defined \Rightarrow .

Lemma 3. We have $\vdash \Rightarrow_{\top} \exists \gamma. \gamma \Rightarrow P(\gamma)$.

Proof. We have to show the allocation rule

$$\vdash \Rightarrow_{\top} \exists \gamma. \gamma \Rightarrow P.$$

From **START-ALLOC** we have a γ such that $\Rightarrow_{\perp} [\overline{S}]^\gamma$ holds and hence from **WEAKEN-MASK** we have $\Rightarrow_{\top} [\overline{S}]^\gamma$. Since we are proving a goal of the form $\Rightarrow_{\top} R$ we may assume $[\overline{S}]^\gamma$. Thus for any P we have $\Rightarrow_{\top} ([\overline{S}]^\gamma \vee [\overline{F}]^\gamma * P)$. Again since our goal is still of the form \Rightarrow_{\top} we may assume $[\overline{S}]^\gamma \vee [\overline{F}]^\gamma * \square P$. The rule **INV-ALLOC** then gives us precisely what we need. \square \square

Lemma 4. We have $\gamma \Rightarrow P * \gamma \Rightarrow Q * \square P \vdash \Rightarrow_{\top} \square Q$ and thus $\gamma \Rightarrow P * \gamma \Rightarrow Q \vdash (\Rightarrow_{\top} \square P) \Leftrightarrow (\Rightarrow_{\top} \square Q)$.

Lemma 4. • We first show

$$\gamma \Rightarrow P * \gamma \Rightarrow Q * \square P \vdash \Rightarrow_{\top} \square Q.$$

We use **INV-OPEN-NOLATER** to open the invariant in $\gamma \Rightarrow P$ and consider two cases:

1. $[\overline{S}]^\gamma$ (the invariant is “uninitialized”): In this case, we use **START-FINISH** to “initialize” the invariant and obtain $[\overline{F}]^\gamma$. Then we duplicate $[\overline{F}]^\gamma$, and use it together with $\square P$ to close the invariant.
2. $[\overline{F}]^\gamma * \square P$ (the invariant is “initialized”): In this case we duplicate $[\overline{F}]^\gamma$, and use a copy to close the invariant.

After closing the invariant, we have obtained $[\overline{F}]^\gamma$. Hence, it is sufficient to prove

$$[\overline{F}]^\gamma * \gamma \Rightarrow P * \gamma \Rightarrow Q * \square P \vdash \Rightarrow_{\top} \square Q.$$

We proceed by using **INV-OPEN-NOLATER** to open the other invariant in $\gamma \Rightarrow Q$, and we again consider two cases:

1. $[\overline{S}]^\gamma$ (the invariant is “uninitialized”): As witnessed by **START-NOT-FINISHED**, this cannot happen, so we derive a contradiction. Notice that this is a key point of the proof: because the two invariants ($\gamma \Rightarrow P$ and $\gamma \Rightarrow Q$) *share* the ghost name γ , initializing one of them is enough to show that the other one has been initialized. Essentially, this is an indirect way of saying that really, we have been opening the same invariant two times.
2. $[\overline{F}]^\gamma * \square Q$ (the invariant is “initialized”): Since $\square Q$ is duplicable we use one copy to close the invariant, and retain another to prove $\Rightarrow_{\top} \square Q$.

- By applying the above twice, we easily obtain

$$\gamma \Rightarrow P * \gamma \Rightarrow Q \vdash (\Rightarrow_{\top} \Box P) \Leftrightarrow (\Rightarrow_{\top} \Box Q)$$

□

□

Intuitively, [Lemma 4](#) shows that we can “convert” a proof from P to Q .

We are now in a position to replay the counterexample from [§10.1](#). The only difference is that because [Lemma 4](#) is slightly weaker than the rule [SPROP-AGREE](#) of [Theorem 1](#), we need to use $\Rightarrow_{\top} \text{False}$ in place of False in the definition of the predicate A : we let $A(\gamma) \triangleq \exists P : \text{iProp}. \Box(P \Rightarrow \Rightarrow_{\top} \text{False}) \wedge \gamma \Rightarrow P$ and replay the proof that we have presented above.

References

- [1] Pierre America and Jan Rutten. “Solving Reflexive Domain Equations in a Category of Complete Metric Spaces”. In: *JCSS* 39.3 (1989), pp. 343–375.
- [2] Lars Birkedal and Aleš Bizjak. *A Taste of Categorical Logic — Tutorial Notes*. Available at <http://users-cs.au.dk/birke/modures/tutorial/categorical-logic-tutorial-notes.pdf>. Oct. 2014.
- [3] Lars Birkedal, Kristian Støvring, and Jacob Thamsborg. “The category-theoretic solution of recursive metric-space equations”. In: *TCS* 411.47 (2010), pp. 4102–4122. DOI: [10.1016/j.tcs.2010.07.010](https://doi.org/10.1016/j.tcs.2010.07.010). URL: <http://dx.doi.org/10.1016/j.tcs.2010.07.010>.
- [4] Mike Dodds et al. “Verifying Custom Synchronization Constructs Using Higher-Order Separation Logic”. In: *TOPLAS* 38.2 (2016), p. 4. DOI: [10.1145/2818638](https://doi.org/10.1145/2818638). URL: <http://doi.acm.org/10.1145/2818638>.
- [5] Ralf Jung et al. “Higher-order ghost state”. In: *ICFP*. 2016, pp. 256–269.
- [6] Aaron Turon, Derek Dreyer, and Lars Birkedal. “Unifying refinement and Hoare-style reasoning in a logic for higher-order concurrency”. In: *ICFP*. 2013, pp. 377–390.