

The Iris 4.0 Reference

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

August 20, 2022

Abstract

This document formally describes 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/iris/iris>. For further information, visit the Iris project website at <https://iris-project.org>.

Contents

1	Iris from the Ground Up	4
2	Algebraic Structures	5
2.1	OFE	5
2.2	COFE	6
2.3	RA	7
2.4	Cameras	8
3	OFE and COFE Constructions	10
3.1	Trivial Pointwise Lifting	10
3.2	Next (Type-Level Later)	10
3.3	Uniform Predicates	10
4	RA and Camera Constructions	11
4.1	Product	11
4.2	Sum	11
4.3	Option	11
4.4	Finite Partial Functions	12
4.5	Agreement	12
4.6	Exclusive Camera	12
4.7	Fractions	13
4.8	Authoritative	13
4.9	STS with Tokens	14
5	Base Logic	16
5.1	Grammar	16
5.2	Types	16
5.3	Proof Rules	17
5.4	Consistency	19
6	Model and Semantics	20
7	Extensions of the Base Logic	23
7.1	Derived Rules about Base Connectives	23
7.2	Derived Rules about Modalities	23
7.3	Persistent Propositions	25
7.4	Timeless Propositions and Except-0	26
7.5	Dynamic Composable Higher-Order Resources	27
8	Language	30
8.1	Concurrent Language	31

9	Program Logic	32
9.1	Later Credits	32
9.2	World Satisfaction, Invariants, Fancy Updates	33
9.3	Weakest Precondition	35
9.4	Invariant Namespaces	40
9.5	Accessors	41
10	Derived Constructions	43
10.1	Cancellable Invariants	43
10.2	Non-atomic (“Thread-Local”) Invariants	43
10.3	Boxes	44
11	Logical Paradoxes	47
11.1	Saved Propositions without a Later	47
11.2	Invariants without a Later	48
12	HeapLang	51
12.1	HeapLang syntax and operational semantics	51
12.2	Syntactic sugar	55

1 Iris from the Ground Up

In “Iris from the Ground Up” [6], we describe Iris 3.1 in a bottom-up way. That paper is hence much more suited as an introduction to the model of Iris than this reference, which mostly contains definitions, not explanations or examples. For a list of changes in Iris since then, please consult our changelog at <https://gitlab.mpi-sws.org/iris/iris/blob/master/CHANGELOG.md>.

2 Algebraic Structures

2.1 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).

Notice that OFEs are just a different presentation of bisected 1-bounded ultrametric spaces, where the family of equivalence relations gives rise to the distance function (two elements that are equal for n steps are no more than 2^{-n} apart).

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 $f : T \overset{\text{ne}}{\rightarrow} U$) 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 **OFE** consists of OFEs as objects, and non-expansive functions as arrows.

Note that **OFE** is bicartesian closed, i.e., it has all sums, products and exponentials as well as an initial and a terminal object. In particular:

Definition 5. Given two OFEs T and U , the set of non-expansive functions $\{f : T \overset{\text{ne}}{\rightarrow} U\}$ 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 $F : \mathbf{OFE} \rightarrow \mathbf{OFE}$ 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 $(-) \xrightarrow{\text{ne}} (-)$ 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.

One very important OFE is the OFE of *step-indexed propositions*: For every step-index, such a proposition either holds or does not hold. Moreover, if a proposition holds for some n , it also has to hold for all smaller step-indices.

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

2.2 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 $(\stackrel{n}{\subseteq} T \times T)_{n \in \mathbb{N}}$ of equivalence relations, a chain $c \in \text{Chains}(T)$ is a function $c : \mathbb{N} \rightarrow T$ such that $\forall n, m. n \leq m \Rightarrow c(m) \stackrel{n}{=} c(n)$.

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

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

Definition 9. The category **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). $SProp$ as defined above is complete, i.e., it is a COFE.

Completeness is necessary to take fixed-points.

Theorem 1 (Banach's fixed-point). Given an inhabited COFE T and a contractive function $f : T \rightarrow T$, there exists a unique fixed-point $\text{fix}_T f$ such that $f(\text{fix}_T f) = \text{fix}_T f$. Moreover, this theorem also holds if f is just non-expansive, and f^k is contractive for an arbitrary k .

Theorem 2 (America and Rutten [1, 3]). Let 1 be the discrete COFE on the unit type: $1 \triangleq \Delta\{\()\}$. Given a locally contractive bifunctor $G : \mathbf{COFE}^{op} \times \mathbf{COFE} \rightarrow \mathbf{COFE}$, and provided that $G(1, 1)$ is inhabited, then there exists a unique¹ COFE T such that $G(T^{op}, T) \cong T$ (i.e., the two are isomorphic in **COFE**).

¹Uniqueness is not proven in Coq.

2.3 RA

Definition 10. A resource algebra (RA) is a tuple $(M, \bar{\vee} : M \rightarrow \text{Prop}, |-| : M \rightarrow M^?, (\cdot) : M \times M \rightarrow M)$ satisfying:

$$\begin{aligned}
& \forall a, b, c. (a \cdot b) \cdot c = a \cdot (b \cdot c) && \text{(RA-ASSOC)} \\
& \forall a, b. a \cdot b = b \cdot a && \text{(RA-COMM)} \\
& \forall a. |a| \in M \Rightarrow |a| \cdot a = a && \text{(RA-CORE-ID)} \\
& \forall a. |a| \in M \Rightarrow ||a|| = |a| && \text{(RA-CORE-IDEM)} \\
& \forall a, b. |a| \in M \wedge a \preceq b \Rightarrow |b| \in M \wedge |a| \preceq |b| && \text{(RA-CORE-MONO)} \\
& \forall a, b. \bar{\vee}(a \cdot b) \Rightarrow \bar{\vee}(a) && \text{(RA-VALID-OP)} \\
\text{where } & M^? \triangleq M \uplus \{\perp\} && a^? \cdot \perp \triangleq \perp \cdot a^? \triangleq a^? \\
& a \preceq b \triangleq \exists c \in M. b = a \cdot c && \text{(RA-INCL)}
\end{aligned}$$

Here, *Prop* is the set of (meta-level) propositions. Think of *Prop* in Coq or \mathbb{B} in classical mathematics.

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 of *valid* elements that is compatible with the composition operation (**RA-VALID-OP**). These valid elements are identified by the *validity predicate* $\bar{\vee}$.

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

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

$$\forall a_i^? \in M^?. \bar{\vee}(a \cdot a_i^?) \Rightarrow \exists b \in B. \bar{\vee}(b \cdot a_i^?)$$

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

The proposition $a \rightsquigarrow B$ says that every element $a_i^?$ compatible with a (we also call such elements *frames*), must also be compatible with some $b \in B$. Notice that $a_i^?$ 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 $b \in B$.

2.4 Cameras

Definition 12. A camera is a tuple $(M : \mathbf{OFE}, \mathcal{V} : M \xrightarrow{\text{ne}} \text{SProp}, |-| : M \xrightarrow{\text{ne}} M^?, (\cdot) : M \times M \xrightarrow{\text{ne}} M)$ satisfying:

$$\begin{aligned}
\forall a, b, c. (a \cdot b) \cdot c &= a \cdot (b \cdot c) && \text{(CAMERA-ASSOC)} \\
\forall a, b. a \cdot b &= b \cdot a && \text{(CAMERA-COMM)} \\
\forall a. |a| \in M &\Rightarrow |a| \cdot a = a && \text{(CAMERA-CORE-ID)} \\
\forall a. |a| \in M &\Rightarrow ||a|| = |a| && \text{(CAMERA-CORE-IDEM)} \\
\forall a, b. |a| \in M \wedge a \preceq b &\Rightarrow |b| \in M \wedge |a| \preceq |b| && \text{(CAMERA-CORE-MONO)} \\
\forall a, b. \mathcal{V}(a \cdot b) &\subseteq \mathcal{V}(a) && \text{(CAMERA-VALID-OP)} \\
\forall n, a, b_1, b_2. n \in \mathcal{V}(a) \wedge a \stackrel{n}{=} b_1 \cdot b_2 &\Rightarrow \\
\exists c_1, c_2. a = c_1 \cdot c_2 \wedge c_1 \stackrel{n}{=} b_1 \wedge c_2 \stackrel{n}{=} b_2 &&& \text{(CAMERA-EXTEND)}
\end{aligned}$$

where

$$\begin{aligned}
a \preceq b &\triangleq \exists c. b = a \cdot c && \text{(CAMERA-INCL)} \\
a \stackrel{n}{\preceq} b &\triangleq \exists c. b \stackrel{n}{=} a \cdot c && \text{(CAMERA-INCLN)}
\end{aligned}$$

This is a natural generalization of RAs over OFEs². 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 $\bar{\mathcal{V}}$ as the “limit” of the step-indexed approximation:

$$\bar{\mathcal{V}}(a) \triangleq \forall n. n \in \mathcal{V}(a)$$

The extension axiom (CAMERA-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:

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

²The reader may wonder why on earth we call them “cameras”. The reason, which may not be entirely convincing, is that “camera” was originally just used as a comfortable pronunciation of “CMRA”, the name used in earlier Iris papers. CMRA was originally supposed to be an acronym for “complete metric resource algebras” (or something like that), but we were never very satisfied with it and thus ended up never spelling it out. To make matters worse, the “complete” part of CMRA is now downright misleading, for whereas previously the carrier of a CMRA was required to be a COFE (complete OFE), we have relaxed that restriction and permit it to be an (incomplete) OFE. For these reasons, we have decided to stick with the name “camera”, for purposes of continuity, but to drop any pretense that it stands for something.

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 camera M is called the unit of M if it satisfies the following conditions:

1. ε is valid:
 $\forall n. n \in \mathcal{V}(\varepsilon)$
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

$$\forall n, a_i^?. n \in \mathcal{V}(a \cdot a_i^?) \Rightarrow \exists b \in B. n \in \mathcal{V}(b \cdot a_i^?)$$

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 camera M is discrete if it satisfies the following conditions:

1. M is a discrete COFE
2. \mathcal{V} ignores the step-index:
 $\forall a \in M. 0 \in \mathcal{V}(a) \Rightarrow \forall n. n \in \mathcal{V}(a)$

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

Definition 16 (Camera homomorphism). A function $f : M_1 \rightarrow M_2$ between two cameras is a camera homomorphism if it satisfies the following conditions:

1. f is non-expansive
2. f commutes with composition:
 $\forall a_1 \in M_1, a_2 \in M_1. f(a_1) \cdot f(a_2) = f(a_1 \cdot a_2)$
3. f commutes with the core:
 $\forall a \in M_1. |f(a)| = f(|a|)$
4. f preserves validity:
 $\forall n, a \in M_1. n \in \mathcal{V}(a) \Rightarrow n \in \mathcal{V}(f(a))$

Definition 17. The category **Camera** consists of cameras as objects, and camera homomorphisms as arrows.

Note that every object/arrow in **Camera** is also an object/arrow of **OFE**. The notion of a locally non-expansive (or contractive) bifunctor naturally generalizes to bifunctors between these categories.

3 OFE and COFE Constructions

3.1 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).

3.2 Next (Type-Level Later)

Given an 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 **OFE** to **OFE**.

3.3 Uniform Predicates

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

$$\begin{aligned} M &\xrightarrow{\text{mon,ne}} SPProp \triangleq \left\{ \Phi : M \xrightarrow{\text{ne}} SPProp \mid \forall n, a, b. a \stackrel{n}{\approx} b \Rightarrow \Phi(a) \stackrel{n}{\subseteq} \Phi(b) \right\} \\ UPred(M) &\triangleq M \xrightarrow{\text{mon,ne}} SPProp / \equiv \\ \Phi &\equiv \Psi \triangleq \forall m, a. m \in \mathcal{V}(a) \Rightarrow (m \in \Phi(a) \iff m \in \Psi(a)) \\ \Phi &\stackrel{n}{=} \Psi \triangleq \forall m \leq n, a. m \in \mathcal{V}(a) \Rightarrow (m \in \Phi(a) \iff m \in \Psi(a)) \end{aligned}$$

You can think of uniform predicates as monotone, step-indexed predicates over a camera that “ignore” invalid elements (as defined by the quotient).

$UPred(-)$ is a locally non-expansive functor from **Camera** to **COFE**.

It is worth noting that the above quotient admits canonical representatives. More precisely, one can show that every equivalence class contains exactly one element P_0 such that:

$$\forall n, a. (\mathcal{V}(a) \stackrel{n}{\subseteq} P_0(a)) \Rightarrow n \in P_0(a) \quad (\text{UPRED-CANONICAL})$$

Intuitively, this says that P_0 trivially holds whenever the resource is invalid. Starting from any element P , one can find this canonical representative by choosing $P_0(a) := \{n \mid n \in \mathcal{V}(a) \Rightarrow n \in P(a)\}$.

Hence, as an alternative definition of $UPred$, we could use the set of canonical representatives. This alternative definition would save us from using a quotient. However, the definitions of the various connectives would get more complicated, because we have to make sure they all verify **UPRED-CANONICAL**, which sometimes requires some adjustments. We would moreover need to prove one more property for every logical connective.

4 RA and Camera Constructions

4.1 Product

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

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

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

4.2 Sum

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

$$\begin{aligned} M_1 +_{\perp} M_2 &\triangleq \text{inl}(a_1 : M_1) \mid \text{inr}(a_2 : M_2) \mid \perp \\ \mathcal{V}(\perp) &\triangleq \emptyset \\ \mathcal{V}(\text{inl}(a)) &\triangleq \mathcal{V}_1(a) \\ \text{inl}(a_1) \cdot \text{inl}(b_1) &\triangleq \text{inl}(a_1 \cdot b_1) \\ |\text{inl}(a_1)| &\triangleq \begin{cases} \perp & \text{if } |a_1| = \perp \\ \text{inl}(|a_1|) & \text{otherwise} \end{cases} \end{aligned}$$

Above, \mathcal{V}_1 refers to the validity of M_1 . The validity, composition and core for inr are defined symmetrically. The remaining cases of the composition and core are all \perp .

Notice that we added the artificial “invalid” (or “undefined”) element \perp to this camera 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:

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

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

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

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

4.3 Option

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

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

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

4.4 Finite Partial Functions

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

We obtain the following frame-preserving updates:

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

Above, $\overline{\mathcal{V}}$ refers to the (full) validity of M .

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

4.5 Agreement

Given some OFE T , we define the camera $Ag(T)$ as follows:

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

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

$Ag(-)$ is a locally non-expansive functor from **OFE** to **Camera**.

We define a non-expansive injection ag into $Ag(T)$ as follows:

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

There are no interesting frame-preserving updates for $Ag(T)$, but we can show the following:

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

4.6 Exclusive Camera

Given an OFE T , we define a camera $Ex(T)$ such that at most one $x \in T$ can be owned:

$$\begin{aligned} Ex(T) &\triangleq \text{ex}(T) \mid \not\downarrow \\ \mathcal{V}(a) &\triangleq \left\{ n \mid a \stackrel{n}{=} \not\downarrow \right\} \end{aligned}$$

All cases of composition go to ζ .

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

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:

$$\frac{x \stackrel{n}{=} y}{\text{ex}(x) \stackrel{n}{=} \text{ex}(y)} \qquad \zeta \stackrel{n}{=} \zeta$$

$Ex(-)$ is a locally non-expansive functor from **OFE** to **Camera**.

We obtain the following frame-preserving update:

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

4.7 Fractions

We define an RA structure on the rational numbers in $(0, 1]$ as follows:

$$\begin{aligned} \text{Frac} &\triangleq \text{frac}(\mathbb{Q} \cap (0, 1]) \mid \zeta \\ \overline{\mathcal{V}}(a) &\triangleq a \neq \zeta \\ \text{frac}(q_1) \cdot \text{frac}(q_2) &\triangleq \text{frac}(q_1 + q_2) \quad \text{if } q_1 + q_2 \leq 1 \\ |\text{frac}(x)| &\triangleq \perp \\ |\zeta| &\triangleq \zeta \end{aligned}$$

All remaining cases of composition go to ζ . Frequently, we will write just x instead of $\text{frac}(x)$.

The most important property of this RA is that 1 has no frame. This is useful in combination with **SUM-SWAP**, and also when used with pairs:

$$\begin{array}{c} \text{PAIR-FRAC-CHANGE} \\ (1, a) \rightsquigarrow (1, b) \end{array}$$

4.8 Authoritative

Given a camera 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.)

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

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 $(a_1, b_1) \rightsquigarrow_{\text{L}} (a_2, b_2)$, if*

$$\forall n, a_i^?. n \in \mathcal{V}(a_1) \wedge a_1 \stackrel{n}{=} b_1 \cdot a_i^? \Rightarrow n \in \mathcal{V}(a_2) \wedge a_2 \stackrel{n}{=} b_2 \cdot a_i^?$$

In other words, the idea is that for every possible frame $a_i^?$ completing b_1 to a_1 , the same frame also completes b_2 to a_2 .

We then obtain

$$\frac{\text{AUTH-UPDATE} \quad (a_1, b_1) \rightsquigarrow_{\text{L}} (a_2, b_2)}{\bullet a_1, \circ b_1 \rightsquigarrow \bullet a_2, \circ b_2}$$

4.9 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 [9]. 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:

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

The STS RA is defined as follows

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

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

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:

$$a \cdot b = \text{frag}(\{s_1\}, \{t_1, t_2\}) \quad |a| = \text{frag}(\{s_1, s_2, s_4\}, \emptyset) \quad |b| = \text{frag}(\{s_1, s_3, s_4\}, \emptyset)$$

$$|a| \cdot |b| = \text{frag}(\{s_1, s_4\}, \emptyset) \neq |a \cdot b| = \text{frag}(\{s_1\}, \emptyset)$$

5 Base Logic

The base logic is parameterized by an arbitrary camera 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:

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

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 .

5.1 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 $\{1, 2\}$.

$$\begin{aligned} \tau ::= & T \mid 0 \mid 1 \mid \tau + \tau \mid \tau \times \tau \mid \tau \rightarrow \tau \\ t, P, \Phi ::= & x \mid F(t_1, \dots, t_n) \mid \mathbf{abort} \ t \mid () \mid (t, t) \mid \pi_i \ t \mid \lambda x : \tau. t \mid t(t) \mid \\ & \mathbf{inj}_i \ t \mid \mathbf{match} \ t \ \mathbf{with} \ \mathbf{inj}_1 \ x. t \mid \mathbf{inj}_2 \ x. t \ \mathbf{end} \mid a \mid |t| \mid t \cdot t \mid \\ & \mathbf{False} \mid \mathbf{True} \mid t =_\tau t \mid P \Rightarrow P \mid P \wedge P \mid P \vee P \mid P * P \mid P \multimap P \mid \\ & \mu x : \tau. t \mid \exists x : \tau. P \mid \forall x : \tau. P \mid \mathbf{Own} \ (t) \mid \mathcal{V}(t) \mid \square P \mid \blacksquare P \mid \triangleright P \mid \dot{\Rightarrow} P \end{aligned}$$

Well-typedness forces recursive definitions to be *guarded*: In $\mu x. t$, the variable x can only appear under the later \triangleright modality. Furthermore, the type of the definition must be *complete*. The type $iProp$ is complete, and if τ is complete, then so is $\tau' \rightarrow \tau$.

Note that the modalities $\dot{\Rightarrow}$, \square , \blacksquare 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.

5.2 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 \frac{\Gamma \vdash t : 0}{\Gamma \vdash \text{abort } t : \tau} \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 \vdash t : \tau_i \quad i \in \{1, 2\}}{\Gamma \vdash \text{inj}_i t : \tau_1 + \tau_2} \\
\\
\frac{\Gamma \vdash t : \tau_1 + \tau_2 \quad \Gamma, x : \tau_1 \vdash t_1 : \tau \quad \Gamma, y : \tau_2 \vdash t_2 : \tau}{\Gamma \vdash \text{match } t \text{ with } \text{inj}_1 x. t_1 \mid \text{inj}_2 y. t_2 \text{ end} : \tau} \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 a : M \quad \frac{\Gamma \vdash a : M}{\Gamma \vdash |a| : M} \quad \frac{\Gamma \vdash a : M \quad \Gamma \vdash b : M}{\Gamma \vdash a \cdot b : M} \\
\\
\Gamma \vdash \text{False} : iProp \quad \Gamma \vdash \text{True} : iProp \quad \frac{\Gamma \vdash t : \tau \quad \Gamma \vdash u : \tau}{\Gamma \vdash t =_{\tau} u : iProp} \quad \frac{\Gamma \vdash P : iProp \quad \Gamma \vdash Q : iProp}{\Gamma \vdash P \Rightarrow Q : iProp} \\
\\
\frac{\Gamma \vdash P : iProp \quad \Gamma \vdash Q : iProp}{\Gamma \vdash P \wedge Q : iProp} \quad \frac{\Gamma \vdash P : iProp \quad \Gamma \vdash Q : iProp}{\Gamma \vdash P \vee Q : iProp} \\
\\
\frac{\Gamma \vdash P : iProp \quad \Gamma \vdash Q : iProp}{\Gamma \vdash P * Q : iProp} \quad \frac{\Gamma \vdash P : iProp \quad \Gamma \vdash Q : iProp}{\Gamma \vdash P \multimap Q : iProp} \\
\\
\frac{\Gamma, x : \tau \vdash t : \tau \quad x \text{ is guarded in } t \quad \tau \text{ is complete and inhabited}}{\Gamma \vdash \mu x : \tau. t : \tau} \quad \frac{\Gamma, x : \tau \vdash P : iProp}{\Gamma \vdash \exists x : \tau. P : iProp} \\
\\
\frac{\Gamma, x : \tau \vdash P : iProp}{\Gamma \vdash \forall x : \tau. P : iProp} \quad \frac{\Gamma \vdash a : M}{\Gamma \vdash \text{Own}(a) : iProp} \quad \frac{\Gamma \vdash a : \tau \quad \tau \text{ is a camera}}{\Gamma \vdash \mathcal{V}(a) : iProp} \quad \frac{\Gamma \vdash P : iProp}{\Gamma \vdash \square P : iProp} \\
\\
\frac{\Gamma \vdash P : iProp}{\Gamma \vdash \blacksquare P : iProp} \quad \frac{\Gamma \vdash P : iProp}{\Gamma \vdash \triangleright P : iProp} \quad \frac{\Gamma \vdash P : iProp}{\Gamma \vdash \dot{\Rightarrow} P : iProp}
\end{array}$$

5.3 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.

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

Furthermore, we have the usual η and β laws for projections, **abort**, sum elimination, 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 plainness modality.

$$\begin{array}{c}
\blacksquare\text{-MONO} \\
\frac{P \vdash Q}{\blacksquare P \vdash \blacksquare Q} \\
\\
\blacksquare\text{-E} \\
\blacksquare P \vdash \square P \\
\\
(\blacksquare P \Rightarrow \blacksquare Q) \vdash \blacksquare(\blacksquare P \Rightarrow Q) \\
\blacksquare((P \Rightarrow Q) \wedge (Q \Rightarrow P)) \vdash P =_{iProp} Q \\
\\
\blacksquare P \vdash \blacksquare\blacksquare P \\
\forall x. \blacksquare P \vdash \blacksquare\forall x. P \\
\blacksquare\exists x. P \vdash \exists x. \blacksquare P
\end{array}$$

Laws for the persistence modality.

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

Laws for the later modality.

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

Laws for resources and validity.

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

Laws for the basic update modality.

$$\begin{array}{l}
\text{UPD-MONO} \\
\frac{P \vdash Q}{\dot{\Rightarrow} P \vdash \dot{\Rightarrow} Q} \\
\text{UPD-I} \\
P \vdash \dot{\Rightarrow} P \\
\text{UPD-TRANS} \\
\dot{\Rightarrow} \dot{\Rightarrow} P \vdash \dot{\Rightarrow} P \\
\text{UPD-FRAME} \\
Q * \dot{\Rightarrow} P \vdash \dot{\Rightarrow} (Q * P) \\
\text{UPD-UPDATE} \\
\frac{a \rightsquigarrow B}{\text{Own}(a) \vdash \dot{\Rightarrow} \exists b \in B. \text{Own}(b)} \\
\text{UPD-PLAINLY} \\
\dot{\Rightarrow} \blacksquare P \vdash P
\end{array}$$

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

5.4 Consistency

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

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

where $(\triangleright)^n$ is short for \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 and \blacksquare , this follows from the elimination rules. For updates, we use the fact that $\dot{\Rightarrow} \text{False} \vdash \dot{\Rightarrow} \blacksquare \text{False} \vdash \text{False}$. However, there is no elimination rule for \triangleright , so we declare that it is impossible to derive a contradiction below any number of laterers.

6 Model and Semantics

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

Semantic domains. The semantic domains are interpreted as follows:

$$\begin{array}{ll}
\llbracket iProp \rrbracket \triangleq UPred(M) & \llbracket \tau + \tau' \rrbracket \triangleq \llbracket \tau \rrbracket + \llbracket \tau' \rrbracket \\
\llbracket \mathbf{M} \rrbracket \triangleq M & \llbracket \tau \times \tau' \rrbracket \triangleq \llbracket \tau \rrbracket \times \llbracket \tau' \rrbracket \\
\llbracket 0 \rrbracket \triangleq \Delta \emptyset & \llbracket \tau \rightarrow \tau' \rrbracket \triangleq \llbracket \tau \rrbracket \xrightarrow{ne} \llbracket \tau' \rrbracket \\
\llbracket 1 \rrbracket \triangleq \Delta \{()\} &
\end{array}$$

For the remaining base types τ defined by the signature \mathcal{S} , we pick an object X_τ in **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 propositions.

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

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

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

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

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_\rho &\triangleq \rho(x) \\
\llbracket \Gamma \vdash F(t_1, \dots, t_n) : \tau_{n+1} \rrbracket_\rho &\triangleq \llbracket F \rrbracket(\llbracket \Gamma \vdash t_1 : \tau_1 \rrbracket_\rho, \dots, \llbracket \Gamma \vdash t_n : \tau_n \rrbracket_\rho) \\
\llbracket \Gamma \vdash \lambda x : \tau. t : \tau \rightarrow \tau' \rrbracket_\rho &\triangleq \lambda u : \llbracket \tau \rrbracket. \llbracket \Gamma, x : \tau \vdash t : \tau \rrbracket_{\rho[x \leftarrow u]} \\
\llbracket \Gamma \vdash t(u) : \tau' \rrbracket_\rho &\triangleq \llbracket \Gamma \vdash t : \tau \rightarrow \tau' \rrbracket_\rho(\llbracket \Gamma \vdash u : \tau \rrbracket_\rho) \\
\llbracket \Gamma \vdash \mu x : \tau. t : \tau \rrbracket_\rho &\triangleq \text{fix}_{\llbracket \tau \rrbracket}(\lambda u : \llbracket \tau \rrbracket. \llbracket \Gamma, x : \tau \vdash t : \tau \rrbracket_{\rho[x \leftarrow u]}) \\
\llbracket \Gamma \vdash \text{abort } t : \tau \rrbracket_\rho &\triangleq \text{abort}_{\llbracket \tau \rrbracket}(\llbracket \Gamma \vdash t : 0 \rrbracket_\rho) \\
\llbracket \Gamma \vdash () : 1 \rrbracket_\rho &\triangleq () \\
\llbracket \Gamma \vdash (t_1, t_2) : \tau_1 \times \tau_2 \rrbracket_\rho &\triangleq (\llbracket \Gamma \vdash t_1 : \tau_1 \rrbracket_\rho, \llbracket \Gamma \vdash t_2 : \tau_2 \rrbracket_\rho) \\
\llbracket \Gamma \vdash \pi_i t : \tau_i \rrbracket_\rho &\triangleq \pi_i(\llbracket \Gamma \vdash t : \tau_1 \times \tau_2 \rrbracket_\rho) \\
\llbracket \Gamma \vdash \text{inj}_i t : \tau_1 + \tau_2 \rrbracket_\rho &\triangleq \text{inj}_i(\llbracket \Gamma \vdash t : \tau_i \rrbracket_\rho) \\
\llbracket \Gamma \vdash \text{match } t \text{ with } \text{inj}_1 x_1. t_1 \mid \text{inj}_2 x_2. t_2 \text{ end} : \tau \rrbracket_\rho &\triangleq \llbracket \Gamma, x_i : \tau_i \vdash t_i : \tau \rrbracket_{\rho[x_i \leftarrow u]} \\
&\quad \text{where } \llbracket \Gamma \vdash t : \tau_1 + \tau_2 \rrbracket_\rho = \text{inj}_i(u)
\end{aligned}$$

$$\begin{aligned}
\llbracket a : \mathbf{M} \rrbracket_\rho &\triangleq a \\
\llbracket \Gamma \vdash |t| : \mathbf{M} \rrbracket_\rho &\triangleq \llbracket \Gamma \vdash t : \mathbf{M} \rrbracket_\rho \\
\llbracket \Gamma \vdash t \cdot u : \mathbf{M} \rrbracket_\rho &\triangleq \llbracket \Gamma \vdash t : \mathbf{M} \rrbracket_\rho \cdot \llbracket \Gamma \vdash u : \mathbf{M} \rrbracket_\rho
\end{aligned}$$

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

and $\rho(x) \in \llbracket \Gamma(x) \rrbracket$. Above, fix is Banach's fixed-point (see [Theorem 1](#)), and $abort_T$ is the unique function $\emptyset \rightarrow T$.

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

Interpretation of entailment

$$\boxed{\llbracket \Gamma \mid \Theta \vdash P \rrbracket : Prop}$$

$$\begin{aligned} \llbracket \Gamma \mid P \vdash Q \rrbracket &\triangleq \forall n \in \mathbb{N}. \forall r \in M. \forall \rho \in \llbracket \Gamma \rrbracket, \\ &n \in \mathcal{V}(r) \wedge n \in \llbracket \Gamma \vdash P : iProp \rrbracket_\rho(r) \Rightarrow n \in \llbracket \Gamma \vdash Q : iProp \rrbracket_\rho(r) \end{aligned}$$

The following soundness theorem connects syntactic and semantic entailment. It is proven by showing that all the syntactic proof rules of [§5](#) can be validated in the model.

$$\Gamma \mid P \vdash Q \Rightarrow \llbracket \Gamma \mid P \vdash Q \rrbracket$$

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

7 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.

7.1 Derived Rules about Base Connectives

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

$$P \Rightarrow Q \vdash P \multimap Q \qquad P * \exists x. Q \dashv\vdash \exists x. P * Q \qquad P * \forall x. Q \vdash \forall x. P * Q$$

Verifying that existential quantifiers commute with separating conjunction requires an intermediate step using a magic wand: From $P * \exists x. Q \vdash \exists x. P * Q$ we can deduce $\exists x. Q \vdash P \multimap \exists x. P * Q$ and then proceed via \exists -elimination.

7.2 Derived Rules about Modalities

Iris comes with 4 built-in modalities (\Box , \blacksquare , $\dot{\Rightarrow}$ and \triangleright) and, as we will see, plenty of derived modalities. However, almost all of them fall into one of two categories (except for \triangleright , as we will see): they are either *always-style* modalities (“something holds in all/many (future) worlds”) or *eventually-style* modalities (“something holds in a possible (future) world”).

Eventually-style modalities are characterized by being easy to “add”/introduce, but hard to “remove”/eliminate. Consider, for example, the basic update modality $\dot{\Rightarrow}$: we have $P \vdash \dot{\Rightarrow} P$ (**UPD-I**), but the inverse direction does not hold. Instead, from **UPD-MONO** and **UPD-TRANS**, we can derive the following elimination principle:

$$\frac{\text{UPD-E} \quad P \vdash \dot{\Rightarrow} Q}{\dot{\Rightarrow} P \vdash \dot{\Rightarrow} Q}$$

In other words, we can remove an $\dot{\Rightarrow}$ in front of an assumption *if* the goal is itself wrapped in $\dot{\Rightarrow}$. Another way to view this rule is to think of it as a *bind rule*. Indeed, together with **UPD-I**, this rule shows that $\dot{\Rightarrow}$ forms a monad.

Always-style modalities, on the other hand, are easy to “remove”/eliminate, but hard to “add”/introduce. The most widely used example of that in Iris is the persistence modality \Box : we have $\Box P \vdash P$ (**\Box -E**), but the inverse direction does not hold. Instead, from **\Box -MONO** and $\Box P \vdash \Box \Box P$, we can derive the following introduction principle:

$$\frac{\Box\text{-I} \quad \Box P \vdash Q}{\Box P \vdash \Box Q}$$

In other words, we can remove an \Box from the goal *if* all our assumptions are wrapped in \Box . This matches the algebraic structure of a comonad.

In particular, both eventually-style and always-style modalities are *idempotent*: we have $\dot{\Rightarrow} \dot{\Rightarrow} P \dashv\vdash \dot{\Rightarrow} P$ and $\Box \Box P \dashv\vdash \Box P$.

Beyond this, all modalities come with plenty of rules that show how they commute around other connectives and modalities. And, of course, they come with a few “defining rules” that give the modalities their individual meaning, *i.e.*, for the update modality, that would be **UPD-UPDATE**.

In the following, we briefly discuss each of the modalities.

Update modality. As already mentioned, the update modality is an eventually-style modality:

$$\frac{\text{UPD-E}}{P \vdash \dot{\Rightarrow} Q} \quad \frac{\text{UPD-IDEMP}}{\dot{\Rightarrow} \dot{\Rightarrow} P \dashv\vdash \dot{\Rightarrow} P}$$

Beyond this (and the obvious variant of **UPD-FRAME** that exploits commutativity of separating conjunction), there are no outstandingly interesting derived rules.

Persistence modality. As already mentioned, the persistence modality is an always-style modality:

$$\frac{\text{□-I}}{\square P \vdash Q} \quad \frac{\text{□-IDEMP}}{\square \square P \dashv\vdash \square P}$$

Some further interesting derived rules include:

$$\begin{aligned} \square(P \wedge Q) \dashv\vdash \square P \wedge \square Q & \quad \square(P \vee Q) \dashv\vdash \square P \vee \square Q & \quad \square \text{True} \dashv\vdash \text{True} & \quad \square \text{False} \dashv\vdash \text{False} \\ \square(P * Q) \dashv\vdash \square P * \square Q & \quad \square P * Q \dashv\vdash \square P \wedge Q & \quad \square(P \multimap Q) \dashv\vdash \square(P \Rightarrow Q) \\ \square(P \Rightarrow Q) \vdash \square P \Rightarrow \square Q & \quad \square(P \multimap Q) \vdash \square P \multimap \square Q \end{aligned}$$

In particular, the persistence modality commutes around conjunction, disjunction, separating conjunction as well as universal and existential quantification. Commuting around conjunction can be derived from the primitive rule that says it commutes around universal quantification (as conjunction is equivalent to a universal quantification of a Boolean), and similar for disjunction. $\text{True} \dashv\vdash \square \text{True}$ (which is basically persistence “commuting around” the nullary operator **True**) can be derived via \square commuting with universal quantification ranging over the empty type. A similar rule holds for **False**.

Moreover, if (at least) one conjunct is below the persistence modality, then conjunction and separating conjunction coincide.

Plainness modality. The plainness modality is very similar to the persistence modality (in fact, we have $\blacksquare P \vdash \square P$, but the inverse does not hold). It is always-style:

$$\frac{\text{■-I}}{\blacksquare P \vdash Q} \quad \blacksquare \blacksquare P \dashv\vdash \blacksquare P$$

It also commutes around separating conjunction, conjunction, disjunction, universal and existential quantification (and **True** and **False**).

The key difference to the persistence modality \Box is that \blacksquare provides a *propositional extensionality* principle:

$$\blacksquare((P \Rightarrow Q) \wedge (Q \Rightarrow P)) \vdash P =_{iProp} Q$$

In contrast, \Box permits using some forms of ghost state ($\text{Own}(a) \vdash \Box \text{Own}(|a|)$).

Having both of these principles for the same modality would lead to a contradiction: imagine we have an RA with elements a, b such that $|a|$ is incompatible with b (i.e., $\neg \bar{\mathcal{V}}(|a| \cdot b)$). Then we can prove:

$$\text{Own}(|a|) \vdash \Box \text{Own}(|a|) \vdash \Box((\text{False} \Rightarrow \text{Own}(b)) \wedge (\text{Own}(b) \Rightarrow \text{False}))$$

The first implication is trivial, the second implication follows because $\Box \text{Own}(|a|) \wedge \text{Own}(b) \vdash \text{Own}(|a|) * \text{Own}(b) \vdash \mathcal{V}(|a| \cdot b)$.

But now, if we had propositional extensionality for \Box the way we do for \blacksquare , we could deduce $\text{False} =_{iProp} \text{Own}(b)$, and that is clearly wrong. This issue arises because \Box , as we have seen, still lets us use some resources from the context, while propositional equality has to hold completely disregarding current resources.

Later modality. The later modality is the “odd one out” in the sense that it is neither eventually-style nor always-style, because it is not idempotent:³ with \triangleright , the number of times the modality is applied matters, and we can get rid of *exactly one* layer of \triangleright in the assumptions only by doing the same in the conclusion (\triangleright -MONO).

Some derived rules:

$$\begin{array}{l} \text{L\"OB} \\ \triangleright(P \Rightarrow P) \vdash P \qquad \triangleright(P \Rightarrow Q) \vdash \triangleright P \Rightarrow \triangleright Q \qquad \triangleright(P * Q) \vdash \triangleright P * \triangleright Q \\ \\ \triangleright(P \wedge Q) \dashv\vdash \triangleright P \wedge \triangleright Q \qquad \triangleright(P \vee Q) \dashv\vdash \triangleright P \vee \triangleright Q \qquad \frac{\tau \text{ is inhabited}}{\triangleright(\exists x : \tau. P) \dashv\vdash \exists x : \tau. \triangleright P} \\ \\ \triangleright \text{True} \dashv\vdash \text{True} \qquad \triangleright(P * Q) \dashv\vdash \triangleright P * \triangleright Q \qquad \triangleright \Box P \dashv\vdash \Box \triangleright P \qquad \triangleright \blacksquare P \dashv\vdash \blacksquare \triangleright P \end{array}$$

Noteworthy here is the fact that L\"ob induction (**L\"OB**) can be derived from \triangleright -introduction and the fact that we can take fixed-points of functions where the recursive occurrences are below \triangleright [8].⁴ Also, \triangleright commutes over separating conjunction, conjunction, disjunction, universal quantification and *non-empty* existential quantification, as well as both the persistence and the plainness modality.

7.3 Persistent Propositions

We call a proposition P *persistent* if $P \vdash \Box P$. These are propositions that “do not own anything”, so we can (and will) treat them like “normal” intuitionistic propositions.

Of course, $\Box P$ is persistent for any P . Furthermore, by the proof rules given in §5.3, **True**, **False**, $t = t'$ as well as $\bar{[a]}$ and $\mathcal{V}(a)$ are persistent. Persistence is preserved by conjunction, disjunction, separating conjunction as well as universal and existential quantification and \triangleright .

³This means \triangleright is neither a monad nor a comonad—it does form an applicative functor, though.

⁴Also see https://en.wikipedia.org/wiki/L%C3%B6b%27s_theorem.

7.4 Timeless Propositions 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:

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

Except-0 satisfies the usual laws of a “monadic” modality (similar to, *e.g.*, the update modalities):

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

In particular, from **EX0-MONO** and **EX0-IDEM** we can derive a “bind”-like elimination rule:

$$\frac{\text{EX0-ELIM}}{P \vdash \diamond Q}{\diamond P \vdash \diamond Q}$$

This modality is useful because there is a class of propositions which we call *timeless* propositions, 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 propositions (using **EX0-ELIM**):

$$\frac{\text{EX0-TIMELESS-STRIP}}{\text{timeless}(P) \quad P \vdash \diamond Q}{\triangleright P \vdash \diamond Q}$$

In fact, it turns out that we can strip away later from timeless propositions even when working under the later modality:

$$\frac{\text{LATER-TIMELESS-STRIP}}{\text{timeless}(P) \quad P \vdash \triangleright Q}{\triangleright P \vdash \triangleright Q}$$

This follows from $\triangleright P \vdash \triangleright \text{False} \vee P$, and then by straightforward disjunction elimination.

The following rules identify the class of timeless propositions:

$\frac{\Gamma \vdash \text{timeless}(P) \quad \Gamma \vdash \text{timeless}(Q)}{\Gamma \vdash \text{timeless}(P \wedge Q)}$	$\frac{\Gamma \vdash \text{timeless}(P) \quad \Gamma \vdash \text{timeless}(Q)}{\Gamma \vdash \text{timeless}(P \vee Q)}$		
$\frac{\Gamma \vdash \text{timeless}(P) \quad \Gamma \vdash \text{timeless}(Q)}{\Gamma \vdash \text{timeless}(P * Q)}$	$\frac{\Gamma \vdash \text{timeless}(P)}{\Gamma \vdash \text{timeless}(\Box P)}$	$\frac{\Gamma \vdash \text{timeless}(Q)}{\Gamma \vdash \text{timeless}(P \Rightarrow Q)}$	
$\frac{\Gamma \vdash \text{timeless}(Q)}{\Gamma \vdash \text{timeless}(P \multimap Q)}$	$\frac{\Gamma, x : \tau \vdash \text{timeless}(P)}{\Gamma \vdash \text{timeless}(\forall x : \tau. P)}$	$\frac{\Gamma, x : \tau \vdash \text{timeless}(P)}{\Gamma \vdash \text{timeless}(\exists x : \tau. P)}$	$\text{timeless}(\text{True})$

$$\begin{array}{c}
\text{timeless}(\text{False}) \qquad \frac{t \text{ or } t' \text{ is a discrete OFE element}}{\text{timeless}(t =_{\tau} t')} \qquad \frac{a \text{ is a discrete OFE element}}{\text{timeless}(\text{Own}(a))} \\
\\
\frac{a \text{ is an element of a discrete camera}}{\text{timeless}(\mathcal{V}(a))}
\end{array}$$

7.5 Dynamic Composable Higher-Order Resources

The base logic described in §5 works over an arbitrary camera M defining the structure of the resources. It turns out that we can generalize this further and permit picking cameras “ $\Sigma(iProp)$ ” that depend on the structure of propositions themselves. Of course, $iProp$ is just the syntactic type of propositions; 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 camera M_1 , and another proof carried out with a *different* camera 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 camera 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 logic with dynamic higher-order ghost state, the user picks a family of locally contractive bifunctors $(\Sigma_i : \mathbf{COFE}^{\text{op}} \times \mathbf{COFE} \rightarrow \mathbf{Camera})_{i \in \mathcal{I}}$. (This is in contrast to the base logic, where the user picks a single, fixed camera 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 camera by lifting the individual cameras pointwise, and it has a unit (using the empty finite partial function). Furthermore, since the Σ_i are locally contractive, so is $ResF$.

Now we can write down the recursive domain equation:

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

Here, $iPreProp$ is a COFE defined as the fixed-point of a locally contractive bifunctor, which exists

and is unique up to isomorphism by [Theorem 2](#), so we obtain some object $iPreProp$ such that:

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

Now we can instantiate the base logic described in [§5](#) with Res as the chosen camera:

$$\llbracket iProp \rrbracket \triangleq UPred(Res)$$

We obtain that $\llbracket iProp \rrbracket = iProp$. Effectively, we just defined a way to instantiate the base logic with Res as the camera of resources, while providing a way for Res to depend on $iPreProp$, which is isomorphic to $\llbracket iProp \rrbracket$.

We thus obtain all the rules of [§5](#), and furthermore, we can use the maps ξ and ξ^{-1} *in the logic* to convert between logical propositions $\llbracket iProp \rrbracket$ and the domain $iPreProp$ which is used in the construction of Res – so from elements of $iPreProp$, we can construct elements of $\llbracket \mathbf{M} \rrbracket$, 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 cameras 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 camera 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:

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

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:

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

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

Below, we will always work within (an instance of) the logic as described here. Whenever a camera 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 Language

A language Λ consists of a set $Expr$ of *expressions* (metavariable e), a set Val of *values* (metavariable v), a set Obs of *observations*⁵ (or “observable events”) and a 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*

$$(-, - \xrightarrow{\vec{\kappa}}_{\text{t}} -, -, -) \subseteq (Expr \times State) \times List(Obs) \times (Expr \times State \times List(Expr))$$

A reduction $e_1, \sigma_1 \xrightarrow{\vec{\kappa}}_{\text{t}} e_2, \sigma_2, \vec{e}$ indicates that, when e_1 in state σ_1 reduces to e_2 with new state σ_2 , the new threads in the list \vec{e} is forked off and the observations $\vec{\kappa}$ are made. We will write $e_1, \sigma_1 \rightarrow_{\text{t}} e_2, \sigma_2$ for $e_1, \sigma_1 \xrightarrow{()}_{\text{t}} e_2, \sigma_2, ()$, *i.e.*, when no threads are forked off and no observations are made.

- All values are stuck:

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

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

$$\exists \vec{\kappa}, e_2, \sigma_2, \vec{e}. e, \sigma \xrightarrow{\vec{\kappa}}_{\text{t}} e_2, \sigma_2, \vec{e}$$

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

$$\text{atomic}(e) \triangleq \forall \sigma_1, \vec{\kappa}, e_2, \sigma_2, \vec{e}. e, \sigma_1 \xrightarrow{\vec{\kappa}}_{\text{t}} e_2, \sigma_2, \vec{e} \Rightarrow \neg \text{red}(e_2, \sigma_2)$$

It is strongly atomic if it reduces in one step to a value:

$$\text{strongly_atomic}(e) \triangleq \forall \sigma_1, \vec{\kappa}, e_2, \sigma_2, \vec{e}. e, \sigma_1 \xrightarrow{\vec{\kappa}}_{\text{t}} e_2, \sigma_2, \vec{e} \Rightarrow \text{expr_to_val}(e_2) \neq \perp$$

We need two notions of atomicity to accommodate both kinds of weakest preconditions that we will define later: If the weakest precondition ensures that the program cannot get stuck, weak atomicity is sufficient. Otherwise, we need strong atomicity.

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:

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

⁵See https://gitlab.mpi-sws.org/iris/iris/merge_requests/173 for how observations are useful to encode prophecy variables.

2. One can perform reductions below K :

$$\forall e_1, \sigma_1, \vec{\kappa}, e_2, \sigma_2, \vec{e}. e_1, \sigma_1 \xrightarrow{\vec{\kappa}}_t e_2, \sigma_2, \vec{e} \Rightarrow K(e_1), \sigma_1 \xrightarrow{\vec{\kappa}}_t K(e_2), \sigma_2, \vec{e}$$

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

$$\begin{aligned} \forall e'_1, \sigma_1, \vec{\kappa}, e_2, \sigma_2, \vec{e}. \text{expr_to_val}(e'_1) = \perp \wedge K(e'_1), \sigma_1 \xrightarrow{\vec{\kappa}}_t e_2, \sigma_2, \vec{e} \Rightarrow \\ \exists e'_2. e_2 = K(e'_2) \wedge e'_1, \sigma_1 \xrightarrow{\vec{\kappa}}_t e'_2, \sigma_2, \vec{e} \end{aligned}$$

8.1 Concurrent Language

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

Machine syntax

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

Machine reduction

$$\boxed{T; \sigma \xrightarrow{\vec{\kappa}}_{\text{tp}} T'; \sigma'}$$

$$\frac{e_1, \sigma_1 \xrightarrow{\vec{\kappa}}_t e_2, \sigma_2, \vec{e}}{T \text{ ++ } [e_1] \text{ ++ } T'; \sigma_1 \xrightarrow{\vec{\kappa}}_{\text{tp}} T \text{ ++ } [e_2] \text{ ++ } T' \text{ ++ } \vec{e}; \sigma_2}$$

We use $\xrightarrow{\vec{\kappa}}_{\text{tp}}^*$ for the reflexive transitive closure of $\xrightarrow{\vec{\kappa}}_{\text{tp}}$, as usual concatenating the lists of observations of the individual steps.

9 Program Logic

This section describes how to build a program logic for an arbitrary language (*c.f.* §8) on top of the base logic. So in the following, we assume that some language Λ was fixed. Furthermore, we work in the logic with higher-order ghost state as described in §7.5.

9.1 Later Credits

Introducing a later modality is easy (see \triangleright -I), but eliminating them can be tricky. Lateres often appear in the middle of our proofs after unfolding a circular construction (*e.g.*, by opening an invariant, see Section 9.2). In these cases, we get to assume $\triangleright P$, but we really want P to continue in the proof. Iris offers us four options to do so. We have seen two of them already: timeless propositions (see Section 7.4) and the commuting rules for later. Together, they can be used to turn $\triangleright P$ into P , or to delay when we have to deal with the later modality by commuting it inward (*e.g.*, over an existential quantifier and a separating conjunction). Another option, which we will encounter in Section 9.3, is taking program steps: every program step allows us to eliminate (at least) one later (see WP-LIFT-STEP). We now introduce the fourth option: *later credits*. Later credits turn the right to eliminate a later into an ownable separation-logic resource $\pounds n$, where n is the number of lateres that we can eliminate.

Resources We assume that the camera

$$\text{LaterCredits} \triangleq \text{Auth}(\mathbb{N}^+)$$

is available (*i.e.*, part of Σ), where \mathbb{N}^+ is the RA derived from the monoid \mathbb{N} with operation $+$, and that an instance of it has been created and made globally available at the beginning of verification under the name γ_{Credits} . We define the following notations for the fragments and authoritative part:

$$\pounds n \triangleq \llbracket \text{on} \rrbracket^{\gamma_{\text{Credits}}} \qquad \pounds_{\bullet} n \triangleq \llbracket \bullet n \rrbracket^{\gamma_{\text{Credits}}}$$

This definition satisfies the following laws:

<p>CREDIT-SPLIT</p> $\pounds(n + m) \Leftrightarrow \pounds n * \pounds m$	<p>CREDIT-TIMELESS</p> $\text{timeless}(\pounds n)$	
<p>CREDIT-SUPPLYBOUND</p> $\pounds_{\bullet} m * \pounds n \vdash m \geq n$	<p>CREDIT-SUPPLYDECR</p> $\pounds_{\bullet}(n + m) * \pounds n \vdash \pounds_{\bullet} m$	<p>CREDIT-SUPPLYEXCL</p> $\pounds_{\bullet} m_1 * \pounds_{\bullet} m_2 \vdash \text{False}$

Later Elimination Update To eliminate lateres by *spending* later credits $\pounds n$, we define a *later-elimination update* $\pounds^{\pounds} P$ on top of the basic update modality. It satisfies all the properties of basic updates, except for UPD-PLAINLY, but with the additional rule

$$\begin{array}{c} \text{CREDIT-UPD-USE} \\ \triangleright P * \pounds 1 \vdash \pounds^{\pounds} P \end{array}$$

CREDIT-UPD-USE allows to *spend* one credit in exchange for stripping one later off of P .

The later-elimination update is defined by guarded recursion:

$$\pounds^{\pounds} \triangleq \mu \text{upd}. \lambda P. \forall n. \pounds_{\bullet} n -* \pounds^{\pounds} (\pounds_{\bullet} n * P) \vee (\exists m < n. \pounds_{\bullet} m * \triangleright \text{upd}(P))$$

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

$$\text{Inv}^\iota(P) \triangleq \llbracket \circ[\iota \leftarrow \text{ag}(\text{next}(\xi(P)))] \rrbracket^{\gamma_{\text{Inv}}}$$

Fancy Updates and View Shifts. Next, we define *fancy updates*, which are essentially the same as the basic updates of the base logic (§5) or later-elimination updates (§9.1), except that they also have access to world satisfaction and can enable and disable invariants.

Depending on how the logic is parameterized with the `UseLaterCredits`, fancy updates are defined on top the basic update or the later-elimination update. This influences which rules fancy updates satisfy: either fancy updates can be used to eliminate later credits, or they satisfy certain interaction laws with the plainly modality.

$$\varepsilon_1 \Rightarrow_{\varepsilon_2} P \triangleq \begin{cases} W * [\overline{\mathcal{E}_1}]^{\gamma_{E_n}} * \vdash \overset{\mathcal{L}}{\Rightarrow} \diamond (W * [\overline{\mathcal{E}_2}]^{\gamma_{E_n}} * P) & \text{if UseLaterCredits} = \text{true} \\ W * [\overline{\mathcal{E}_1}]^{\gamma_{E_n}} * \vdash \overset{\mathcal{L}}{\Rightarrow} \diamond (W * [\overline{\mathcal{E}_2}]^{\gamma_{E_n}} * P) & \text{if UseLaterCredits} = \text{false} \end{cases}$$

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. Masks are sets of natural numbers, *i.e.*, they are subsets of \mathbb{N} .⁶ We use \top as symbol for the largest possible mask, \mathbb{N} , and \perp for the smallest possible mask \emptyset . We will write $\Rightarrow_{\mathcal{E}} P$ for $\varepsilon \Rightarrow_{\mathcal{E}} P$. Fancy updates satisfy the following basic proof rules:

$$\begin{array}{c} \text{FUP-MONO} \\ \frac{P \vdash Q}{\varepsilon_1 \Rightarrow_{\varepsilon_2} P \vdash \varepsilon_1 \Rightarrow_{\varepsilon_2} Q} \end{array} \quad \begin{array}{c} \text{FUP-INTRO-MASK} \\ \frac{\mathcal{E}_2 \subseteq \mathcal{E}_1}{P \vdash \varepsilon_1 \Rightarrow_{\varepsilon_2} \varepsilon_2 \Rightarrow_{\varepsilon_1} P} \end{array} \quad \begin{array}{c} \text{FUP-TRANS} \\ \varepsilon_1 \Rightarrow_{\varepsilon_2} \varepsilon_2 \Rightarrow_{\varepsilon_3} P \vdash \varepsilon_1 \Rightarrow_{\varepsilon_3} P \end{array} \quad \begin{array}{c} \text{FUP-UPD} \\ \vdash P \vdash \Rightarrow_{\mathcal{E}} P \end{array}$$

$$\begin{array}{c} \text{FUP-FRAME} \\ Q * \varepsilon_1 \Rightarrow_{\varepsilon_2} P \vdash \varepsilon_1 \uplus \varepsilon_f \Rightarrow_{\varepsilon_2 \uplus \varepsilon_f} Q * P \end{array} \quad \begin{array}{c} \text{FUP-UPDATE} \\ \frac{a \rightsquigarrow B}{\text{Own}(a) \vdash \Rightarrow_{\mathcal{E}} \exists b \in B. \text{Own}(b)} \end{array} \quad \begin{array}{c} \text{FUP-TIMELESS} \\ \frac{\text{timeless}(P)}{\triangleright P \vdash \Rightarrow_{\mathcal{E}} P} \end{array}$$

$$\begin{array}{c} \text{FUP-CREDIT-USE} \\ \frac{\text{UseLaterCredits} = \text{true}}{\mathcal{L} 1 * \triangleright \varepsilon_1 \Rightarrow_{\varepsilon_2} P \vdash \varepsilon_1 \Rightarrow_{\varepsilon_2} P} \end{array}$$

(There are no rules related to invariants here. Those rules will be discussed later, in §9.4.)

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

$$\begin{aligned} P \varepsilon_1 \Rightarrow_{\varepsilon_2}^* Q &\triangleq P * \varepsilon_1 \Rightarrow_{\varepsilon_2} Q \\ P \varepsilon_1 \Rightarrow_{\varepsilon_2} Q &\triangleq \square(P * \varepsilon_1 \Rightarrow_{\varepsilon_2} Q) \\ P \Rightarrow_{\mathcal{E}} Q &\triangleq P \varepsilon \Rightarrow_{\mathcal{E}} Q \end{aligned}$$

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:

⁶Actually, in the Coq development masks are restricted to a class of sets of natural numbers that contains all finite sets and is closed under union, intersection, difference and complement. The restriction is necessary for engineering reasons to still obtain representation independence: two masks should be *propositionally* equal iff they contain the same invariant names.

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

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

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

9.3 Weakest Precondition

Finally, we can define the core piece of the program logic, the proposition 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 $S : State \times \mathbb{N} \times List(Obs) \times \mathbb{N} \rightarrow iProp$ that interprets the machine state as an Iris proposition, a predicate $\Phi_F : Val \rightarrow iProp$ that serves as postcondition for forked-off threads, and a function $n_{\triangleright} : \mathbb{N} \rightarrow \mathbb{N}$ specifying the number of additional lateres and later credits used for each physical step. The state interpretation can depend on the current physical state, the number of steps since the beginning of the execution, the list of *future* observations as well as the total number of *forked* threads (that is one less than the total number of threads). It should be monotone with respect to the step counter: $S(\sigma, n_s, \vec{\kappa}, n_t) \Rightarrow_{\emptyset} S(\sigma, n_s, \vec{\kappa}, n_t + 1)$. 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. Finally, weakest precondition takes a parameter $s \in \{\text{NotStuck}, \text{Stuck}\}$ indicating whether program execution is allowed to get stuck.

$$\begin{aligned}
wp(S, \Phi_F, s) &\triangleq \mu wp_rec. \lambda \mathcal{E}, e, \Phi. \\
&(\exists v. \text{expr_to_val}(e) = v \wedge \models_{\mathcal{E}} \Phi(v)) \vee \\
&\left(\text{expr_to_val}(e) = \perp \wedge \forall \sigma, n_s, \vec{\kappa}, \vec{\kappa}', n_t. S(\sigma, n_s, \vec{\kappa} \uplus \vec{\kappa}', n_t) \ \varepsilon \Rightarrow_{\emptyset}^* \right. \\
&\quad (s = \text{NotStuck} \Rightarrow \text{red}(e, \sigma)) * \forall e', \sigma', \vec{e}. (e, \sigma \xrightarrow{\vec{\kappa}}_{\mathfrak{t}} e', \sigma', \vec{e}) * \mathfrak{L}(n_{\triangleright}(n_s) + 1) \multimap \\
&\quad (\models_{\emptyset} \triangleright \models_{\emptyset})^{n_{\triangleright}(n_s)+1} \models_{\mathcal{E}} S(\sigma', \vec{\kappa}', n + |\vec{e}|) * wp_rec(\mathcal{E}, e', \Phi) * \\
&\quad \left. \bigstar_{e'' \in \vec{e}} wp_rec(\top, e'', \Phi_F) \right)
\end{aligned}$$

$$wp_{s; \mathcal{E}}^{S; \Phi_F} e \{v. P\} \triangleq wp(S, \Phi_F, s)(\mathcal{E}, e, \lambda v. P)$$

The S and Φ_F will always be set by the context; typically, when instantiating Iris with a language, we also pick the corresponding state interpretation S and fork-postcondition Φ_F . All proof rules

leave S and Φ_F unchanged. If we leave away the mask \mathcal{E} , we assume it to default to \top . If we leave away the stuckness s , it defaults to **NotStuck**.

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

$$\begin{array}{c}
\text{WP-VALUE} \\
P[v/x] \vdash \text{wp}_{s;\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 \quad (s_2 = \text{Stuck} \vee s_1 = s_2)}{\Gamma \mid \text{wp}_{s_1;\mathcal{E}_1} e \{x. P\} \vdash \text{wp}_{s_2;\mathcal{E}_2} e \{x. Q\}} \\
\\
\text{FUP-WP} \qquad \qquad \qquad \text{WP-FUP} \\
\frac{}{\mathbb{H}_{\mathcal{E}} \text{wp}_{s;\mathcal{E}} e \{x. P\} \vdash \text{wp}_{s;\mathcal{E}} e \{x. P\}} \qquad \frac{}{\text{wp}_{s;\mathcal{E}} e \{x. \mathbb{H}_{s;\mathcal{E}} P\} \vdash \text{wp}_{s;\mathcal{E}} e \{x. P\}} \\
\\
\text{WP-ATOMIC} \\
\frac{s = \text{NotStuck} \Rightarrow \text{atomic}(e) \quad s = \text{Stuck} \Rightarrow \text{strongly_atomic}(e)}{\mathcal{E}_1 \mathbb{H}_{\mathcal{E}_2} \text{wp}_{s;\mathcal{E}_2} e \{x. \mathcal{E}_2 \mathbb{H}_{\mathcal{E}_1} P\} \vdash \text{wp}_{s;\mathcal{E}_1} e \{x. P\}} \\
\\
\text{WP-FRAME} \qquad \qquad \qquad \text{WP-FRAME-STEP} \\
\frac{Q * \text{wp}_{s;\mathcal{E}} e \{x. P\} \vdash \text{wp}_{s;\mathcal{E}} e \{x. Q * P\}}{\text{wp}_{s;\mathcal{E}_2} e \{x. P\} * \mathcal{E}_1 \mathbb{H}_{\mathcal{E}_2} \triangleright \mathcal{E}_2 \mathbb{H}_{\mathcal{E}_1} Q \vdash \text{wp}_{s;\mathcal{E}_1} e \{x. Q * P\}} \quad \frac{\text{expr_to_val}(e) = \perp \quad \mathcal{E}_2 \subseteq \mathcal{E}_1}{\text{wp}_{s;\mathcal{E}_2} e \{x. P\} * \mathcal{E}_1 \mathbb{H}_{\mathcal{E}_2} \triangleright \mathcal{E}_2 \mathbb{H}_{\mathcal{E}_1} Q \vdash \text{wp}_{s;\mathcal{E}_1} e \{x. Q * P\}} \\
\\
\text{WP-FRAME-N-STEPS} \\
\frac{\text{expr_to_val}(e) = \perp \quad \mathcal{E}_2 \subseteq \mathcal{E}_1}{(\forall \sigma, n_s, \vec{\kappa}, n_t. S(\sigma, n_s, \vec{\kappa}, n_t) \not\Rightarrow_{\mathcal{E}_1, \emptyset} n \leq n_{\triangleright}(n_s) + 1) \wedge} \\
\text{wp}_{s;\mathcal{E}_2} e \{x. P\} * \mathcal{E}_1 \mathbb{H}_{\mathcal{E}_2} (\triangleright \mathbb{H}_{\emptyset})^n \mathcal{E}_2 \mathbb{H}_{\mathcal{E}_1} Q \\
\vdash \text{wp}_{s;\mathcal{E}_1} e \{x. Q * P\} \\
\\
\text{WP-BIND} \\
\frac{K \text{ is a context}}{\text{wp}_{s;\mathcal{E}} e \{x. \text{wp}_{s;\mathcal{E}} K(\text{val_to_expr}(x)) \{y. P\}\} \vdash \text{wp}_{s;\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. This basically just copies the second branch (the non-value case) of the definition of weakest preconditions.

$$\begin{array}{c}
\text{WP-LIFT-STEP} \\
\frac{\text{expr_to_val}(e_1) = \perp}{\forall \sigma_1, \vec{\kappa}, \vec{\kappa}', n. S(\sigma_1, n_s, \vec{\kappa} + \vec{\kappa}', n_t) \mathcal{E} \Rightarrow_{\emptyset} (s = \text{NotStuck} \Rightarrow \text{red}(e_1, \sigma_1)) *} \\
\forall e_2, \sigma_2, \vec{e}. (e_1, \sigma_1 \xrightarrow{\vec{\kappa}}_t e_2, \sigma_2, \vec{e}) * \mathcal{L}(n_{\triangleright}(n_s) + 1) * (\mathbb{H}_{\emptyset} \triangleright \mathbb{H}_{\emptyset})^{n_{\triangleright}(n_s)} \emptyset \mathbb{H}_{\mathcal{E}} \\
\left(S(\sigma_2, n_s + 1, \vec{\kappa}', n_t + |\vec{e}|) * \text{wp}_{s;\mathcal{E}}^{S;\Phi_F} e_2 \{x. P\} * \bigstar_{e_f \in \vec{e}} \text{wp}_{s;\top}^{S\Phi_F} e_f \{\Phi_F\} \right) \\
\vdash \text{wp}_{s;\mathcal{E}}^{S\Phi_F} e_1 \{x. P\}
\end{array}$$

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. The most general form of the adequacy statement is about proving properties of an arbitrary program execution.

Theorem 3 (Adequacy). *Assume we are given some $\vec{e}_1, \sigma_1, \vec{\kappa}, T_2, \sigma_2$ such that $(\vec{e}_1, \sigma_1) \xrightarrow{\vec{\kappa}, *}_{\text{tp}} (T_2, \sigma_2)$. Moreover, assume we are given a stuckness parameter s and meta-level property p that we want to show. To verify that p holds, it is sufficient to show the following Iris entailment:*

$$\text{True} \vdash \text{Irr}_{\top} \exists S, \vec{\Phi}, \Phi_F. S(\sigma_1, 0, \vec{\kappa}, 0) * \left(\bigstar_{e, \Phi \in \vec{e}_1, \vec{\Phi}} \text{wp}_{s; \top}^{S; \Phi_F} e \{x. \Phi(x)\} \right) * \left(C_s^{S; \vec{\Phi}; \Phi_F}(T_2, \sigma_2) \vdash_{\text{Irr}_{\emptyset}} \hat{p} \right)$$

where C describes states that are consistent with the state interpretation and postconditions:

$$\begin{aligned} C_s^{S; \vec{\Phi}; \Phi_F}(T_2, \sigma_2) \triangleq & \exists \vec{e}_2, T'_2. T_2 = \vec{e}_2 \text{ ++ } T'_2 * \\ & |\vec{e}_1| = |\vec{e}_2| * \\ & (s = \text{NotStuck} \Rightarrow \forall e \in T_2. \text{expr_to_val}(e) \neq \perp \vee \text{red}(e, \sigma_2)) * \\ & S(\sigma_2, (), |T'_2|) * \\ & \left(\bigstar_{e, \Phi \in \vec{e}_2, \vec{\Phi}} \text{expr_to_val}(e) \neq \perp * \Phi(\text{expr_to_val}(e)) \right) * \\ & \left(\bigstar_{e \in T'_2} \text{expr_to_val}(e) \neq \perp * \Phi_F(\text{expr_to_val}(e)) \right) \end{aligned}$$

The \hat{p} here arises because we need a way to talk about p inside Iris. To this end, we assume that the signature \mathcal{S} contains some assertion \hat{p} :

$$\hat{p} : iProp \in \mathcal{F}$$

Furthermore, we assume that the interpretation $\llbracket \hat{p} \rrbracket$ of \hat{p} reflects p (also see §6):

$$\begin{aligned} \llbracket \hat{p} \rrbracket & : \llbracket iProp \rrbracket \\ \llbracket \hat{p} \rrbracket & \triangleq \lambda_. \{n \mid p\} \end{aligned}$$

The signature can of course state arbitrary additional properties of \hat{p} , as long as they are proven sound.

In other words, to show that p holds, we have to prove an entailment in Iris that, starting from the empty context, chooses some state interpretation, postconditions for the initial threads, forked-thread postcondition and stuckness and then proves:

- the initial state interpretation,
- a weakest-precondition,
- and a view shift showing the desired \hat{p} under the extra assumption $C(T_2, \sigma_2)$.

Notice that the state interpretation and the postconditions are chosen *after* doing a fancy update, which allows them to depend on the names of ghost variables that are picked in that initial fancy update. This gives us a chance to allocate some “global” ghost state that state interpretation and postcondition can refer to (*e.g.*, the name $\gamma_{Credits}$).

$C_s^{S;\vec{\Phi};\Phi_F}(T_2, \sigma_2)$ says that:

- The final thread-pool T_2 contains the final state of the initial threads \vec{e}_2 , and any number of additional forked threads in T'_2 .
- If this is a stuck-free weakest precondition, then all threads in the final thread-pool are either values or are reducible in the final state σ_2 .
- The state interpretation S holds for the final state.
- If one of the initial threads reduced to a value, the corresponding post-condition $\Phi \in \vec{\Phi}$ holds for that value.
- If any other thread reduced to a value, the forked-thread post-condition Φ_F holds for that value.

As an example for how to use this adequacy theorem, let us say we wanted to prove that a program e_1 for which we derived a **NotStuck** weakest-precondition cannot get stuck:

Corollary 1 (Stuck-freedom). *Assume we are given some e_1 such that the following holds:*

$$\text{True} \vdash \forall \sigma_1, \vec{\kappa}. \models_{\top} \exists S, \Phi, \Phi_F. S(\sigma_1, 0, \vec{\kappa}, 0) * \text{wp}_{\text{NotStuck}; \top}^{S; \Phi_F} e_1 \{x. \Phi(x)\}$$

Then it is the case that:

$$\forall \sigma_1, \vec{\kappa}, T_2, \sigma_2. ([e_1], \sigma_1) \xrightarrow{\vec{\kappa}}_{\text{tp}}^* (T_2, \sigma_2) \Rightarrow \forall e \in T_2. \text{expr_to_val}(e) \neq \perp \vee \text{red}(e, \sigma_2)$$

To prove the conclusion of this corollary, we assume some $\sigma_1, \vec{\kappa}, T_2, \sigma_2$ and $([e_1], \sigma_1) \xrightarrow{\vec{\kappa}}_{\text{tp}}^* (T_2, \sigma_2)$, and we instantiate the main theorem with this execution and $p \triangleq \forall e \in T_2. \text{expr_to_val}(e) \neq \perp \vee \text{red}(e, \sigma_2)$. We can then show the premise of adequacy using the Iris entailment that we assumed in the corollary and:

$$\text{True} \vdash C_{\text{NotStuck}}^{S; [\Phi]; \Phi_F}(T_2, \sigma_2) \top \Rightarrow_{\emptyset} p$$

This proof, just like the following, also exploits that we can freely swap between meta-level universal quantification ($\forall x. \text{True} \vdash P$) and quantification in Iris ($\text{True} \vdash \forall x. P$).

Similarly we could show that the postcondition makes adequate statements about the possible final value of the main thread:

Corollary 2 (Adequate postcondition). *Assume we are given some e_1 and a set $V \subseteq \text{Val}$ such that the following holds (assuming we can talk about sets like V inside the logic):*

$$\text{True} \vdash \forall \sigma_1, \vec{\kappa}. \models_{\top} \exists S, \Phi_F. S(\sigma_1, 0, \vec{\kappa}, 0) * \text{wp}_{s; \top}^{S; \Phi_F} e_1 \{x. x \in V\}$$

Then it is the case that:

$$\forall \sigma_1, \vec{\kappa}, v_2, T_2, \sigma_2. ([e_1], \sigma_1) \xrightarrow{\vec{\kappa}}_{\text{tp}}^* ([\text{val_to_expr}(v_2)] \# T_2, \sigma_2) \Rightarrow v_2 \in V$$

To show this, we assume some $\sigma_1, \vec{\kappa}, v_2, T_2, \sigma_2$ such that $([e_1], \sigma_1) \xrightarrow{\vec{\kappa}}_{\text{tp}}^* ([\text{val_to_expr}(v_2)] \uparrow T_2, \sigma_2)$, and we instantiate adequacy with this execution and $p \triangleq v_2 \in \text{Val}$. Then we only have to show:

$$\text{True} \vdash C_s^{S;[(\lambda v. v \in \text{Val})];\Phi_F}([\text{val_to_expr}(v_2)] \uparrow T_2, \sigma_2) \dashv\equiv_{\emptyset} v_2 \in \text{Val}$$

As a final example, we could use adequacy to show that the state σ of the program is always in some set $\Sigma \subseteq \text{State}$:

Corollary 3 (Adequate state interpretation). *Assume we are given some e_1 and a set $\Sigma \subseteq \text{State}$ such that the following holds (assuming we can talk about sets like Σ inside the logic):*

$$\text{True} \vdash \forall \sigma_1, \vec{\kappa}. \exists_{\top} \exists S, \Phi, \Phi_F. S(\sigma_1, 0, \vec{\kappa}, 0) * \text{wp}_{s;\top}^{S;\Phi_F} e_1 \{\Phi\} * (\forall \sigma_2, n_s, n_t. S(\sigma_2, n_s, (), n_t) \dashv\equiv_{\emptyset} \sigma_2 \in \Sigma)$$

Then it is the case that:

$$\forall \sigma_1, \vec{\kappa}, T_2, \sigma_2. ([e_1], \sigma_1) \xrightarrow{\vec{\kappa}}_{\text{tp}}^* (T_2, \sigma_2) \Rightarrow \sigma_2 \in \Sigma$$

To show this, we assume some $\sigma_1, \vec{\kappa}, T_2, \sigma_2$ such that $([e_1], \sigma_1) \xrightarrow{\vec{\kappa}}_{\text{tp}}^* (T_2, \sigma_2)$, and we instantiate adequacy with this execution and $p \triangleq \sigma_2 \in \Sigma$. Then we have to show:

$$(\forall \sigma_2, n_s, n_t. S(\sigma_2, n_s, (), n_t) \dashv\equiv_{\emptyset} \sigma_2 \in \Sigma) \vdash C_s^{S;[\Phi];\Phi_F}(T_2, \sigma_2) \dashv\equiv_{\emptyset} \sigma_2 \in \Sigma$$

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:

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

We assume the state interpretation S to be fixed by the context.

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.

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

9.4 Invariant Namespaces

In §9.2, we defined a proposition $\text{Inv}^l(P)$ expressing knowledge (*i.e.*, the proposition is persistent) that P is maintained as invariant with name l . The concrete name l 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 $\mathcal{N}.l$ for the namespace $[l] \uparrow \mathcal{N}$. (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 §9.2, 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:

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

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

$$\text{Inv}^{\mathcal{N}}(P) \triangleq \exists l \in \mathcal{N}^\dagger. \text{Inv}^l(P)$$

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

$$\begin{array}{c} \text{INV-PERSIST} \\ \text{Inv}^{\mathcal{N}}(P) \vdash \square \text{Inv}^{\mathcal{N}}(P) \end{array} \quad \begin{array}{c} \text{INV-ALLOC} \\ \triangleright P \vdash \text{Inv}^{\mathcal{N}}(P) \end{array} \quad \begin{array}{c} \text{INV-OPEN} \\ \frac{\mathcal{N} \subseteq \mathcal{E}}{\text{Inv}^{\mathcal{N}}(P) \ \varepsilon \Rightarrow_{\varepsilon \setminus \mathcal{N}} \triangleright P * (\triangleright P \ \varepsilon \setminus \mathcal{N} \Rightarrow *_{\varepsilon} \text{True})} \end{array}$$

$$\begin{array}{c} \text{INV-OPEN-TIMELESS} \\ \frac{\mathcal{N} \subseteq \mathcal{E} \quad \text{timeless}(P)}{\text{Inv}^{\mathcal{N}}(P) \ \varepsilon \Rightarrow_{\varepsilon \setminus \mathcal{N}} P * (P \ \varepsilon \setminus \mathcal{N} \Rightarrow *_{\varepsilon} \text{True})} \end{array}$$

9.5 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 propositions of the form

$$P \quad \varepsilon_1 \Rightarrow_{\varepsilon_2} \exists x. Q * (\forall y. Q' \quad \varepsilon_2 \Rightarrow_{\varepsilon_1} * R)$$

One way to think about such propositions is as follows: Given some accessor, if during our verification we have the proposition P and the mask ε_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 ε_2 . Additionally, opening the accessor provides us with $\forall y. Q' \quad \varepsilon_2 \Rightarrow_{\varepsilon_1} * R$, 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 ε_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:

$$\frac{\text{ACC-VS} \quad P \quad \varepsilon_1 \Rightarrow_{\varepsilon_2} \exists x. Q * (\forall y. Q' \quad \varepsilon_2 \Rightarrow_{\varepsilon_1} * R) \quad \forall x. Q * P_F \Rightarrow_{\varepsilon_2} \exists y. Q' * P_F}{P * P_F \Rightarrow_{\varepsilon_1} R * P_F}$$

$$\frac{\text{ACC-HT} \quad P \quad \varepsilon_1 \Rightarrow_{\varepsilon_2} \exists x. Q * (\forall y. Q' \quad \varepsilon_2 \Rightarrow_{\varepsilon_1} * R) \quad \forall x. \{Q * P_F\} e \{\exists y. Q' * P_F\}_{\varepsilon_2} \quad \text{atomic}(e)}{\{P * P_F\} e \{R * P_F\}_{\varepsilon_1}}$$

Furthermore, in the special case that $\varepsilon_1 = \varepsilon_2$, 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 symmetric case where $P = R$ and $Q = Q'$, we use the following notation that avoids repetition:

$$P \quad \varepsilon_1 \propto_{\varepsilon_2} x. Q \triangleq P \quad \varepsilon_1 \Rightarrow_{\varepsilon_2} \exists x. Q * (Q \quad \varepsilon_2 \Rightarrow_{\varepsilon_1} * P)$$

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

Accessor-style invariants. In fact, the user-visible notion of invariants $\boxed{P}^{\mathcal{N}}$ is defined via **INV-OPEN**:

$$\boxed{P}^{\mathcal{N}} \triangleq \square \forall \mathcal{E}. \varepsilon \Vdash_{\varepsilon \setminus \mathcal{N}} \triangleright P * (\triangleright P \quad \varepsilon \setminus \mathcal{N} \Rightarrow_{\varepsilon} * \text{True})$$

All the invariant laws shown above for $\text{Inv}^{\mathcal{N}}(P)$ also hold for $\boxed{P}^{\mathcal{N}}$, but we can also show some additional laws that would otherwise not hold:

$$\frac{\text{INV-COMBINE} \quad \mathcal{N}_1 \# \mathcal{N}_2 \quad \mathcal{N}_1 \cup \mathcal{N}_2 \subseteq \mathcal{N}}{\boxed{P_1}^{\mathcal{N}_1} * \boxed{P_2}^{\mathcal{N}_2} \vdash \boxed{P_1 * P_2}^{\mathcal{N}}} \quad \text{INV-SPLIT} \quad \boxed{P_1 * P_2}^{\mathcal{N}} \vdash \boxed{P_1}^{\mathcal{N}} * \boxed{P_2}^{\mathcal{N}}$$

$$\text{INV-ALTER} \quad \triangleright \square(P \multimap Q * (Q \multimap P)) \vdash \boxed{P}^{\mathcal{N}} \multimap \boxed{Q}^{\mathcal{N}}$$

10 Derived Constructions

10.1 Cancellable Invariants

Iris invariants as described in §9.2 are persistent—once established, they hold forever. However, based on them, it is possible to *encode* a form of invariants that can be “cancelled” again.

First, we need some ghost state:

$$CInvTok \triangleq Frac$$

Now we define:

$$[CInv : \gamma]_q \triangleq [\underline{q}]_1^{\gamma}$$

$$CInv^{\gamma, \mathcal{N}}(P) \triangleq \boxed{P \vee [\underline{1}]_1^{\gamma}}^{\mathcal{N}}$$

It is then straightforward to prove:

$$\begin{array}{c} \text{CINV-NEW} \\ \triangleright P \Rightarrow_{\perp} \exists \gamma. [CInv : \gamma]_1 * \square CInv^{\gamma, \mathcal{N}}(P) \end{array} \qquad \begin{array}{c} \text{CINV-ACC} \\ CInv^{\gamma, \mathcal{N}}(P) \vdash [CInv : \gamma]_q \quad \mathcal{N}^{\alpha_{\emptyset}} \triangleright P \end{array}$$

$$\begin{array}{c} \text{CINV-CANCEL} \\ CInv^{\gamma, \mathcal{N}}(P) \vdash [CInv : \gamma]_1 \Rightarrow_{\mathcal{N}} \triangleright P \end{array}$$

Cancellable invariants are useful, for example, when reasoning about data structures that will be deallocated: Every reference to the data structure comes with a fraction of the token, and when all fractions have been gathered, **CINV-CANCEL** is used to cancel the invariant, after which the data structure can be deallocated.

10.2 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 tokens⁷ that guard access to non-atomic invariants. Having the token $[Nalnv : p.\mathcal{E}]$ 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:

$$PId \triangleq GName$$

$$NaTok \triangleq \wp^{\text{fin}}(InvName) \times \wp(InvName)$$

⁷Very much like the tokens that are used to encode “normal”, atomic invariants

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:

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

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

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

We easily obtain:

$$\begin{array}{ll} \text{NAINV-NEW-POOL} & \text{NAINV-TOK-SPLIT} \\ \text{True} \Rightarrow_{\perp} \exists p. [\text{Nalnv} : p] & [\text{Nalnv} : p.\mathcal{E}_1 \uplus \mathcal{E}_2] \Leftrightarrow [\text{Nalnv} : p.\mathcal{E}_1] * [\text{Nalnv} : p.\mathcal{E}_2] \\ \\ \text{NAINV-NEW-INV} & \text{NAINV-ACC} \\ \triangleright P \Rightarrow_{\mathcal{N}} \square \text{Nalnv}^{p.\mathcal{N}}(P) & \text{Nalnv}^{p.\mathcal{N}}(P) \vdash [\text{Nalnv} : p.\mathcal{N}] \propto_{\mathcal{N}} \triangleright P \end{array}$$

from which we can derive

$$\frac{\mathcal{N} \subseteq \mathcal{E}}{\text{Nalnv}^{p.\mathcal{N}}(P) \vdash [\text{Nalnv} : p.\mathcal{E}] \propto_{\mathcal{N}} \triangleright P * [\text{Nalnv} : p.\mathcal{E} \setminus \mathcal{N}]}$$

10.3 Boxes

The idea behind the *boxes* is to have a proposition 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 propositions 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 a proposition P . A box consists of a bunch of *slices* which decompose P into a separating conjunction of the propositions Q_{ι} governed by the individual slices. Each slice is either *full* (it right now contains Q_{ι}), or *empty* (it does not contain anything currently). The proposition 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 propositions are: $\text{BoxSlice}(\mathcal{N}, P, \iota)$, saying that there is a slice in namespace \mathcal{N} with name ι 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 $\text{BoxState} \triangleq \{\text{full}, \text{empty}\}$.

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

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

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

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

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

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

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

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 camera as follows:

$$\begin{array}{l} \text{BoxState} \triangleq \text{full} + \text{empty} \\ \text{Box} \triangleq \text{Auth}(\text{Ex}(\text{BoxState})?) \times \text{Ag}(\blacktriangleright i\text{Prop})? \end{array}$$

Now we can define the propositions:

$$\begin{array}{l} \text{SliceInv}(\iota, P) \triangleq \exists b. \boxed{[\bullet b, \varepsilon]}^{\iota} * ((b = \text{full}) \Rightarrow P) \\ \text{BoxSlice}(\mathcal{N}, P, \iota) \triangleq \boxed{[\varepsilon, P]}^{\iota} * \boxed{\text{SliceInv}(\iota, P)}^{\mathcal{N}} \\ \text{Box}(\mathcal{N}, P, f) \triangleq \exists Q : \mathbb{N} \rightarrow i\text{Prop}. \triangleright \left(P = \bigstar_{\iota \in \text{dom}(f)} Q(\iota) \right) * \\ \bigstar_{\iota \in \text{dom}(f)} \boxed{[\circ f(\iota), Q(\iota)]}^{\iota} * \boxed{\text{SliceInv}(\iota, Q(\iota))}^{\mathcal{N}} \end{array}$$

Derived rules. Here are some derived rules:

SLICE-INSERT-FULL

$$\triangleright Q * \triangleright^b \text{Box}(\mathcal{N}, P, f) \Rightarrow_{\mathcal{N}} \exists \iota \notin \text{dom}(f). \square \text{BoxSlice}(\mathcal{N}, Q, \iota) * \triangleright^b \text{Box}(\mathcal{N}, P * Q, f[\iota \leftarrow \text{full}])$$

SLICE-DELETE-FULL

$$\frac{f(\iota) = \text{full}}{\text{BoxSlice}(\mathcal{N}, Q, \iota) \vdash \triangleright^b \text{Box}(\mathcal{N}, P, f) \Rightarrow_{\mathcal{N}} \triangleright Q * \exists P'. \triangleright^b (\triangleright (P = P' * Q) * \text{Box}(\mathcal{N}, P', f[\iota \leftarrow \perp]))}$$

SLICE-SPLIT

$$\frac{f(\iota) = s}{\text{BoxSlice}(\mathcal{N}, Q_1 * Q_2, \iota) \vdash \triangleright^b \text{Box}(\mathcal{N}, P, f) \Rightarrow_{\mathcal{N}} \exists \iota_1 \notin \text{dom}(f), \iota_2 \notin \text{dom}(f). \iota_1 \neq \iota_2 \wedge \square \text{BoxSlice}(\mathcal{N}, Q_1, \iota_1) * \square \text{BoxSlice}(\mathcal{N}, Q_2, \iota_2) * \triangleright^b \text{Box}(\mathcal{N}, P, f[\iota \leftarrow \perp][\iota_1 \leftarrow s][\iota_2 \leftarrow s])}$$

SLICE-MERGE

$$\frac{\iota_1 \neq \iota_2 \quad f(\iota_1) = f(\iota_2) = s}{\text{BoxSlice}(\mathcal{N}, Q_1, \iota_1), \text{BoxSlice}(\mathcal{N}, Q_2, \iota_2) \vdash \triangleright^b \text{Box}(\mathcal{N}, P, f) \Rightarrow_{\mathcal{N}} \exists \iota \notin \text{dom}(f) \setminus \{\iota_1, \iota_2\}. \square \text{BoxSlice}(\mathcal{N}, Q_1 * Q_2, \iota) * \triangleright^b \text{Box}(\mathcal{N}, P, f[\iota_1 \leftarrow \perp][\iota_2 \leftarrow \perp][\iota \leftarrow s])}$$

11 Logical Paradoxes

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

11.1 Saved Propositions without a Later

As a preparation for the proof about invariants in §11.2, 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 $\dot{\Rightarrow}$ satisfying the rules in §5.

Theorem 4. *If there exists a type $GName$ and a proposition $_ \dot{\Rightarrow} _ : GName \rightarrow iProp \rightarrow iProp$ associating names $\gamma : GName$ to propositions and satisfying:*

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

then $\vdash \dot{\Rightarrow} \text{False}$.

The type $GName$ should be thought of as the type of “locations” and $\gamma \dot{\Rightarrow} P$ should be read as stating that location γ “stores” proposition P . Notice that these are immutable locations, so the maps-to proposition 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 4** is the following proposition:

Definition 22. $A(\gamma) \triangleq \exists P : iProp. \Box \neg P \wedge \gamma \dot{\Rightarrow} P$.

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 $P \triangleq \gamma \dot{\Rightarrow} A(\gamma)$: this says that the proposition named γ is the proposition saying that it, itself, does not hold. In other words, P says that the proposition named γ expresses its own negation. Unsurprisingly, that leads to a contradiction, as is shown in the following lemma:

Lemma 2. *We have $\gamma \dot{\Rightarrow} A(\gamma) \vdash \Box \neg A(\gamma)$ and $\gamma \dot{\Rightarrow} A(\gamma) \vdash A(\gamma)$.*

Proof.

- First we show $\gamma \dot{\Rightarrow} A(\gamma) \vdash \Box \neg A(\gamma)$. Since $\gamma \dot{\Rightarrow} A(\gamma)$ is persistent it suffices to show $\gamma \dot{\Rightarrow} A(\gamma) \vdash \neg A(\gamma)$. Suppose $\gamma \dot{\Rightarrow} A(\gamma)$ and $A(\gamma)$. Then by definition of A there is a P such that $\Box \neg P$ and $\gamma \dot{\Rightarrow} P$. 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 4](#) is simple.

Proof of [Theorem 4](#). Using the previous lemmas we have

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

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

11.2 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 that it also applies to previous, less powerful versions of Iris.

Theorem 5. *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 a proposition $\Box[\cdot] : InvName \rightarrow iProp \rightarrow iProp$ satisfying:

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

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

$$\gamma \Rightarrow P \triangleq \exists l. \boxed{[\underline{S}]_1^{\gamma} \vee [\underline{F}]_1^{\gamma} * \square P}^l.$$

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

We use the tokens $[\underline{S}]_1^{\gamma}$ and $[\underline{F}]_1^{\gamma}$ to model invariants that can be initialized “lazily”: $[\underline{S}]_1^{\gamma}$ indicates that the invariant is still not initialized, whereas the duplicable $[\underline{F}]_1^{\gamma}$ 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 $\vdash \Rightarrow_1 \exists \gamma. \gamma \Rightarrow P(\gamma)$.

Proof. We have to show the allocation rule

$$\vdash \Rightarrow_1 \exists \gamma. \gamma \Rightarrow P.$$

From **START-ALLOC** we have a γ such that $\Rightarrow_0 [\underline{S}]_1^{\gamma}$ holds and hence from **WEAKEN-MASK** we have $\Rightarrow_1 [\underline{S}]_1^{\gamma}$. Since we are proving a goal of the form $\Rightarrow_1 R$ we may assume $[\underline{S}]_1^{\gamma}$. Thus for any P we have $\Rightarrow_1 ([\underline{S}]_1^{\gamma} \vee [\underline{F}]_1^{\gamma} * P)$. Again since our goal is still of the form \Rightarrow_1 we may assume $[\underline{S}]_1^{\gamma} \vee [\underline{F}]_1^{\gamma} * P$. The rule **INV-ALLOC** then gives us precisely what we need. \square

Lemma 4. We have $\gamma \Rightarrow P * \gamma \Rightarrow Q * \square P \vdash \Rightarrow_1 \square Q$ and thus $\gamma \Rightarrow P * \gamma \Rightarrow Q \vdash (\Rightarrow_1 \square P) \Leftrightarrow (\Rightarrow_1 \square Q)$.

Proof.

- We first show

$$\gamma \Rightarrow P * \gamma \Rightarrow Q * \square P \vdash \Rightarrow_1 \square Q.$$

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

1. $[\underline{S}]_1^{\gamma}$ (the invariant is “uninitialized”): In this case, we use **START-FINISH** to “initialize” the invariant and obtain $[\underline{F}]_1^{\gamma}$. Then we duplicate $[\underline{F}]_1^{\gamma}$, and use it together with $\square P$ to close the invariant.
2. $[\underline{F}]_1^{\gamma} * \square P$ (the invariant is “initialized”): In this case we duplicate $[\underline{F}]_1^{\gamma}$, and use a copy to close the invariant.

After closing the invariant, we have obtained $[\underline{F}]_1^{\gamma}$. Hence, it is sufficient to prove

$$[\underline{F}]_1^{\gamma} * \gamma \Rightarrow P * \gamma \Rightarrow Q * \square P \vdash \Rightarrow_1 \square Q.$$

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

1. $[\underline{S}]_1^{\gamma}$ (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. $\overline{\overline{\overline{\mathbb{F}}_1}}^\gamma * \Box Q$ (the invariant is “initialized”): Since $\Box Q$ is duplicable we use one copy to close the invariant, and retain another to prove $\overline{\overline{\overline{\mathbb{F}}_1}} \Box Q$.
- By applying the above twice, we easily obtain

$$\gamma \Rightarrow P * \gamma \Rightarrow Q \vdash (\overline{\overline{\overline{\mathbb{F}}_1}} \Box P) \Leftrightarrow (\overline{\overline{\overline{\mathbb{F}}_1}} \Box Q)$$

□

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 [§11.1](#). The only difference is that because [Lemma 4](#) is slightly weaker than the rule [SPROP-AGREE](#) of [Theorem 4](#), we need to use $\overline{\overline{\overline{\mathbb{F}}_1}} \text{False}$ in place of False in the definition of the predicate A : we let $A(\gamma) \triangleq \exists P : iProp. \Box(P \Rightarrow \overline{\overline{\overline{\mathbb{F}}_1}} \text{False}) \wedge \gamma \Rightarrow P$ and replay the proof that we have presented above.

12 HeapLang

So far, we have treated the programming language we work on entirely generically. In this section we present the default language that ships with Iris, HeapLang. HeapLang is an ML-like languages with a higher-order heap, unstructured concurrency, some common atomic operations, and prophecy variables. It is an instance of the language interface (§8), so we only define a per-thread small-step operational semantics—the thread-pool semantics are given in in §8.1.

12.1 HeapLang syntax and operational semantics

The grammar of HeapLang, and in particular its set *Expr* of *expressions* and *Val* of *values*, is defined as follows:

$$\begin{aligned}
 v, w \in \text{Val} &::= z \mid \mathbf{true} \mid \mathbf{false} \mid () \mid \mathfrak{A} \mid \ell \mid p \mid & (z \in \mathbb{Z}, \ell \in \text{Loc}, p \in \text{ProphId}) \\
 &\quad \mathbf{rec}_v f(x) = e \mid (v, w)_v \mid \mathbf{inl}(v)_v \mid \mathbf{inr}(v)_v \\
 e \in \text{Expr} &::= v \mid x \mid \mathbf{rec}_e f(x) = e \mid e_1(e_2) \mid \\
 &\quad \odot_1 e \mid e_1 \odot_2 e_2 \mid \mathbf{if } e \mathbf{ then } e_1 \mathbf{ else } e_2 \mid \\
 &\quad (e_1, e_2)_e \mid \mathbf{fst}(e) \mid \mathbf{snd}(e) \mid \\
 &\quad \mathbf{inl}(e)_e \mid \mathbf{inr}(e)_e \mid \mathbf{match } e \mathbf{ with inl } \Rightarrow e_1 \mid \mathbf{inr} \Rightarrow e_2 \mathbf{ end} \mid \\
 &\quad \mathbf{Alloc}(e_1, e_2) \mid \mathbf{Free}(e) \mid !e \mid e_1 \leftarrow e_2 \mid \mathbf{CmpXchg}(e_1, e_2, e_3) \mid \mathbf{Xchg}(e_1, e_2) \mid \mathbf{FAA}(e_1, e_2) \mid \\
 &\quad \mathbf{fork } \{e\} \mid \mathbf{newproph} \mid \mathbf{resolve with } e_1 \mathbf{ at } e_2 \mathbf{ to } e_3 \\
 \odot_1 &::= - \mid \dots \quad (\text{list incomplete}) \\
 \odot_2 &::= + \mid - \mid +_L \mid = \mid \dots \quad (\text{list incomplete})
 \end{aligned}$$

(Note that **match** contains a literal `|` that is not part of the BNF but part of HeapLang syntax.)

To simplify the formalization, the only binders occur in **rec**. **match** has a closure in each arm which will be applied to the value of the left/right variant, respectively. (See the syntactic sugar defined later.)

Recursive abstractions, pairs, and the injections exist both as a value form and an expression form. The expression forms will reduce to the value form once all their arguments are values. Conceptually, one can think of that as corresponding to “boxing” that most functional language implementations do. We will leave away the disambiguating subscript when it is clear from the context or does not matter. All of this lets us define `val_to_expr` as simply applying the value injection (the very first syntactic form of *Expr*), which makes a lot of things in Coq much simpler. `expr_to_val` is defined recursively in the obvious way.

Alloc takes as first argument the number of heap cells to allocate (must be strictly positive), and as second argument the default value to use for these heap cells. This lets one allocate arrays. `+L` implements pointer arithmetic (the left operand must be a pointer, the right operand an integer), which is used to access array elements.

For our set of states and observations, we pick

$$\begin{aligned}
\ell \ni Loc &\triangleq \mathbb{Z} \\
p \ni ProphId &\triangleq \mathbb{Z} \\
\sigma \ni State &\triangleq \left\{ \begin{array}{l} \text{HEAP} : Loc \xrightarrow{\text{fin}} Val, \\ \text{PROPHS} : \wp(\text{ProphId}) \end{array} \right\} \\
\kappa \ni Obs &\triangleq \text{ProphId} \times (Val \times Val)
\end{aligned}$$

The HeapLang operational semantics is defined via the use of *evaluation contexts*:

$$\begin{aligned}
K \in Ctx &::= \bullet \mid Ctx_{>} \\
K_{>} \in Ctx_{>} &::= e(K) \mid K(v) \mid \\
&\quad \odot_1 K \mid e \odot_2 K \mid K \odot_2 v \mid \mathbf{if} K \mathbf{then} e_1 \mathbf{else} e_2 \mid \\
&\quad (e, K) \mid (K, v) \mid \mathbf{fst}(K) \mid \mathbf{snd}(K) \mid \\
&\quad \mathbf{inl}(K) \mid \mathbf{inr}(K) \mid \mathbf{match} K \mathbf{with} \mathbf{inl} \Rightarrow e_1 \mid \mathbf{inr} \Rightarrow e_2 \mathbf{end} \mid \\
&\quad \mathbf{Alloc}(e, K) \mid \mathbf{Alloc}(K, v) \mid \mathbf{Free}(K) \mid !K \mid e \leftarrow K \mid K \leftarrow v \mid \\
&\quad \mathbf{CmpXchg}(e_1, e_2, K) \mid \mathbf{CmpXchg}(e_1, K, v_3) \mid \mathbf{CmpXchg}(K, v_2, v_3) \mid \\
&\quad \mathbf{Xchg}(e, K) \mid \mathbf{Xchg}(K, v) \mid \mathbf{FAA}(e, K) \mid \mathbf{FAA}(K, v) \mid \\
&\quad \mathbf{resolve} \mathbf{with} e_1 \mathbf{at} e_2 \mathbf{to} K \mid \mathbf{resolve} \mathbf{with} e_1 \mathbf{at} K \mathbf{to} v_3 \mid \\
&\quad \mathbf{resolve} \mathbf{with} K_{>} \mathbf{at} v_2 \mathbf{to} v_3
\end{aligned}$$

Note that we use right-to-left evaluation order. This means in a curried function call $f(x)(y)$, we know syntactically the arguments will all evaluate before f gets to do anything, which makes specifying curried calls a lot easier.

The **resolve** evaluation context for the leftmost expression (the nested expression that executes atomically together with the prophecy resolution) is special: it must not be empty; only further nested evaluation contexts are allowed. **resolve** takes care of reducing the expression once the nested contexts are taken care of, and at that point it requires the expression to reduce to a value in exactly one step. Hence we define $Ctx_{>}$ for non-empty evaluation contexts. For more details on prophecy variables, see [7].

This lets us define the primitive reduction relation in terms of a “head step” reduction; see Figure 1 and Figure 2. Comparison (both for **CmpXchg** and the binary comparison operator) is a bit tricky and uses a helper judgment (Figure 3). Basically, two values can only be compared if at least one of them is “compare-safe”. Compare-safe values are basic literals (integers, Booleans, locations, unit) as well as **inl** and **inr** of those literals. The intention of this is to forbid directly comparing large values such as pairs, which could not be done in a single atomic step on a real machine.

Per-thread reduction

$$e_1, \sigma_1 \xrightarrow{\vec{\kappa}}_t e_2, \sigma_2, \vec{e}$$

$$\frac{e_1, \sigma_1 \xrightarrow{\vec{\kappa}}_h e_2, \sigma_2, \vec{e}}{K[e_1], \sigma_1 \xrightarrow{\vec{\kappa}}_t K[e_2], \sigma_2, \vec{e}}$$

“Head” reduction (pure)

$$e_1, \sigma_1 \xrightarrow{\vec{\kappa}}_h e_2, \sigma_2, \vec{e}$$

“Boxing” reductions

$$\begin{aligned} (\mathbf{rec}_e f(x) = e, \sigma) &\xrightarrow{\epsilon}_h (\mathbf{rec}_v f(x) = e, \sigma, \epsilon) \\ ((v_1, v_2)_e, \sigma) &\xrightarrow{\epsilon}_h ((v_1, v_2)_v, \sigma, \epsilon) \\ (\mathbf{inl}(v)_e, \sigma) &\xrightarrow{\epsilon}_h (\mathbf{inl}(v)_v, \sigma, \epsilon) \\ (\mathbf{inr}(v)_e, \sigma) &\xrightarrow{\epsilon}_h (\mathbf{inr}(v)_v, \sigma, \epsilon) \end{aligned}$$

Pure reductions

$$\begin{aligned} ((\mathbf{rec}_v f(x) = e)(v), \sigma) &\xrightarrow{\epsilon}_h (e[(\mathbf{rec} f(x) = e)/f][v/x], \sigma, \epsilon) \\ (- \odot z, \sigma) &\xrightarrow{\epsilon}_h (-z, \sigma, \epsilon) \\ (z_1 + \odot z_2, \sigma) &\xrightarrow{\epsilon}_h (z_1 + z_2, \sigma, \epsilon) \\ (z_1 - \odot z_2, \sigma) &\xrightarrow{\epsilon}_h (z_1 - z_2, \sigma, \epsilon) \\ (\ell +_L z, \sigma) &\xrightarrow{\epsilon}_h (\ell + z, \sigma, \epsilon) \\ (v_1 = \odot v_2, \sigma) &\xrightarrow{\epsilon}_h (\mathbf{true}, \sigma, \epsilon) && \text{if } v_1 \cong v_2 \\ (v_1 = \odot v_2, \sigma) &\xrightarrow{\epsilon}_h (\mathbf{false}, \sigma, \epsilon) && \text{if } v_1 \not\cong v_2 \\ (\mathbf{if true then } e_1 \mathbf{ else } e_2, \sigma) &\xrightarrow{\epsilon}_h (e_1, \sigma, \epsilon) \\ (\mathbf{if false then } e_1 \mathbf{ else } e_2, \sigma) &\xrightarrow{\epsilon}_h (e_2, \sigma, \epsilon) \\ (\mathbf{fst}((v_1, v_2)_v), \sigma) &\xrightarrow{\epsilon}_h (v_1, \sigma, \epsilon) \\ (\mathbf{snd}((v_1, v_2)_v), \sigma) &\xrightarrow{\epsilon}_h (v_2, \sigma, \epsilon) \\ (\mathbf{match inl}(v)_v \mathbf{ with inl } \Rightarrow e_1 \mid \mathbf{inr} \Rightarrow e_2 \mathbf{ end}, \sigma) &\xrightarrow{\epsilon}_h (e_1(v), \sigma, \epsilon) \\ (\mathbf{match inr}(v)_v \mathbf{ with inl } \Rightarrow e_1 \mid \mathbf{inr} \Rightarrow e_2 \mathbf{ end}, \sigma) &\xrightarrow{\epsilon}_h (e_2(v), \sigma, \epsilon) \end{aligned}$$

Figure 1: HeapLang pure and boxed reduction rules.

The \odot subscript indicates that this is the HeapLang operator, not the mathematical operator.

“Head” reduction (impure)

$$e_1, \sigma_1 \xrightarrow{\vec{\kappa}}_h e_2, \sigma_2, \vec{e}$$

Heap reductions

$$\begin{array}{ll}
(\mathbf{Alloc}(z, v), \sigma) \xrightarrow{\epsilon}_h (\ell, \sigma : \text{HEAP}[[\ell, \ell + z] \leftarrow v], \epsilon) & \text{if } z > 0 \text{ and } \forall i < z. \sigma.\text{HEAP}(\ell + i) = \perp \\
(\mathbf{Free}(\ell), \sigma) \xrightarrow{\epsilon}_h ((), \sigma : \text{HEAP}[\ell \leftarrow \perp], \epsilon) & \text{if } \sigma.\text{HEAP}(\ell) = v \\
(!\ell, \sigma) \xrightarrow{\epsilon}_h (v, \sigma, \epsilon) & \text{if } \sigma.\text{HEAP}(\ell) = v \\
(\ell \leftarrow w, \sigma) \xrightarrow{\epsilon}_h ((), \sigma : \text{HEAP}[\ell \leftarrow w], \epsilon) & \text{if } \sigma.\text{HEAP}(\ell) = v \\
(\mathbf{CmpXchg}(\ell, w_1, w_2), \sigma) \xrightarrow{\epsilon}_h ((v, \mathbf{true}), \sigma : \text{HEAP}[\ell \leftarrow w_2], \epsilon) & \text{if } \sigma.\text{HEAP}(\ell) = v \text{ and } v \cong w_1 \\
(\mathbf{CmpXchg}(\ell, w_1, w_2), \sigma) \xrightarrow{\epsilon}_h ((v, \mathbf{false}), \sigma, \epsilon) & \text{if } \sigma.\text{HEAP}(\ell) = v \text{ and } v \not\cong w_1 \\
(\mathbf{Xchg}(\ell, w)) \xrightarrow{\epsilon}_h (v, \sigma : \text{HEAP}[\ell \leftarrow w], \epsilon) & \text{if } \sigma.\text{HEAP}(\ell) = v \\
(\mathbf{FAA}(\ell, z_2)) \xrightarrow{\epsilon}_h (z_1, \sigma : \text{HEAP}[\ell \leftarrow z_1 + z_2], \epsilon) & \text{if } \sigma.\text{HEAP}(\ell) = z_1
\end{array}$$

Special reductions

$$\begin{array}{ll}
(\mathbf{fork} \{e\}, \sigma) \xrightarrow{\epsilon}_h ((), \sigma, e) & \\
(\mathbf{newproph}, \sigma) \xrightarrow{\epsilon}_h (p, \sigma : \text{PROPHS} \uplus \{p\}, \epsilon) & \text{if } p \notin \sigma.\text{PROPHS}
\end{array}$$

$$\frac{(e, \sigma) \xrightarrow{\vec{\kappa}}_h (v, \sigma', \vec{e}')}{(\mathbf{resolve\ with\ } e \text{ at } p \text{ to } w, \sigma) \xrightarrow{\vec{\kappa} + [(p, (v, w))]}_h (v, \sigma', \vec{e}')}$$

Figure 2: HeapLang impure reduction rules.

Here, $\sigma : \text{HEAP}$ denotes σ with the `HEAP` field updated as indicated. $[\ell, \ell + z)$ in the `Alloc` rule indicates that we update all locations in this (left-closed, right-open) interval.

Value (in)equality

$$v \cong w, v \not\cong w$$

$$\begin{array}{l}
v \cong w \triangleq v = w \wedge (\text{val_compare_safe}(v) \vee \text{val_compare_safe}(w)) \\
v \not\cong w \triangleq v \neq w \wedge (\text{val_compare_safe}(v) \vee \text{val_compare_safe}(w)) \\
\text{lit_compare_safe}(z) \quad \text{lit_compare_safe}(\mathbf{true}) \quad \text{lit_compare_safe}(\mathbf{false}) \\
\text{lit_compare_safe}(\ell) \quad \text{lit_compare_safe}() \\
\frac{\text{lit_compare_safe}(v)}{\text{val_compare_safe}(v)} \quad \frac{\text{lit_compare_safe}(v)}{\text{val_compare_safe}(\mathbf{inl}(v))} \quad \frac{\text{lit_compare_safe}(v)}{\text{val_compare_safe}(\mathbf{inr}(v))}
\end{array}$$

Figure 3: HeapLang value comparison judgment.

12.2 Syntactic sugar

We recover many of the common language operations as syntactic sugar.

$$\begin{aligned}\lambda x. e &\triangleq \mathbf{rec_}(x) = e \\ \mathbf{let } x = e \mathbf{ in } e' &\triangleq (\lambda x. e')(e) \\ e; e' &\triangleq \mathbf{let } _ = e \mathbf{ in } e' \\ \mathbf{skip} &\triangleq (); () \\ e_1 \ \&\& \ e_2 &\triangleq \mathbf{if } e_1 \mathbf{ then } e_2 \mathbf{ else false} \\ e_1 \ || \ e_2 &\triangleq \mathbf{if } e_1 \mathbf{ then true else } e_2 \\ \mathbf{match } e \mathbf{ with inl}(x) \Rightarrow e_1 \ | \ \mathbf{inr}(y) \Rightarrow e_2 \mathbf{ end} &\triangleq \mathbf{match } e \mathbf{ with inl} \Rightarrow \lambda x. e_1 \ | \ \mathbf{inr} \Rightarrow \lambda y. e_2 \mathbf{ end} \\ \mathbf{??}e &\triangleq \mathbf{Alloc}(1, e) \\ \mathbf{CAS}(e_1, e_2, e_3) &\triangleq \mathbf{snd}(\mathbf{CmpXchg}(e_1, e_2, e_3)) \\ \mathbf{resolve } e_1 \mathbf{ to } e_2 &\triangleq \mathbf{resolve with skip at } e_1 \mathbf{ to } e_2\end{aligned}$$

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] Ralf Jung et al. “Iris from the Ground Up”. In: *Submitted to JFP* (2017).
- [7] Ralf Jung et al. “The future is ours: prophecy variables in separation logic”. In: *PACMPL* 4.POPL (2020), 45:1–45:32. DOI: [10.1145/3371113](https://doi.org/10.1145/3371113).
- [8] Martin H. Löb. “Solution of a Problem of Leon Henkin”. In: *The Journal of Symbolic Logic* 20.2 (1955), pp. 115–118. ISSN: 00224812. URL: <http://www.jstor.org/stable/2266895>.
- [9] 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.