

Destabilizing Iris (Appendix)

SIMON SPIES, MPI-SWS, Germany

NIKLAS MÜCK, MPI-SWS, Germany

HAOYI ZENG, Saarland University, Germany

MICHAEL SAMMLER, ETH Zurich, Switzerland and ISTA, Austria

ANDREA LATTUADA, MPI-SWS, Germany

PETER MÜLLER, ETH Zurich, Switzerland

DEREK DREYER, MPI-SWS, Germany

CONTENTS

Contents	1
A Resource Algebras	2
A.1 Resource Algebra Combinators	4
A.2 Modeling the Unstable Points-To	5
B Base Logic	8
C Many-Sorted First-Order Logic	10
C.1 General First-Order Logic	10
C.2 Interpreted Theories	12
D From First-Order Logic to Iris	13
D.1 Connecting <i>hProp</i> -Assertions and Meta-Level Assertions	14
D.2 Constructing a First-Order Model for λ_{dyn}	14
D.3 Putting Everything Together with the Ambient Heap	16
E Case Studies	17
E.1 Best of Both Worlds (#1)	17
E.2 Leveraging SMT-Solvers (#2)	21
E.3 Iterative Verification (#3)	24
E.4 Polymorphic Hashmap (#4)	26
E.5 Iris Examples (#5)	29
References	30

Resource Algebra

An *partially stable resource algebra* $A = (M, \cdot, \varepsilon, |_{\text{core}}, \mathcal{V}, |_{\text{st}}, |_{\text{unst}})$ is a *unital resource algebra*:

$$\begin{aligned}
 a \cdot (b \cdot c) &= (a \cdot b) \cdot c & a \cdot b &= b \cdot a \\
 |a|_{\text{core}} \cdot a &= a & ||a|_{\text{core}}|_{\text{core}} &= |a|_{\text{core}} & a \leqslant b &\Rightarrow |a|_{\text{core}} \leqslant |b|_{\text{core}} & (a \cdot b) \in \mathcal{V} &\Rightarrow a \in \mathcal{V} \\
 \varepsilon &\in \mathcal{V} & \varepsilon \cdot a &= a & |\varepsilon|_{\text{core}} &= \varepsilon
 \end{aligned}$$

where

$$\begin{aligned}
 a \leqslant b &\triangleq \exists c. a \cdot c = b \\
 a \rightsquigarrow B &\triangleq \forall a_f. (a \cdot a_f) \in \mathcal{V} \Rightarrow \exists b \in B. (b \cdot a_f) \in \mathcal{V}
 \end{aligned}$$

with Stable and Unstable Elements

$$\begin{aligned}
 ||a|_{\text{st}}|_{\text{st}} &= |a|_{\text{st}} & (\text{RA-STABLE-IDEMP}) \\
 |a \cdot b|_{\text{st}} &= |a|_{\text{st}} \cdot |b|_{\text{st}} & (\text{RA-STABLE-DISTR}) \\
 ||a|_{\text{core}}|_{\text{st}} &= |a|_{\text{core}} & (\text{RA-CORE-STABLE}) \\
 |a|_{\text{st}} \cdot |a|_{\text{unst}} &= a & (\text{RA-DECOMPOSE}) \\
 |a|_{\text{unst}} \cdot a &= a & (\text{RA-UNSTABLE-DUPL}) \\
 ||a|_{\text{unst}}|_{\text{unst}} &= |a|_{\text{unst}} & (\text{RA-UNSTABLE-IDEMP}) \\
 a \leqslant b &\Rightarrow |a|_{\text{unst}} \leqslant |b|_{\text{unst}} & (\text{RA-UNSTABLE-MONO}) \\
 |a \cdot b|_{\text{unst}}|_{\text{unst}} &= ||a|_{\text{unst}} \cdot b|_{\text{unst}} & (\text{RA-UNSTABLE-FLIP}) \\
 a \in \mathcal{V} &\Rightarrow |a|_{\text{unst}} \cdot b \in \mathcal{V} \Rightarrow a \cdot |b|_{\text{unst}} \in \mathcal{V} & (\text{RA-UNSTABLE-EXTENSION})
 \end{aligned}$$

where

$$a \rightsquigarrow_{\text{st}} B \triangleq \forall a_f. (a \cdot a_f) \in \mathcal{V} \Rightarrow \exists b \in B. (b \cdot |a_f|_{\text{st}}) \in \mathcal{V}$$

Fig. 1. The partially stable resource algebra, additions in violet.

A Resource Algebras

As a foundation for Daenerys, we take the resource algebra model of Iris [5, 6] and generalize it with *unstable resources*. To do so, we extend the definition of a resource algebra (see Fig. 1) with two projections, $|a|_{\text{st}}$ and $|a|_{\text{unst}}$.

The stable projection. Resources can be decomposed in a stable- and unstable-part (**RA-DECOMPOSE**). The stable-projection $|a|_{\text{st}}$ defines which part of the resource will be preserved by the corresponding updates ($\rightsquigarrow_{\text{st}}$):

$$\frac{b \rightsquigarrow_{\text{st}} b'}{b \cdot a \rightsquigarrow_{\text{st}} b' \cdot |a|_{\text{st}}}$$

As discussed in the paper, the stable projection corresponds to the frame modality \boxplus in the logic.

The stable projection is idempotent (**RA-STABLE-IDEMP**), distributes over composition (**RA-STABLE-DISTR**), and preserves the core (**RA-CORE-STABLE**). Idempotence ensures that the frame modality is idempotent (**FRAME-IDEMP** in Fig. 6). Distributivity ensures that we can combine the separating

$$\begin{aligned}
\text{ExFrac}(X) \quad \ni \quad & \varepsilon \mid \text{stab}_q(x) \mid \text{unst}(x) \\
\\
& \text{stab}_q(x) \in \mathcal{V} \Leftrightarrow q \leq 1 \quad \text{unst}(x) \in \mathcal{V} \quad |\text{stab}_q(x)|_{\text{st}} = \text{stab}_q(x) \quad |\text{unst}(x)|_{\text{st}} = \varepsilon \\
& |\text{stab}_q(x)|_{\text{unst}} = \text{unst}(x) \quad |\text{unst}(x)|_{\text{unst}} = \text{unst}(x) \quad \text{stab}_{q_1}(x) \cdot \text{stab}_{q_2}(x) = \text{stab}_{q_1+q_2}(x) \\
& \text{stab}_q(x) = \text{stab}_q(x) \cdot \text{unst}(x) \quad \text{unst}(x) = \text{unst}(x) \cdot \text{unst}(x) \\
& \text{stab}_{q_1}(x) \cdot \text{stab}_{q_2}(y) \in \mathcal{V} \Leftrightarrow (q_1 + q_2 \leq 1 \wedge x = y) \quad \text{unst}(x) \cdot \text{unst}(y) \in \mathcal{V} \Rightarrow x = y \\
& \text{stab}_1(x) \rightsquigarrow_{\text{st}} \text{stab}_1(y) \quad \text{UnstableComplete}(\text{stab}_1(x))
\end{aligned}$$

Fig. 2. The resource algebra $\text{ExFrac}(X)$.

conjunction of two frame modalities (**FRAME-SEP** in Fig. 6). The preservation of the core ensures that persistent assertions are always frameable (**FRAME-PERS** in Fig. 6).

Note that in resource algebras from standard Iris, all elements are stable (*i.e.*, $|a|_{\text{st}} = a$) such that $(\rightsquigarrow_{\text{st}})$ and (\rightsquigarrow) coincide. More specifically, one can trivially turn regular Iris unital resource algebras into partially stable resource algebra by picking $|a|_{\text{st}} = a$ and $|a|_{\text{unst}} = \varepsilon$.

The unstable projection. The unstable-projection $|a|_{\text{unst}}$ defines a potentially temporary part of the resource. As discussed in the paper, the unstable projection corresponds to the unstable modality \ast in the logic. It is idempotent (**RA-UNSTABLE-IDEMP**) and monotone (**RA-UNSTABLE-MONO**). Idempotence ensures that the unstable modality is idempotent (**UNSTABLE-IDEMP** in Fig. 6). Monotonicity ensures (together with **RA-UNSTABLE-DUPL**) that the unstable modality distributes over separating conjunction (**UNSTABLE-SEP** in Fig. 6).

In addition, it imposes three key axioms, **RA-UNSTABLE-DUPL**, **RA-UNSTABLE-FLIP**, and **RA-UNSTABLE-EXTENSION**. The axiom **RA-UNSTABLE-DUPL** guarantees that the unstable part of a resource is *duplicable*, meaning we can create as many copies of it as we would like. It is the basis for duplicating unstable assertions (**UNSTABLE-DUPL** in Fig. 6). The axiom **RA-UNSTABLE-FLIP** allows us to flip the unstable projection inside an unstable projection. It is a weaker form of distributivity of the unstable projection (*i.e.*, weaker than $|a \cdot b|_{\text{unst}} = |a|_{\text{unst}} \cdot |b|_{\text{unst}}$) that still suffices to prove the implication rule for unstable propositions **UNSTABLE-IMPL** in Fig. 6. The axiom **RA-UNSTABLE-EXTENSION** ensures that the unstable part $|a|_{\text{unst}}$ is a “complete snapshot” of a in the sense that there cannot be an element b that is valid with $|a|_{\text{unst}}$ but whose unstable part $|b|_{\text{unst}}$ is not valid with a . This rule is needed—in addition to **RA-UNSTABLE-FLIP**—to prove that the implication between two unstable assertions is unstable (*i.e.*, **UNSTABLE-IMPL** in Fig. 6). The issue is that Iris’s implication $P \Rightarrow Q$ is upclosed with respect to larger resources and, to prove **UNSTABLE-IMPL**, we end up needing to extend the validity predicate for a resource. However, so long as the unstable projection provides a “complete snapshot” as given by **RA-UNSTABLE-EXTENSION**, we can show that the extension does not suddenly break validity.

$$\begin{array}{l}
\text{Sil}(M) \quad \ni \quad \varepsilon \mid \text{orig}(a) \mid \text{sil}(a) \\
\\
\text{orig}(a) \in \mathcal{V} \Leftrightarrow a \in \mathcal{V} \qquad \qquad \qquad \text{sil}(a) \in \mathcal{V} \Leftrightarrow a \in \mathcal{V} \\
|\text{orig}(a)|_{\text{st}} = \text{orig}(a) \qquad \qquad \qquad |\text{sil}(a)|_{\text{st}} = \varepsilon \\
|\text{orig}(a)|_{\text{unst}} = \text{sil}(a) \qquad \qquad \qquad |\text{sil}(a)|_{\text{unst}} = \text{sil}(a) \\
\\
\text{orig}(a) = \text{orig}(a) \cdot \text{sil}(a) \qquad \text{orig}(a \cdot b) = \text{orig}(a) \cdot \text{orig}(b) \qquad \text{sil}(a) = \text{sil}(a) \cdot \text{sil}(a) \\
\\
\text{sil}(a) \cdot \text{sil}(b) \in \mathcal{V} \Leftrightarrow (\exists c. c \in \mathcal{V} \wedge a \leq c \wedge b \leq c) \qquad a \leq b \Leftrightarrow \text{sil}(a) \leq \text{sil}(b) \\
\\
(a \rightsquigarrow b) \Rightarrow (\text{orig}(a) \rightsquigarrow_{\text{st}} \text{orig}(b)) \\
\\
(a, b) \rightsquigarrow_{\text{L}} (a', b') \Rightarrow (\text{orig}(a), \text{orig}(b)) \rightsquigarrow_{\text{st}}^{\text{L}} (\text{orig}(a'), \text{orig}(b')) \qquad \text{UnstableComplete}(\text{orig}(a))
\end{array}$$

Fig. 3. The resource algebra $\text{Sil}(M)$.

A.1 Resource Algebra Combinators

We discuss several interesting instances of our partially stable resource algebra.

Exclusive with fractions. The exclusive resource algebra with fractions $\text{ExFrac}(X)$ generalizes the $\text{Ex}(\mathbb{N})$ resource algebra discussed in the paper by enabling fractional ownership of the exclusive element. We have two elements $\text{stab}_q(x)$ (the counter part of $\text{ex}(x)$) and $\text{unst}(x)$ (the counter part of $\text{tmp}(x)$). Their key rules are depicted in Fig. 2. The element $\text{stab}_q(x)$ carries fractional ownership, and it is always stable. We can update it for fraction 1. The element $\text{unst}(x)$ is an unstable temporary copy.

Authoritative Resource Algebra. We define an version of the authoritative resource algebra $\text{Auth}(A)$ for partially stable resource algebras, written $\text{PSAuth}(A)$. Compared to the standard Iris authoritative resource algebra, this version has an additional, unstable element $\ominus x$, which is an unstable copy of the full element $\bullet x$. That is,

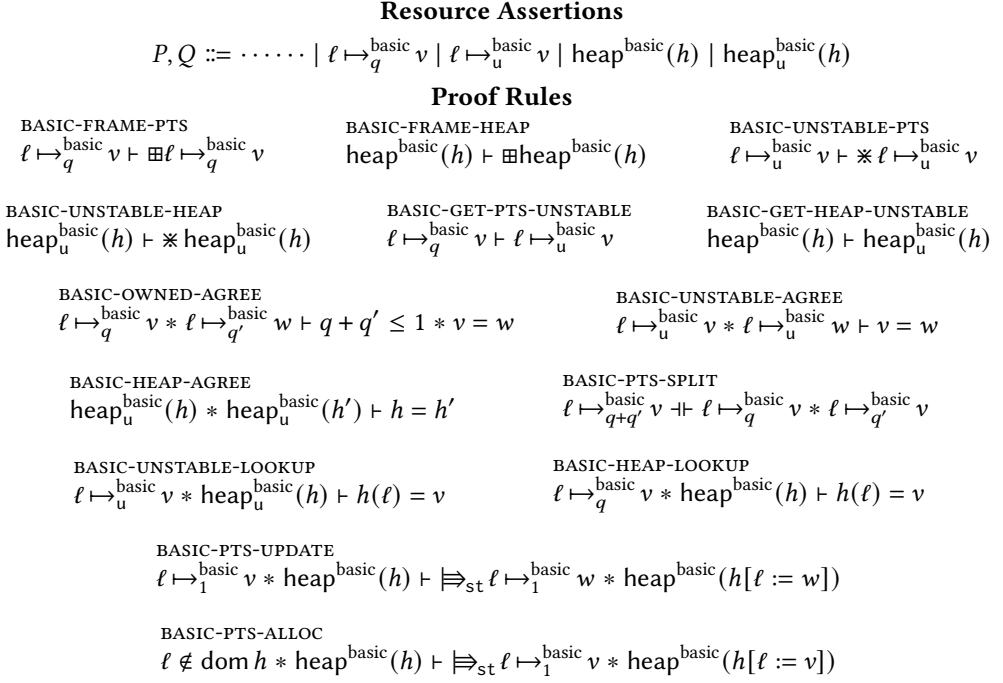
$$\bullet x = \bullet x \cdot \ominus x \qquad \ominus x = \ominus x \cdot \ominus x \qquad |\bullet x|_{\text{unst}} = \ominus x \qquad \ominus x \cdot \ominus y \in \mathcal{V} \Rightarrow x = y$$

To fulfill the property **RA-UNSTABLE-EXTENSION**, we require that full elements $\bullet x$ are maximal in the sense that

$$\text{UnstableComplete}(x) \triangleq \forall y. y \in \mathcal{V} \Rightarrow |y|_{\text{unst}} \leq |x|_{\text{unst}} \Rightarrow y \leq x$$

Silhouette Resource Algebra. We define the $\text{Sil}(M)$ resource algebra, which completes a standard unital resource algebra M to one with unstable resources.¹ The resource algebra is depicted in Fig. 3. Its two key elements are $\text{orig}(a)$ which simply embeds elements $a \in M$, and $\text{sil}(a)$, which is an unowned copy of a . When we have two unowned copies $\text{sil}(a) \cdot \text{sil}(b)$, we do not add them together. Instead, there must be some c that includes a and b . The $\text{orig}(a)$ elements are unstable complete, so we can use them in an PSAuth -construction.

¹This resource algebra supports only non-step-indexed (i.e., “discrete”) resource algebras as instantiations of M .

Fig. 4. The basic theory of heaps obtained by choosing the resource algebra $B\text{Heap}$

A.2 Modeling the Unstable Points-To

Equipped with these resource algebra combinators, we can discuss the model of the unstable points-to assertion $\ell \mapsto_u v$. We first discuss a basic version and then how, in Daenerys, we get the full version that supports Iris’s discardable fractions (for additional “persistent points-to” $\ell \mapsto_{\square} v$).

A basic heap. We obtain the basic heap by choosing the resource algebra $S\text{Heap} \triangleq \text{PSAuth}(Loc \xrightarrow{\text{fin}} \text{ExFrac}(Val))$. In Iris, the corresponding heap theory would consist only of two assertions: $\text{heap}^{\text{basic}}(h)$ to keep track of the full heap, and the points-to assertions $\ell \mapsto_q^{\text{basic}} v$ for $0 < q \leq 1$ to keep track of individual locations in h . Here, we extend it with two new, unstable assertions: $\text{heap}_u^{\text{basic}}(h)$, an unstable temporary copy of $\text{heap}(h)$ and $\ell \mapsto_u^{\text{basic}} v$, an unstable temporary copy of the points-to.

Let us first discuss their rules, depicted in Fig. 4, before we discuss the definition of the connectives below. As usual, points-to must agree on their values (**BASIC-OWNED-AGREE**), we can split-and combine points-to assertions as needed (**BASIC-PTS-SPLIT**), and a points-to assertion for ℓ determines the value of ℓ in h (**BASIC-HEAP-LOOKUP**). We can update an entry in the heap with a *stable update* (**BASIC-PTS-UPDATE**) and allocate a new entry (**BASIC-PTS-ALLOC**). (In Iris, these rules would be standard frame preserving updates.) The points-to for $q > 0$ and the heap are frameable (**BASIC-FRAME-PTS** and **BASIC-FRAME-HEAP**), and their unstable counterparts are unstable (**BASIC-UNSTABLE-PTS** and **BASIC-UNSTABLE-HEAP**). We can always get an unstable copy from the owned versions (**BASIC-GET-PTS-UNSTABLE** and **BASIC-GET-HEAP-UNSTABLE**), and the unstable versions agree: the unstable points-to agree on the value in the heap (**BASIC-UNSTABLE-AGREE**), two unstable heaps agree on the entire heap (**BASIC-HEAP-AGREE**), and we can look up a location in the unstable heap (**BASIC-UNSTABLE-LOOKUP**).

To obtain this resource theory, we define

$$\begin{aligned} \ell \mapsto_q^{\text{basic}} v &\triangleq [\circ \{ \ell \mapsto \text{stab}_q(v) \}]^{\gamma_{\text{heap}}} & \ell \mapsto_u^{\text{basic}} v &\triangleq [\circ \{ \ell \mapsto \text{unst}(v) \}]^{\gamma_{\text{heap}}} \\ \text{heap}^{\text{basic}}(h) &\triangleq [\bullet \{ \ell \mapsto \text{stab}_1(v) \mid \ell \mapsto v \in h \}]^{\gamma_{\text{heap}}} \\ \text{heap}_u^{\text{basic}}(h) &\triangleq [\ominus \{ \ell \mapsto \text{stab}_1(v) \mid \ell \mapsto v \in h \}]^{\gamma_{\text{heap}}} \end{aligned}$$

Here, we write $[\bar{a}]^\gamma$ for Iris’s ownership connective, a version of $\text{Own}(a)$ shown in the paper for supporting multiple resource algebras at the same time. The variable γ is the name of the concrete “instance” of the resource; for heaps, we can fix a global name such as γ_{heap} .

This resource algebra is a variation of a standard technique for constructing a resource algebra for heaps in Iris. For the regular points-to $\ell \mapsto_q^{\text{basic}} v$, we use a fragment \circ containing a singleton map that maps ℓ to the *stable resource* $\text{stab}_q(v)$ (from the resource algebra in Fig. 2). Analogously, for the unstable points-to $\ell \mapsto_u^{\text{basic}} v$, we use a fragment \circ containing a singleton map that maps ℓ to the *unstable resource* $\text{unst}(v)$ (from the resource algebra in Fig. 2). For the full heap $\text{heap}^{\text{basic}}(h)$, we use the authoritative element \bullet containing the entire heap (where every element has been allocated with fraction 1). For the unstable copy $\text{heap}_u^{\text{basic}}(h)$, we use the new unstable authoritative element \ominus of our *PSAuth* resource algebra.

Supporting discardable fractions. As mentioned above, the version of the heap theory that we use in Daenerys goes beyond the basic heaps by supporting discardable fractions (*i.e.*, the “persistent points-tos” $\ell \mapsto_\square v$). The corresponding theory is depicted in Fig. 5.

Similar to the standard Iris points-tos with discardable fractions, its construction in terms of resource algebra combinators is somewhat more involved. It can be found in the accompanying Rocq development. We end up using the resource algebra

$$\text{Heap}(\text{Loc}, \text{Val}) \triangleq \text{PSView}((\text{Loc} \xrightarrow{\text{fin}} \text{Val}), \text{Sil}(\text{Loc} \xrightarrow{\text{fin}} \text{DFrac} \times \text{Ag}(\text{Val})))$$

where *PSView* is a version of Iris’s view resource algebra *extended with unstable elements* (generalizing the *PSAuth*-resource algebra from §A.1) and *Sil* is the silhouette resource algebra (from §A.1).

Resource Assertions

$$P, Q ::= \dots \mid \ell \mapsto_{dq} v \mid \ell \mapsto_u v \mid \text{heap}(h) \mid \text{heap}_u(h) \quad dq \in \{\square\} \cup \mathbb{Q}^+$$

Proof Rules

FRAME-PTS

$$\ell \mapsto_{dq} v \vdash \boxplus \ell \mapsto_{dq} v$$

PERSISTENT-PTS

$$\ell \mapsto_{\square} v \vdash \square(\ell \mapsto_{\square} v)$$

FRAME-HEAP

$$\text{heap}(h) \vdash \boxplus \text{heap}(h)$$

UNSTABLE-PTS

$$\ell \mapsto_u v \vdash \ast \ell \mapsto_u v$$

UNSTABLE-HEAP

$$\text{heap}_u(h) \vdash \ast \text{heap}_u(h)$$

GET-PTS-UNSTABLE

$$\ell \mapsto_{dq} v \vdash \ell \mapsto_u v$$

GET-HEAP-UNSTABLE

$$\text{heap}(h) \vdash \text{heap}_u(h)$$

OWNED-AGREE

$$\ell \mapsto_{dq} v \ast \ell \mapsto_{dq'} w \vdash (dq \cdot dq') \in \mathcal{V} \ast v = w$$

UNSTABLE-AGREE

$$\ell \mapsto_u v \ast \ell \mapsto_u w \vdash v = w$$

HEAP-AGREE

$$\text{heap}_u(h) \ast \text{heap}_u(h') \vdash h = h'$$

PTS-SPLIT

$$\ell \mapsto_{q_1+q_2} v \dashv \vdash \ell \mapsto_{q_1} v \ast \ell \mapsto_{q_2} v$$

UNSTABLE-LOOKUP

$$\ell \mapsto_u v \ast \text{heap}_u(h) \vdash h(\ell) = v$$

HEAP-LOOKUP

$$\ell \mapsto_{dq} v \ast \text{heap}(h) \vdash h(\ell) = v$$

PTS-UPDATE

$$\ell \mapsto_1 v \ast \text{heap}(h) \vdash \boxRightarrow_{\text{st}} \ell \mapsto_1 w \ast \text{heap}(h[\ell := w])$$

PTS-ALLOC

$$\ell \notin \text{dom } h \ast \text{heap}(h) \vdash \boxRightarrow_{\text{st}} \ell \mapsto_1 v \ast \text{heap}(h[\ell := v])$$

PTS-DEALLOC

$$\text{heap}(h) \ast \ell \mapsto_1 v \vdash \boxRightarrow_{\text{st}} \text{heap}(h[\ell := /])$$

PTS-PERSIST

$$\ell \mapsto_{dq} v \vdash \boxRightarrow_{\text{st}} \ell \mapsto_{\square} v$$

Fig. 5. The full theory of heaps with discardable fractions obtained by choosing the resource algebra *Heap*.

$\frac{\text{UPD-MONO} \quad P \vdash Q}{\models_{\text{st}} P \vdash \models_{\text{st}} Q}$		$\text{UPD-BUPD} \quad \models P \vdash \models_{\text{st}} P$	$\text{UPD-INTRO} \quad P \vdash \models_{\text{st}} P$	$\text{UPD-TRANS} \quad \models_{\text{st}} \models_{\text{st}} P \vdash \models_{\text{st}} P$
$\text{UPD-FRAME} \quad (\boxplus P) * (\models_{\text{st}} Q) \models_{\text{st}} (\boxplus P) * Q$		$\frac{\text{UPD-OWN} \quad a \rightsquigarrow_{\text{st}} b}{\text{Own}(a) \vdash \models_{\text{st}} \text{Own}(b)}$		$\frac{\text{FRAME-MONO} \quad P \vdash Q}{\boxplus P \vdash \boxplus Q}$
		$\text{FRAME-ELIM} \quad \boxplus P \vdash P$		
$\text{FRAME-IDEMP} \quad \boxplus P \vdash \boxplus \boxplus P$	$\text{FRAME-EXISTS} \quad \boxplus (\exists x. P x) \dashv\vdash \exists x. (\boxplus P x)$	$\text{FRAME-ALL} \quad \boxplus (\forall x. P x) \dashv\vdash \forall x. (\boxplus P x)$		$\text{FRAME-LATER} \quad \boxplus (\triangleright P) \dashv\vdash \triangleright (\boxplus P)$
$\text{FRAME-SEP} \quad (\boxplus P) * (\boxplus Q) \vdash \boxplus (P * Q)$		$\text{FRAME-PERS} \quad \Box P \vdash \boxplus P$	$\text{FRAME-OWN} \quad \text{Own}(a) \vdash \boxplus \text{Own}(a _{\text{st}})$	$\frac{\text{UNSTABLE-MONO} \quad P \vdash Q}{* P \vdash * Q}$
$\text{UNSTABLE-ELIM} \quad * P \vdash P$	$\text{UNSTABLE-IDEMP} \quad * P \vdash * * P$	$\text{UNSTABLE-EXISTS} \quad * (\exists x. P x) \dashv\vdash \exists x. (* P x)$	$\text{UNSTABLE-ALL} \quad * (\forall x. P x) \dashv\vdash \forall x. (* P x)$	
$\text{UNSTABLE-LATER} \quad * (\triangleright P) \dashv\vdash \triangleright (* P)$		$\text{UNSTABLE-SEP} \quad (* P) * (* Q) \vdash * (P * Q)$		$\text{UNSTABLE-OWN} \quad \text{Own}(a) \vdash * \text{Own}(a _{\text{unst}})$
		$\text{UNSTABLE-IMPL} \quad * P \Rightarrow * Q \vdash * (* P \Rightarrow Q)$		$\text{UNSTABLE-DUPL} \quad (* P) \wedge Q \dashv\vdash (* P) * Q$

Fig. 6. Base logic rules for the update modality $\models_{\text{st}} P$, frame modality $\boxplus P$, and unstable modality $* P$

B Base Logic

The new modalities \models_{st} , \boxplus , and $*$ satisfy the rules in Fig. 6.

Let us start with the *stable update* \models_{st} . It satisfies several standard rules for Iris updates that make it compositional (a monad to be precise): it is monotone (**UPD-MONO**), we can always introduce it (**UPD-INTRO**), and we can always compose updates via transitivity (**UPD-TRANS**). For the last two rules to hold (which are important for making updates practically usable and integrating them into the weakest precondition), it is important that the underlying update $a \rightsquigarrow_{\text{st}} a$ is *reflexive* (for **UPD-INTRO**) and *transitive* (for **UPD-TRANS**). Besides these rules, we can always turn a stable update on resources $a \rightsquigarrow_{\text{st}} a$ into one on the assertion level (**UPD-OWN**). The standard update \models is included in \models_{st} (**UPD-BUPD**), which means existing Iris ghost theories (e.g., the theory $\text{Auth}(\mathbb{N}, \text{max})$) can still be used as is—especially considering that \models_{st} is used in the definition of the weakest precondition. Lastly, the rule **UPD-FRAME** allows us to frame assertions R into a stable update *if we can put them underneath a frame modality* \boxplus .

Let us now turn to the *frame modality* \boxplus . The frame modality is a co-monad (e.g., like Iris’s persistency modality): it is monotone (**FRAME-MONO**), we can always eliminate it (**FRAME-ELIM**), and it is idempotent (**FRAME-IDEMP**). The rule **FRAME-ELIM** means we never have to worry about “getting rid” of a frame modality (from our assumptions). The rule **FRAME-IDEMP** means we can always “add another” frame modality, which can be useful for proof rules that remove a frame modality. Moreover, the frame modality distributes over existential quantification (**FRAME-EXISTS**), universal quantification (**FRAME-ALL**), and the later modality (**FRAME-LATER**). It *does not* distribute over the separating conjunction in both directions. More specifically, we can combine two frameable assertions (**FRAME-SEP**), but the opposite direction does not hold (i.e., turning $\boxplus (P * Q)$ into $(\boxplus P) *$

($\boxplus Q$). For example, it would be unsound to turn $\boxplus(\ell \mapsto v * \ell \mapsto_u v)$ into $(\boxplus \ell \mapsto v) * (\boxplus \ell \mapsto_u v)$, because $(\boxplus \ell \mapsto v)$ could be used to update ℓ invalidating $(\boxplus \ell \mapsto_u v)$. Furthermore, to make sure that persistent propositions $\Box P$ keep their usual meaning (*i.e.*, once created, they persist across updates and state changes), we make sure that persistent assertions are frameable (**FRAME-PERS**). Finally, when we own a resource a , we always frameably own its stable part (**FRAME-OWN**).

Let us now turn to the *unstable modality* $*$. Like the frame modality, it is a co-monad: it is monotone (**UNSTABLE-MONO**), we can always eliminate it (**UNSTABLE-ELIM**), and it is idempotent (**UNSTABLE-IDEMP**). Thus, also for the unstable modality, we never have to worry about “getting rid” of it (from our assumptions), and we can always “add another”. Moreover, the unstable modality distributes over existential quantification (**UNSTABLE-EXISTS**), universal quantification (**UNSTABLE-ALL**), and the later modality (**UNSTABLE-LATER**). Importantly, the unstable modality makes separating conjunction and ordinary conjunction coincide (**UNSTABLE-DUPL**). This rule ensures that we do not have to give up ownership when we prove an unstable assertions. From it, we can derive, *e.g.*, $*P \vdash *P * *P$. The unstable modality *does* distribute over separation conjunction in both directions: **UNSTABLE-SEP** gives one direction, and the other direction can be derived from **UNSTABLE-MONO** and **UNSTABLE-DUPL**. Furthermore, we have the rule **UNSTABLE-IMPL**, which effectively means that the implication between two unstable assertions $*P$ and $*Q$ is itself an unstable proposition. We use it to justify adding $F_1 \Rightarrow F_2$ to the *hProp*-fragment. Finally, when we own a resource a , we always unstably own its unstable part (**UNSTABLE-OWN**).

Terms	t, s	$::=$	$x \mid f \vec{t}$	$(f \in \mathcal{C})$
Formulas	π, χ	$::=$	$\text{True} \mid \text{False} \mid t \dot{=} s \mid p \vec{t} \mid \pi_1 \dot{\wedge} \pi_2 \mid \pi_1 \dot{\vee} \pi_2$ $\mid \pi_1 \dot{\Rightarrow} \pi_2 \mid \exists x : S. \pi \mid \forall x : S. \pi$	$(p \in \mathcal{P})$

Typing

TYPE-VAR $\frac{x : S \in \Sigma}{\Sigma \vdash_{\text{term}} x : S}$	TYPE-FUNC $\frac{\vdash_{\text{func}} f : \vec{S} \rightarrow T \quad \Sigma \vdash_{\text{term}} \vec{t} : \vec{S}}{\Sigma \vdash_{\text{term}} f \vec{t} : T}$	TYPE-TRUE $\Sigma \vdash_{\text{form}} \text{True}$	TYPE-FALSE $\Sigma \vdash_{\text{form}} \text{False}$
TYPE-AND $\frac{\Sigma \vdash_{\text{form}} \pi_1 \quad \Sigma \vdash_{\text{form}} \pi_2}{\Sigma \vdash_{\text{form}} \pi_1 \dot{\wedge} \pi_2}$	TYPE-OR $\frac{\Sigma \vdash_{\text{form}} \pi_1 \quad \Sigma \vdash_{\text{form}} \pi_2}{\Sigma \vdash_{\text{form}} \pi_1 \dot{\vee} \pi_2}$	TYPE-IMPL $\frac{\Sigma \vdash_{\text{form}} \pi_1 \quad \Sigma \vdash_{\text{form}} \pi_2}{\Sigma \vdash_{\text{form}} \pi_1 \dot{\Rightarrow} \pi_2}$	
TYPE-ALL $\frac{\Sigma, x : S \vdash_{\text{form}} \pi}{\Sigma \vdash_{\text{form}} \dot{\forall} x : S. \pi}$	TYPE-EXISTS $\frac{\Sigma, x : S \vdash_{\text{form}} \pi}{\Sigma \vdash_{\text{form}} \dot{\exists} x : S. \pi}$	TYPE-EQ $\frac{\Sigma \vdash_{\text{term}} t : S \quad \Sigma \vdash_{\text{term}} s : S}{\Sigma \vdash_{\text{form}} t \dot{=} s}$	
TYPE-PRED $\frac{\vdash_{\text{pred}} p : \vec{S} \rightarrow \text{Prop} \quad \Sigma \vdash_{\text{term}} \vec{t} : \vec{S}}{\Sigma \vdash_{\text{form}} p \vec{t}}$			

Fig. 7. Terms and formulas of first-order logic.

C Many-Sorted First-Order Logic

In this section, we describe the first-order logic that we consider.

C.1 General First-Order Logic

First-order terms and formulas. We define first-order terms t and formulas π in Fig. 7. We define them parametrically over a signature $\Omega = (\mathcal{S}, \mathcal{C}, \mathcal{P}, \vdash_{\text{func}}, \vdash_{\text{pred}})$, where \mathcal{S} is a collection of first-order sort symbols; \mathcal{C} a collection of first-order function symbols; \mathcal{P} a collection of first-order predicate symbols; \vdash_{func} assigns argument sorts and a result sort to every function symbol $f \in \mathcal{C}$; and \vdash_{pred} assigns argument sorts to every predicate symbol $p \in \mathcal{P}$. We write $\vdash_{\text{func}} f : \vec{S} \rightarrow T \in \mathcal{C}$ to mean $f \in \mathcal{C}$ and $\vdash_{\text{func}} f : \vec{S} \rightarrow T$. We write $\vdash_{\text{pred}} p : \vec{S} \rightarrow \text{Prop} \in \mathcal{P}$ to mean $p \in \mathcal{P}$ and $\vdash_{\text{pred}} p : \vec{S} \rightarrow \text{Prop}$.

Definition C.1. We say a signature $\Omega_1 = (\mathcal{S}_1, \mathcal{C}_1, \mathcal{P}_1, \vdash_{\text{func}_1}, \vdash_{\text{pred}_1})$ extends a signature $\Omega_2 = (\mathcal{S}_2, \mathcal{C}_2, \mathcal{P}_2, \vdash_{\text{func}_2}, \vdash_{\text{pred}_2})$, written $\Omega_1 \supseteq_{\text{sig}} \Omega_2$, iff. (1) $\mathcal{S}_1 \supseteq \mathcal{S}_2$, (2) $\mathcal{C}_1 \supseteq \mathcal{C}_2$, and (3) $\mathcal{P}_1 \supseteq \mathcal{P}_2$ and the typing judgments agree:

$$\vdash_{\text{pred}_2} p : \vec{S} \rightarrow \text{Prop} \Rightarrow \vdash_{\text{pred}_1} p : \vec{S} \rightarrow \text{Prop} \quad \vdash_{\text{func}_2} f : \vec{S} \rightarrow T \Rightarrow \vdash_{\text{func}_1} f : \vec{S} \rightarrow T$$

$$\begin{aligned}
\llbracket x \rrbracket_{\text{term}}^M(\epsilon) &= \epsilon(x) \\
\llbracket f \tilde{t} \rrbracket_{\text{term}}^M(\epsilon) &= \llbracket f \rrbracket_{\text{func}}(\llbracket \tilde{t} \rrbracket_{\text{term}}^M(\epsilon)) \\
\llbracket \text{True} \rrbracket_{\text{form}}^M(\epsilon) &= \text{True} \\
\llbracket \text{False} \rrbracket_{\text{form}}^M(\epsilon) &= \text{False} \\
\llbracket \pi_1 \wedge \pi_2 \rrbracket_{\text{form}}^M(\epsilon) &= \llbracket \pi_1 \rrbracket_{\text{form}}^M(\epsilon) \wedge \llbracket \pi_2 \rrbracket_{\text{form}}^M(\epsilon) \\
\llbracket \pi_1 \dot{\vee} \pi_2 \rrbracket_{\text{form}}^M(\epsilon) &= \llbracket \pi_1 \rrbracket_{\text{form}}^M(\epsilon) \vee \llbracket \pi_2 \rrbracket_{\text{form}}^M(\epsilon) \\
\llbracket \pi_1 \dot{\Rightarrow} \pi_2 \rrbracket_{\text{form}}^M(\epsilon) &= \llbracket \pi_1 \rrbracket_{\text{form}}^M(\epsilon) \Rightarrow \llbracket \pi_2 \rrbracket_{\text{form}}^M(\epsilon) \\
\llbracket \dot{\forall} x : S. \pi \rrbracket_{\text{form}}^M(\epsilon) &= \forall d : D_S. \llbracket \pi \rrbracket_{\text{form}}^M(\epsilon, x \mapsto d) \\
\llbracket \dot{\exists} x : S. \pi \rrbracket_{\text{form}}^M(\epsilon) &= \exists d : D_S. \llbracket \pi \rrbracket_{\text{form}}^M(\epsilon, x \mapsto d) \\
\llbracket t \dot{=} s \rrbracket_{\text{form}}^M(\epsilon) &= \llbracket t \rrbracket_{\text{term}}^M(\epsilon) = \llbracket s \rrbracket_{\text{term}}^M(\epsilon) \\
\llbracket p \tilde{t} \rrbracket_{\text{form}}^M(\epsilon) &= \llbracket p \rrbracket_{\text{pred}}(\llbracket \tilde{t} \rrbracket_{\text{term}}^M(\epsilon))
\end{aligned}$$

Fig. 8. The first-order Tarski semantics for a given model $M = (D_{_}, \llbracket _ \rrbracket_{\text{func}}, \llbracket _ \rrbracket_{\text{pred}})$

Tarski semantics. We define the semantics of first-order terms and formulas, a standard Tarski-semantics, in Fig. 8. To define the semantics, we assume a model $M = (D_{_}, \llbracket _ \rrbracket_{\text{func}}, \llbracket _ \rrbracket_{\text{pred}})$ where D maps each $S \in \mathcal{S}$ to a non-empty set D_S , the *domain* for S ; $\llbracket _ \rrbracket_{\text{func}}$ maps each function $\vdash_{\text{func}} f : \vec{S} \rightarrow T \in \mathcal{C}$ to a meta-level function $\llbracket f \rrbracket_{\text{func}} : \vec{D}_S \rightarrow D_T$; and $\llbracket _ \rrbracket_{\text{pred}}$ maps each predicate $\vdash_{\text{pred}} p : \vec{S} \rightarrow \text{Prop} \in \mathcal{P}$ to a meta-level predicate $\llbracket p \rrbracket_{\text{pred}} : \vec{D}_S \rightarrow \text{Prop}$.

Definition C.2. We say a well-formed formula $\vdash_{\text{form}} \pi$ holds in a model $M = (D_{_}, \llbracket _ \rrbracket_{\text{func}}, \llbracket _ \rrbracket_{\text{pred}})$, written $M \models \pi$, iff. $\llbracket \pi \rrbracket_{\text{form}}^M(\emptyset)$ holds.

Note the distinction between the meta-level logic and the object language: The terms and formulas in Fig. 7 are just syntax. We give them meaning via the Tarski-semantics in Fig. 8, which interprets every symbol using the corresponding connective from the ambient meta-logic.

Definition C.3. Let $\Omega_1 = (\mathcal{S}_1, \mathcal{C}_1, \mathcal{P}_1, \vdash_{\text{func}1}, \vdash_{\text{pred}1})$ and $\Omega_2 = (\mathcal{S}_2, \mathcal{C}_2, \mathcal{P}_2, \vdash_{\text{func}2}, \vdash_{\text{pred}2})$ such that $\Omega_1 \sqsupset_{\text{sig}} \Omega_2$. We say a model $M^1 = (D_{_}^1, \llbracket _ \rrbracket_{\text{func}}^1, \llbracket _ \rrbracket_{\text{pred}}^1)$ for signature Ω_1 extends a model $M^2 = (D_{_}^2, \llbracket _ \rrbracket_{\text{func}}^2, \llbracket _ \rrbracket_{\text{pred}}^2)$ for signature Ω_2 , written $M^1 \sqsupset_{\text{model}} M^2$, iff.

- (1) for every $S \in \mathcal{S}_2$, there is a function $i_S : D_S^2 \rightarrow D_S^1$,
- (2) for every $S \in \mathcal{S}_2$, there is a function $r_S : D_S^1 \rightarrow D_S^2$,
- (3) for every $S \in \mathcal{S}_2$, i_S and r_S form a bijection, i.e., $\forall (a \in D_S^2), (b \in D_S^1). i_S(a) = b \text{ iff. } a = r_S(b)$,
- (4) for every $\vdash_{\text{func}} f : \vec{S} \rightarrow T \in \mathcal{C}_2$, we have $\forall \vec{a} \in \vec{D}_S. r_T(\llbracket f \rrbracket_{\text{func}}^1(i_{\vec{S}}(\vec{a}))) = \llbracket f \rrbracket_{\text{func}}^2(\vec{a})$,
- (5) for every $\vdash_{\text{pred}2} p : \vec{S} \rightarrow \text{Prop} \in \mathcal{P}_2$, we have $\forall \vec{a} \in \vec{D}_S. \llbracket p \rrbracket_{\text{pred}}^1(i_{\vec{S}}(\vec{a})) \text{ iff. } \llbracket p \rrbracket_{\text{pred}}^2(\vec{a})$.

The model extension $M^1 \sqsupset_{\text{model}} M^2$ ensures that, for the part described by Ω_2 , the domains in are in bijection, the function interpretations are the same (up to the bijection), and the predicates are equivalent (up to the bijection).

Sort Interpretations				
$D_{\text{unit}} \triangleq \{()\}$	$D_{\text{bool}} \triangleq \mathbb{B}$	$D_{\text{int}} \triangleq \mathbb{Z}$	$D_{\text{bv } n} \triangleq \{m \in \mathbb{Z} \mid 0 \leq m < 2^n\}$	\dots
Function Interpretations				
$\llbracket () \rrbracket_{\text{func}}() \triangleq ()$	$\llbracket \text{true} \rrbracket_{\text{func}}() \triangleq \text{true}$	$\llbracket \text{false} \rrbracket_{\text{func}}() \triangleq \text{true}$	$\llbracket n \rrbracket_{\text{func}}() \triangleq n$	
$\llbracket +_{\text{int}} \rrbracket_{\text{func}}(n, m) \triangleq n + m$	$\llbracket \cdot_{\text{int}} \rrbracket_{\text{func}}(n, m) \triangleq n \cdot m$	$\llbracket -_{\text{int}} \rrbracket_{\text{func}}(n, m) \triangleq n - m$		
$\llbracket \text{neg} \rrbracket_{\text{func}}(a) \triangleq \text{if } a \text{ then false else true}$	$\llbracket \text{xor}_{\text{bv } n} \rrbracket_{\text{func}}(u_1, u_2) \triangleq u_1 \oplus u_2$			
$\llbracket ==_{\text{bv } n} \rrbracket_{\text{func}}(u_1, u_2) \triangleq \text{if } u_1 = u_2 \text{ then true else false}$	$\llbracket \text{if}_S \rrbracket_{\text{func}}(a, d_1, d_2) \triangleq \text{if } a \text{ then } d_1 \text{ else } d_2$			
\dots				
Predicate Interpretations				
$\llbracket \leq_{\text{int}} \rrbracket_{\text{pred}}(n, m) \triangleq n \leq m$	$\llbracket <_{\text{int}} \rrbracket_{\text{pred}}(n, m) \triangleq n < m$	\dots		

Fig. 9. The base model M_{base} .

C.2 Interpreted Theories

An SMT-solver has built-in knowledge for several mathematical theories such as integers and bitvectors. To be sound, we have to make sure that we agree with the SMT-solver on these theories. To this end, we define the following *base signature* Ω_{base} (and *base model* M_{base} below):

Sorts	$S, T ::=$	$\text{unit} \mid \text{bool} \mid \text{int} \mid \text{bv } n \mid \dots$
Functions	$f, g ::=$	$() \mid \text{true} \mid \text{false} \mid n \mid +_{\text{int}} \mid \cdot_{\text{int}} \mid -_{\text{int}} \mid \text{neg} \mid \text{xor}_{\text{bv } n} \mid ==_{\text{bv } n} \mid \text{if}_S \mid \dots$
Predicates	$p, q ::=$	$\leq_{\text{int}} \mid <_{\text{int}} \mid \dots$

The signature introduces several standard base sorts (such as integers int and bit vectors $\text{bv } n$), along with canonical functions and predicates on them (e.g., addition on integers, or comparison on integers, etc.). We then define a base model M_{base} , a standard interpretation for these sorts, functions, and predicates, depicted in Fig. 9.

With the base model in hand, we can then define the main validity judgment $\models \pi$:

Definition C.4. Let Ω_{smt} be a signature extending the base signature Ω_{base} , i.e., $\Omega_{\text{smt}} \supseteq_{\text{sig}} \Omega_{\text{base}}$.

$$\models \pi \triangleq (\vdash_{\text{form}} \pi) \Rightarrow \forall (M : \Omega_{\text{smt}}). M \sqsupseteq_{\text{model}} M_{\text{base}} \Rightarrow M \models \pi$$

That is, π holds true in *all models extending the base model* M_{base} . Since the Tarski semantics is only defined for well-formed formulas, we assume the formula is well-formed in this definition.

$$\begin{aligned}
\langle x \rangle_T^Y &= \gamma(x) & \langle f \tilde{t} \rangle_T^Y &= \langle f \rangle_C \langle \tilde{t} \rangle_T^Y & \langle \text{True} \rangle_F^Y &\triangleq \text{True} & \langle \text{False} \rangle_F^Y &\triangleq \text{False} \\
\langle t \doteq_S s \rangle_F^Y &\triangleq \langle t \rangle_T^Y \equiv \langle s \rangle_T^Y & \langle p \tilde{t} \rangle_F^Y &\triangleq \langle p \rangle_P (\langle \tilde{t} \rangle_T^Y) & \langle \pi_1 \wedge \pi_2 \rangle_F^Y &\triangleq \langle \pi_1 \rangle_F^Y \wedge \langle \pi_2 \rangle_F^Y \\
\langle \pi_1 \dot{\vee} \pi_2 \rangle_F^Y &\triangleq \langle \pi_1 \rangle_F^Y \vee \langle \pi_2 \rangle_F^Y & \langle \pi_1 \Rightarrow \pi_2 \rangle_F^Y &\triangleq \langle \pi_1 \rangle_F^Y \Rightarrow \langle \pi_2 \rangle_F^Y \\
\langle \forall x : S. \pi \rangle_F^Y &\triangleq \forall v. v \in \mathcal{V}[\![\langle S \rangle_S]\!] \Rightarrow \langle \pi \rangle_F^{Y, x \mapsto v} & \langle \exists x : S. \pi \rangle_F^Y &\triangleq \exists v. v \in \mathcal{V}[\![\langle S \rangle_S]\!] \wedge \langle \pi \rangle_F^{Y, x \mapsto v}
\end{aligned}$$

Fig. 10. Translations of terms to expressions $\langle t \rangle_T^Y$ and formulas to $hProp$ -assertions $\langle \pi \rangle_F^Y$, given a translation of sorts to types $\langle S \rangle_S$, function constants to values $\langle f \rangle_C$, and predicates to $hProp$ -predicates $\langle p \rangle_P$.

Sort Translations

$$\langle \text{unit} \rangle_S \triangleq \text{unit} \quad \langle \text{bool} \rangle_S \triangleq \text{bool} \quad \langle \text{int} \rangle_S \triangleq \text{int} \quad \langle \text{bv } n \rangle_S \triangleq \text{bv } n \quad \dots$$

Function Translations

$$\begin{aligned}
\langle () \rangle_C &\triangleq () & \langle \text{true} \rangle_C &\triangleq \text{true} & \langle \text{false} \rangle_C &\triangleq \text{true} & \langle n \rangle_C &\triangleq n & \langle +_{\text{int}} \rangle_C &\triangleq \lambda(x, y). x + y \\
\langle \cdot_{\text{int}} \rangle_C &\triangleq \lambda(x, y). x * y & \langle -_{\text{int}} \rangle_C &\triangleq \lambda(x, y). x - y & \langle \text{neg} \rangle_C &\triangleq \lambda x. \sim x \\
\langle \text{xor}_{\text{bv } n} \rangle_C &\triangleq \lambda(x, y). x \text{ xor } y & \langle ==_{\text{bv } n} \rangle_C &\triangleq \lambda(x, y). x == y \\
\langle \text{if}_S \rangle_C &\triangleq \lambda(x, y_1, y_2). \text{if } x \text{ then } y_1 \text{ else } y_2 & & & & & \dots
\end{aligned}$$

Predicate Translations

$$\begin{aligned}
\langle \leq_{\text{int}} \rangle_P(v_1, v_2) &\triangleq \exists n_1, n_2. v_1 = n_1 \wedge v_2 = n_2 \wedge n_1 \leq n_2 \\
\langle <_{\text{int}} \rangle_P(v_1, v_2) &\triangleq \exists n_1, n_2. v_1 = n_1 \wedge v_2 = n_2 \wedge n_1 < n_2 \quad \dots
\end{aligned}$$

Fig. 11. Translations of sorts to semantic types $\langle S \rangle_S$, function constants to λ_{dyn} -values $\langle f \rangle_C$, and predicates to $hProp$ -predicates $\langle p \rangle_P$.

D From First-Order Logic to Iris

The translation from first-order logic to $hProp$ -assertions $\langle \pi \rangle_F^Y$ is depicted in Fig. 10. We complete it with translations for sorts, functions, and predicates, depicted in Fig. 11. Note that the translation maps (1) *sorts* S to *semantic types* τ in our logical relation, (2) *functions* f to λ_{dyn} -*values* v in our programming language, and (3) *predicates* Φ to $hProp$ -*predicates* $F(v)$ in the fragment of almost-pure assertions.² The translation can be extended, as needed, with uninterpreted sorts (e.g., mapping `buffer` to `buf bv 64`) or uninterpreted functions (e.g., mapping `checksum` to `checksum`).

We will now use it to show the following correspondence:

Let π be well-formed. If $\models \pi$ holds, then $(\langle \pi \rangle_F^0 \Rightarrow \text{wp } e \{v. Q(v)\}) \vdash \text{wp } e \{v. Q(v)\}$ holds in Iris.

As an aside, note that we assume that π is well-formed (i.e., $\vdash_{\text{form}} \pi$ according to the typing judgment in Fig. 7) here. Well-formedness ensures that the semantics defined in Fig. 8 is meaningful (e.g., we do not apply a function that expects Boolean arguments to an integer). To ease the presentation, we have glossed over well-formedness in the paper.

²For better usability, in Rocq, the translation is actually aware of, e.g., binary operations such that $\langle 41 +_{\text{int}} 1 \rangle_T$ is translated to the λ_{dyn} -expression $41 + 1$ instead of the λ_{dyn} -expression $(\lambda(x, y). x + y) (41, 1)$. We gloss over this detail in the following.

$$\begin{aligned}
\llbracket \phi \rrbracket(h) &\triangleq \phi & \llbracket e \Downarrow v \rrbracket(h) &\triangleq (e, h) \rightsquigarrow_{\text{det}}^* (v, h) & \llbracket \ell \mapsto_u v \rrbracket(h) &\triangleq h(\ell) = v \\
\llbracket F \wedge G \rrbracket(h) &\triangleq \llbracket F \rrbracket(h) \wedge \llbracket G \rrbracket(h) & \llbracket F \vee G \rrbracket(h) &\triangleq \llbracket F \rrbracket(h) \vee \llbracket G \rrbracket(h) \\
\llbracket F \Rightarrow G \rrbracket(h) &\triangleq \llbracket F \rrbracket(h) \Rightarrow \llbracket G \rrbracket(h) & \llbracket \forall x. F x \rrbracket(h) &\triangleq \forall x. \llbracket F x \rrbracket(h) & \llbracket \exists x. F x \rrbracket(h) &\triangleq \exists x. \llbracket F x \rrbracket(h)
\end{aligned}$$

Fig. 12. Translation $\llbracket _ \rrbracket(_) : hProp \rightarrow Heap \rightarrow Prop$ between $hProp$ -assertions F and pure assertions $\llbracket F \rrbracket(h)$.

Proof overview. We will proceed in three steps: First, we connect $hProp$ -assertions to meta-level, pure assertions (§D.1). Then, we connect the Tarski-semantics from Fig. 8 to $hProp$ -assertions (§D.2). Finally, we put everything together to derive the correspondence above in Theorem D.8 (§D.3).

D.1 Connecting $hProp$ -Assertions and Meta-Level Assertions

As the first step, we define a translation—for a fixed heap h —from $hProp$ -assertions to Rocq assertions. The translation, $\llbracket _ \rrbracket(h)$, is depicted in Fig. 12. We then show the following correspondence:

LEMMA D.1. *For any heap h , we have $\text{heap}_u(h) * (*_{\ell \mapsto v \in h} \ell \mapsto_u v) \vdash (F \Leftrightarrow \llbracket F \rrbracket(h))$.*

PROOF. By induction on F .³ For most connectives the proof is straightforward. (For implication $F_1 \Rightarrow F_2$, it is important that we prove an equivalence instead of merely an implication between the two.) An interesting case is evaluation $e \Downarrow v$. Here, we must show:

$$\text{heap}_u(h) * (*_{\ell \mapsto v \in h} \ell \mapsto_u v) \vdash (e \Downarrow v \Leftrightarrow \llbracket e \Downarrow v \rrbracket(h))$$

Considering the definition of $e \Downarrow v \triangleq \exists h'. (e, h') \rightsquigarrow_{\text{det}}^* (v, h') * (*_{\ell \mapsto w \in h'} \ell \mapsto_u w)$, the forward direction “ \Rightarrow ” follows from heap_u including all $\ell \mapsto_u v$ assertions concealed under the definition of $e \Downarrow v$. The backward direction “ \Rightarrow ” follows, because we have unstable points-to assertions $\ell \mapsto_u v$ in our assumption for every entry in h . \square

D.2 Constructing a First-Order Model for λ_{dyn}

Equipped with the translation $\llbracket _ \rrbracket(h)$, we can now construct a first-order logic model using λ_{dyn} . We will use the model to instantiate $\models \pi$ (in §D.3).

The model M_{dyn}^h . We construct the model for an arbitrary, but fixed heap h , and write M_{dyn}^h to indicate the dependence on the heap h . First, we must instantiate the domains:

$$D_S \triangleq \{v \mid \llbracket v \rrbracket \in \mathcal{V}[\llbracket \langle S \rangle_S \rrbracket](h)\}$$

We instantiate the domains with sets of values v , where—using our translation $\llbracket _ \rrbracket(h)$ —we know the value v is in the logical relation at type $\langle S \rangle_S$. Since first-order domains must be non-empty, we only do so for sorts where the corresponding value relation $\mathcal{V}[\llbracket \langle S \rangle_S \rrbracket]$ is non-empty.

For the predicates p , we simply choose:

$$\llbracket p \rrbracket_{\text{pred}} \triangleq \llbracket \langle p \rangle_p \rrbracket(h)$$

The first-order functions f are more interesting. We must instantiate the functions with meta-level functions for the interpretation $\llbracket _ \rrbracket_{\text{func}}$. To this end, we prove the following lemma, which turns program functions f into meta-level functions f :

³To obtain impredicative quantification in $hProp$, in Rocq, we actually make Lemma D.1 a defining property of $hProp$ -assertions, together with $F \vdash * F$, and then prove that the $hProp$ -connectives (e.g., $e \Downarrow v$, $\ell \mapsto_u v$, $\forall x. F(x)$, ...) satisfy these properties.

LEMMA D.2. *Let $\vdash f \in \mathcal{V}[\tau_1 \rightarrow \tau_2]$. For any h , there exists a function $f : \text{Val} \rightarrow \text{Val}$ such that for all values v where $\lfloor v \in \tau_1 \rfloor(h)$, we have $(f v, h) \rightsquigarrow_{\text{det}}^* (f(v), h)$ and $\lfloor f(v) \in \tau_2 \rfloor(h)$.*

PROOF. Given the definition of $\lfloor _ \rfloor(h)$ and the equivalence in Lemma D.1, this lemma follows by first proving $\lfloor f \in \mathcal{V}[\tau_1 \rightarrow \tau_2] \rfloor(h)$, which unfolds to

$$\forall v. \lfloor v \in \tau_1 \rfloor(h) \Rightarrow \exists w. (f v, h) \rightsquigarrow_{\text{det}}^* (w, h) \wedge \lfloor w \in \tau_2 \rfloor(h)$$

and then applying the axiom of choice to extract w . \square

To define $\llbracket f \rrbracket_{\text{func}}$, we apply Lemma D.2 (generalized to arbitrary function arities) to $f \triangleq \langle f \rangle_C$.

Proving the extension. Having defined the model, we must then show that it is an extension of the base model M_{base} :

LEMMA D.3. $M_{\text{dyn}}^h \sqsupseteq_{\text{model}} M_{\text{base}}$ for any heap h

PROOF SKETCH. Recall that the signature of M_{base} is Ω_{base} . For the sorts S in Ω_{base} , it is straightforward to construct a bijection between the domains D_S of the base model M_{base} and well-typed values in our semantic types $\mathcal{V}[\langle S \rangle_S]$. The remaining equivalences of $M_{\text{dyn}}^h \sqsupseteq_{\text{model}} M_{\text{base}}$ then follow by case analysis on our interpreted sorts, functions, and predicates. Here, we must make sure that the implementation of, e.g., multiplication $\langle \cdot_{\text{int}} \rangle_C \triangleq \lambda xy. x * y$ (a value in λ_{dyn}) matches (up to the bijection) the semantics of multiplication, i.e., $\llbracket \cdot_{\text{int}} \rrbracket_{\text{func}}(n, m) \triangleq n \cdot m$ (in ordinary mathematics on integers). \square

Proving the equivalence. Equipped with the model, we can now show that formulas π interpreted in the first-order logic model M_{dyn}^h are equivalent to their translation $\langle \pi \rangle_F$ at the meta-level if we apply $\lfloor _ \rfloor(h)$. To state the relationship formally, we connect a first-order variable mapping ϵ (to elements of D_S in M_{dyn}^h) to a variable substitution γ (from variables to values) with:

$$\text{agree}(\Sigma, \gamma, \epsilon) \triangleq \forall x : S \in \Sigma. \exists v. \gamma(x) = v \wedge \epsilon(x) = v$$

We first show for terms that their semantics corresponds to the evaluation in λ_{dyn} :

LEMMA D.4. *Let $\Sigma \vdash_{\text{term}} t : S$. If $\text{agree}(\Sigma, \gamma, \epsilon)$, then for every heap h ,*

$$\llbracket t \rrbracket_{\text{term}}^{M_{\text{dyn}}^h}(\epsilon) = v \quad \text{iff.} \quad (\langle t \rangle_F^\gamma, h) \rightsquigarrow_{\text{det}}^* (v, h)$$

PROOF. By induction on the typing $\Sigma \vdash_{\text{term}} t : S$. \square

Then we use this result to show the desired equivalence:

LEMMA D.5. *Let $\Sigma \vdash_{\text{form}} \pi$. If $\text{agree}(\Sigma, \gamma, \epsilon)$, then for every heap h ,*

$$\llbracket \pi \rrbracket_{\text{form}}^{M_{\text{dyn}}^h}(\epsilon) \quad \text{iff.} \quad \lfloor \langle \pi \rangle_F^\gamma \rfloor(h)$$

PROOF. By induction on the typing $\Sigma \vdash_{\text{form}} \pi$, using Lemma D.4. \square

As a corollary, we can connect valid formulas and the meta-level version of their translation:

COROLLARY D.6. *Let π be well-formed. If $\models \pi$, then $\lfloor \langle \pi \rangle_F^\emptyset \rfloor(h)$ for any heap h .*

PROOF. Fix a heap h . We instantiate $\models \pi$ with M_{dyn}^h , which extends the M_{base} (Lemma D.3). Thus, we obtain $\llbracket \pi \rrbracket_{\text{form}}^{M_{\text{dyn}}^h}(\emptyset)$. With Lemma D.5 it follows that $\lfloor \langle \pi \rangle_F^\emptyset \rfloor(h)$ holds. \square

D.3 Putting Everything Together with the Ambient Heap

To put everything together, as the last step, we need access to the ambient heap. That is, note that the equivalence in [Lemma D.1](#) imposes a precondition, namely $\text{heap}_u(h) * (*_{\ell \mapsto v \in h} \ell \mapsto_u v)$, requiring effectively unstable resources for the entire current heap. We call this precondition the *ambient heap* and define:

$$\begin{aligned} \text{ambient_heap_is}(h) &\triangleq \text{heap}_u(h) * (*_{\ell \mapsto v \in h} \ell \mapsto_u v) \\ \text{ambient_heap} &\triangleq \exists h. \text{ambient_heap_is}(h) \end{aligned}$$

Let us now first show that given an ambient heap, we can bring π from first-order logic to Iris:

LEMMA D.7. *Let π be well-formed. If $\models \pi$ holds, then $\text{ambient_heap} \vdash \langle \pi \rangle_F^\emptyset$*

PROOF. Let $\models \pi$. It suffices to prove for all heaps h

$$\text{ambient_heap_is}(h) \vdash \langle \pi \rangle_F^\emptyset$$

By [Lemma D.1](#), it then suffices to prove

$$(\langle \pi \rangle_F^\emptyset \Leftrightarrow \lfloor \langle \pi \rangle_F^\emptyset \rfloor(h)) \vdash \langle \pi \rangle_F^\emptyset$$

By [Lemma D.6](#) and $\models \pi$, we know $\lfloor \langle \pi \rangle_F^\emptyset \rfloor(h)$ holds true for any heap h . Thus, we can use the equivalence $\langle \pi \rangle_F^\emptyset \Leftrightarrow \lfloor \langle \pi \rangle_F^\emptyset \rfloor(h)$ to establish $\langle \pi \rangle_F^\emptyset$. \square

We can obtain an ambient heap at any point when proving a weakest precondition for λ_{dyn} :

$$\begin{array}{c} \text{GET-AMBIENT-HEAP} \\ (\forall h. \text{ambient_heap_is } h \multimap \text{wp } e \{v. Q(v)\}) \vdash \text{wp } e \{v. Q(v)\} \end{array}$$

Thus, putting both together, we obtain:

THEOREM D.8.

Let π be well-formed. If $\models \pi$ holds, then $(\langle \pi \rangle_F^\emptyset \Rightarrow \text{wp } e \{v. Q(v)\}) \vdash \text{wp } e \{v. Q(v)\}$ holds in Iris.

PROOF. We acquire the ambient heap via [GET-AMBIENT-HEAP](#). Then, we apply [Lemma D.7](#) on the ambient heap to obtain $\langle \pi \rangle_F^\emptyset$. Finally, we use $\langle \pi \rangle_F^\emptyset \Rightarrow \text{wp } e \{v. Q(v)\}$ to obtain $\text{wp } e \{v. Q(v)\}$. \square

Group	Case Study	Iris [5]	ViperCore [3]	Viper [7]	Daenerys
#1	Channel Library	●	○	○	●
	Checksum Exchange	○	○	●	●
#2	Popcount 32-bit Integer	●	●	●	●
	Popcount Buffer <i>à la</i> Redis [8]	●	●	●	●
	Priority Bit Map <i>à la</i> RefinedC [9]	●	●	●	●
#3	Iterative Linked-List	○	○	●	●
#4	Polymorphic Hashmap	○	○	●	●
#5	Iris Concurrent Logical Relation [11]	●	○	○	●
	Barrier [4], Reader-Writer Lock, Spinlock	●	○	○	●
Foundational Model		●	●	○	●

Fig. 13. Evaluation of Daenerys. We compare with the other approaches Iris, ViperCore, and Viper-based verifiers and mark whether they support the case study. We write ● for yes, ○ for no, and ● for case studies that *could conceivably be done* but require significant manual effort in Iris.

Implementation

```

chan()  $\triangleq$  ref(None)
recv(c)  $\triangleq$  let v = !c in
  match v with
  | None  $\Rightarrow$  recv(c)
  | Some(l)  $\Rightarrow$  if CAS(c, Some(l), None) then !l else recv(c)
send(c, x)  $\triangleq$  let l = ref(x) in
  if CAS(c, None, Some(l)) then () else send(c, x)

```

Specification

```

{True} chan() {c. ischan(c,  $\Phi$ )}           persistent(ischan(c,  $\Phi$ ))

{ischan(c,  $\Phi$ ) *  $\boxplus \Phi(v)$ } send(c, v) {_. True}   {ischan(c,  $\Phi$ )} recv(c) {v.  $\boxplus \Phi(v)$ }

```

Fig. 14. The channel implementation and specification.

E Case Studies

Below, we describe in more detail the case studies from the main paper, depicted in Fig. 13.

E.1 Best of Both Worlds (#1)

This example consists of two parts: verifying the channel implementation and verifying the actual exchange between client and worker thread. We first discuss the channel implementation and then proceed with how to verify the exchange.

Channel Library. The implementation of the channel library is depicted in Fig. 14. Each channel is represented as a reference to an option. The option is either None if there is currently no value being transferred over the channel, or Some(ℓ) where ℓ is a reference storing the value that is

currently being transferred from the sender to the receiver. (We add the additional indirection via a reference wrapping the value, because `HeapLang` and by extension λ_{dyn} forbid “compare-and-set” operations on references that store larger, composite values such as a pair.)

The three operations for channels are `chan`, `send`, and `recv`.

- (1) The operation `chan` initializes a new channel by creating the channel reference. The reference initially stores `None` (i.e., no value is currently being transferred).
- (2) The operation `send` wraps the value it wants to send x in a reference l , and then tries to store `Some(l)` in the channel reference c . It does so via an atomic “compare-and-set” operation that tests whether the channel is currently empty (i.e., currently storing `None`) and, if so, updates it to `Some(l)`. If the update succeeds, `send` returns, and if it fails, `send` loops to wait for a point where the channel is empty again.
- (3) The operation `recv` reads the contents v of the channel reference c . If they are `None`, then no value is currently being transferred and the `recv` operation loops to wait for a value. If there is some value currently being transferred (wrapped in a reference l) then `recv` tries to update c to `None` with a “compare-and-set” to claim the value. If the update succeeds, then the `recv` can return the contents of l (i.e., the value being transferred). If the update fails, another call to `recv` must have been successful at receiving the value. In this case, `recv` loops to wait for the next value.

To verify the channel operations (depicted in Fig. 14), we use an Iris invariant:

$$\text{ischan}(c, \Phi) \triangleq \exists \ell. c = \ell * \boxed{\exists v. \ell \mapsto v * (v = \text{None} \vee \exists \ell', w. v = \text{Some}(\ell') * \ell' \mapsto w * \boxplus \Phi(w))}^N$$

It captures the two possible cases the channel reference ℓ can be in: either (1) no value is currently being transferred and the reference is storing `None`, or (2) some value w is being transferred, wrapped in some reference ℓ' . In the latter case, the invariant stores that Φ holds for the value being currently transferred. It does so underneath a frame modality to ensure that the ownership being transferred via the channel is stable. (Formally, this modality is needed to make sure that the contents of the invariant are frameable, a side condition of invariants in Daenerys.) With this invariant, verifying the specifications in Fig. 14 is a straightforward Iris proof. Moreover, the representation predicate $\text{ischan}(c, \Phi)$ is persistent—and can be shared freely between threads and different program parts—because invariants are persistent.

What is interesting about the invariant above is that it stores an arbitrary Iris predicate Φ , which can include Hoare-triples (as we will see below) or even other invariants. The reason why such generality is allowed in Iris is that its invariants are *impredicative* [10] (i.e., they that can contain arbitrary Iris propositions). Impredicativity allows us to state a general and modular specification for the channels in Fig. 14 (e.g., for every channel, a new predicate Φ can be chosen). It does, however, also come at a cost: to justify soundness of its invariants, Iris uses step-indexing [1, 2], which adds significant complexity to its underlying model. Fortunately, in Daenerys, we effectively inherit the support for step-indexing from Iris and can, hence, use impredicative invariants to verify the channel implementation.

Checksum Exchange. Let us now turn to the exchange between the worker thread and the client (that uses the channel implementation to communicate between both). Recall the implementation:

```

wrk(i, o)  $\triangleq$  let (p, c) = recv(i) in let b = p() in let s = c(b) in send(o, (b, s)); wrk(i, o)
client()  $\triangleq$  let (i, o) = (chan(), chan()) in fork {wrk(i, o)} ;
    send(i, (produceA, checksumA)); let (b, s) = recv(o) in assert(s == checksumA(b));
    send(i, (produceB, checksumB)); let (b, s) = recv(o) in assert(s == checksumB(b))

```

The worker wrk (1) receives on the input channel i a workload p (i.e., a function that will produce a buffer) and a checksum function c , (2) produces the buffer b by executing p , (3) computes the checksum s of b , (4) sends both b and s back via the output channel o , and (5) repeats the entire process. The client client creates an input and an output channel, spawns the worker thread, and then sends the worker two different workloads: First, it sends produceA with checksum function checksumA , receives the result b and s , and ensures that s matches checksumA of b . Then, it repeats the same process with a *different workload and checksum* implementation.

Let us now discuss how we can verify this implementation given the channel specifications in Fig. 14. We start by introducing the two channel predicates, one for the input channel and one for the output channel:

$$\begin{aligned}\Phi_{\text{inp}}(p, c) &\triangleq \{\text{True}\} p() \{v. \exists b, \vec{u}. v = b * b \mapsto \vec{u}\} * \Box(c \in \mathcal{V}[\text{buf } u64 \rightarrow u64]) * \gamma \models_{1/2}(p, c) \\ \Phi_{\text{outp}}(b, s) &\triangleq \exists \vec{u}. p, c. b \mapsto \vec{u} * c(b) \Downarrow s * \gamma \models_{1/2}(p, c)\end{aligned}$$

On the input channel (Φ_{inp}), the client sends a pair of *two functions* p and c . For p , it sends a Hoare-triple that p will compute a buffer. For c , it sends the fact that c is a well-typed function. In addition, it sends a custom piece of ghost state, a “fractional ghost variable” $\gamma \models_{1/2}(p, c)$, to track which workload the worker is currently working on. The fractional ghost variables $\gamma \models_q z$ behave like fractional points-to assertions $\ell \mapsto_q v$, but are not physically part of the program. As we will see below, the client keeps one half $\gamma \models_{1/2}(p, c)$ and sends the other half to the worker (via Φ_{inp}).

On the output channel (Φ_{outp}), the worker sends a pair of buffer b and checksum s . For b , it sends the ownership of the buffer $b \mapsto \vec{u}$. For s , it sends $c(b) \Downarrow s$, meaning s is the result of computing the checksum c on the buffer b . Along with them, it returns the ghost variable $\gamma \models_{1/2}(p, c)$.

Given these two predicates, the specifications that we have proven are:

$$\{\text{True}\} \text{client}() \{_ . \text{True}\} \quad \{\text{ischan}(i, \Phi_{\text{inp}}) * \text{ischan}(o, \Phi_{\text{outp}})\} \text{wrk}(i, o) \{_ . \text{True}\}$$

The client is safe to execute (i.e., pre- and postcondition True), which means the asserts inside must succeed. The worker thread is safe to execute given an input channel adhering to Φ_{inp} and an output channel adhering to Φ_{outp} .

To illustrate the basic structure of the proof, we give a proof outline of the client in Fig. 15. Initially, we allocate the two channels (using the specification of chan in Fig. 14). We additionally allocate a new ghost variable γ initially storing an arbitrary pair of values. Next, we fork-off the worker thread. To do so, we give it a copy of the representation predicate of the input/output channel $\text{ischan}(i, \Phi_{\text{inp}}) * \text{ischan}(o, \Phi_{\text{outp}})$ and keep another copy to ourselves.

Next, we start with the first checksum-and-produce combination. We update the ghost variable and split it into two halves: $\gamma \models_{1/2}(\text{produceA}, \text{checksumA}) * \gamma \models_{1/2}(\text{produceA}, \text{checksumA})$. Moreover, we prove the following specifications for produceA and checksumA :

$$\{\text{True}\} \text{produceA}() \{v. \exists b, \vec{u}. v = b * b \mapsto \vec{u}\} \quad \models \text{checksumA} : \text{buf } u64 \rightarrow u64$$

Together, we can then send $\Phi_{\text{inp}}(\text{produceA}, \text{checksumA})$ to the worker thread via the input channel. (The assertion $\Phi_{\text{inp}}(\text{produceA}, \text{checksumA})$ is frameable, meaning $\Phi_{\text{inp}}(\text{produceA}, \text{checksumA}) \vdash \boxplus \Phi_{\text{inp}}(\text{produceA}, \text{checksumA})$, so we can cross the frame modality in the specification of send .)

Next, we receive on the output channel the resulting combination of buffer b and checksum s , meaning $\Phi_{\text{outp}}(b, s)$. We have two halves of the ghost variable γ : $\gamma \models_{1/2}(\text{produceA}, \text{checksumA})$ and $\gamma \models_{1/2}(p, c)$ at this point. Just like for a fractional points-to, this means the two must agree (so $p = \text{produceA}$ and $c = \text{checksumA}$). Thus, the client has just received the knowledge of $\text{checksumA}(b) \Downarrow s$. It can use it to justify that the assert succeeds.

Finally, we can combine the ghost variable again into one, and proceed with the second checksum-and-produce combination. It is completely analogous to the first one.

```

{True}
let (i, o) = (chan(), chan()) in
{ischan(i,  $\Phi_{\text{inp}}$ ) * ischan(o,  $\Phi_{\text{outp}}$ ) *  $\gamma \models_1 (\_, \_)$ }
fork {wrk(i, o)};
{ischan(i,  $\Phi_{\text{inp}}$ ) * ischan(o,  $\Phi_{\text{outp}}$ ) *  $\gamma \models_1 (\_, \_)$ }
{ischan(i,  $\Phi_{\text{inp}}$ ) * ischan(o,  $\Phi_{\text{outp}}$ ) *  $\gamma \models_{1/2}(\text{produceA}, \text{checksumA})$  *  $\gamma \models_{1/2}(\text{produceA}, \text{checksumA})$ }
send(i, (produceA, checksumA));
{ischan(i,  $\Phi_{\text{inp}}$ ) * ischan(o,  $\Phi_{\text{outp}}$ ) *  $\gamma \models_{1/2}(\text{produceA}, \text{checksumA})$ }
let (b, s) = recv(o) in
{
  ischan(i,  $\Phi_{\text{inp}}$ ) * ischan(o,  $\Phi_{\text{outp}}$ ) *  $\gamma \models_{1/2}(\text{produceA}, \text{checksumA})$  *
  {  $\boxplus (\exists \vec{u}, p, c. b \mapsto \vec{u} * c(b) \Downarrow s * \gamma \models_{1/2}(p, c))$  }
  ischan(i,  $\Phi_{\text{inp}}$ ) * ischan(o,  $\Phi_{\text{outp}}$ ) *  $\gamma \models_{1/2}(\text{produceA}, \text{checksumA})$  *
  {  $(\exists \vec{u}. b \mapsto \vec{u} * \text{checksumA}(b) \Downarrow s * \gamma \models_{1/2}(\text{produceA}, \text{checksumA}))$  }
}
assert(s == checksumA(b));
{ischan(i,  $\Phi_{\text{inp}}$ ) * ischan(o,  $\Phi_{\text{outp}}$ ) *  $\gamma \models_1 (\_, \_)$ }
{ischan(i,  $\Phi_{\text{inp}}$ ) * ischan(o,  $\Phi_{\text{outp}}$ ) *  $\gamma \models_{1/2}(\text{produceB}, \text{checksumB})$  *  $\gamma \models_{1/2}(\text{produceB}, \text{checksumB})$ }
send(i, (produceB, checksumB));
{ischan(i,  $\Phi_{\text{inp}}$ ) * ischan(o,  $\Phi_{\text{outp}}$ ) *  $\gamma \models_{1/2}(\text{produceB}, \text{checksumB})$ }
let (b, s) = recv(o) in
{
  ischan(i,  $\Phi_{\text{inp}}$ ) * ischan(o,  $\Phi_{\text{outp}}$ ) *  $\gamma \models_{1/2}(\text{produceB}, \text{checksumB})$  *
  {  $\boxplus (\exists \vec{u}, p, c. b \mapsto \vec{u} * c(b) \Downarrow s * \gamma \models_{1/2}(p, c))$  }
  ischan(i,  $\Phi_{\text{inp}}$ ) * ischan(o,  $\Phi_{\text{outp}}$ ) *  $\gamma \models_{1/2}(\text{produceB}, \text{checksumB})$  *
  {  $(\exists \vec{u}. b \mapsto \vec{u} * \text{checksumB}(b) \Downarrow s * \gamma \models_{1/2}(\text{produceB}, \text{checksumB}))$  }
}
assert(s == checksumB(b))
{True}

```

Fig. 15. Proof outline of the client.

Summary. In summary, in this example, we get the best of both worlds: Iris and IDF. We use Iris’s impredicative invariants to verify the channel implementation above, and then we use the evaluation assertion $e \Downarrow v$ of Daenerys to send information about the current result of the checksum function over the output channel. The latter allows us to avoid proving functional correctness of *two different* checksum implementations. Instead, we can simply send the assertion $c(b) \Downarrow s$ from worker to client.

E.2 Leveraging SMT-Solvers (#2)

We consider three different examples:

Popcount 32-bit Integer. For the Popcount 32-bit Integer-example, we consider the following implementation to compute the number of ones in a 32-bit integer:

```
pc32(u)  $\triangleq$  let v = u - (u  $\gg$  1 & 0x55555555) in
    let w = (v & 0x33333333) + ((v  $\gg$  2) & 0x33333333) in
    (((w + (w  $\gg$  4)) & 0x0F0F0F0F) * 0x01010101)  $\gg$  24
```

We show that it implements the following, naïve implementation:

$$\text{ones}(u) \triangleq (u \gg 31) \& 0x1 + \dots + (u \gg 0) \& 0x1$$

More specifically, for a bitvector u where $0 \leq u < 2^{32}$, the specification that we prove is

$$\{\text{True}\} \text{pc32}(u) \{v. v \equiv \text{ones}(u)\}$$

For that, we give the SMT solver an unfolding of the definitions of pc32 and ones and ask the following query:

$$\begin{aligned} \pi_{\text{pc32}} &\triangleq (\forall x. \text{aux1}(x) = x - ((x \gg 1) \& 0x55555555)) \Rightarrow \\ &(\forall x. \text{aux2}(x) = (x \& 0x33333333) + ((x \gg 2) \& 0x33333333)) \Rightarrow \\ &(\forall x. \text{aux3}(x) = ((x + (x \gg 4)) \& 0x0F0F0F0F)) \Rightarrow \\ &(\forall x. \text{aux}(x) = \text{aux3}(\text{aux2}(\text{aux1}(x)))) \Rightarrow \\ &(\forall x. \text{pc32}(x) = (\text{aux}(x) \& 0x01010101) \gg 24) \Rightarrow \\ &(\forall x. \text{ones}(x) = (x \gg 31) \& 0x1 + \dots + (x \gg 0) \& 0x1) \Rightarrow \\ &\forall x. \text{True} \Rightarrow \forall res. \text{pc32}(x) = res \Rightarrow res = \text{ones}(x) \end{aligned}$$

In Rocq we assume $\models \pi_{\text{pc32}}$, instantiate the premises of the implication by manually proving the unfoldings (which is straightforward), and then we automatically derive the Hoare triple from the conclusion of the query.

Popcount Buffer. In the Popcount Buffer example (inspired by a similar implementation in Redis [8]), the implementation counts the ones in a buffer⁴ by considering seven 32-bit integers at a time (in a loop over the buffer). Afterwards, for the remaining 0-6 integers, it looks up the corresponding ones in a table.

```
aux1(x)  $\triangleq$  x - ((x  $\gg$  1) & 0x55555555)
aux2(x)  $\triangleq$  (x & 0x33333333) + ((x  $\gg$  2) & 0x33333333)
aux3(x)  $\triangleq$  ((x + (x  $\gg$  4)) & 0x0F0F0F0F)
aux(x)  $\triangleq$  aux3(aux2(aux1(x)))
```

```
pc7( $x_1, x_v, x_3, x_4, x_5, x_6, x_7$ )  $\triangleq$  ((aux( $x_1$ ) + aux( $x_2$ ) + aux( $x_3$ ) + aux( $x_4$ ) + aux( $x_5$ ) + aux( $x_6$ ) + aux( $x_7$ ))
    * 0x01010101)  $\gg$  24
```

```
bitsinbyte  $\triangleq$  #[0, 1, 1, 2, 1, 2, 2, 3, 1, 2, 2, 3, 2, 3, 3, 4, 1, 2, 2, 3, 2, 3, 3, 4, 2, 3, 3, 4, 3, 4, 4, 5, 1, 2, 2,
    3, 2, 3, 3, 4, 2, 3, 3, 4, 3, 4, 4, 5, 2, 3, 3, 4, 3, 4, 4, 5, 3, 4, 4, 5, 4, 5, 5, 6, 1, 2, 2, 3, 2, 3, 3, 4,
    2, 3, 3, 4, 3, 4, 4, 5, 2, 3, 3, 4, 3, 4, 4, 5, 3, 4, 4, 5, 4, 5, 5, 6, 2, 3, 3, 4, 3, 4, 4, 5, 3, 4, 4, 5, 4,
```

⁴To access the elements of the buffer, we use a function `buf_get(·, buffer()) $\tau \times \text{int} \rightarrow \tau$` , which returns a default value in case the index is out of bounds.

```

5, 5, 6, 3, 4, 4, 5, 4, 5, 5, 6, 4, 5, 5, 6, 5, 6, 6, 7, 1, 2, 2, 3, 2, 3, 3, 4, 2, 3, 3, 4, 3, 4, 4, 5, 2, 3,
3, 4, 3, 4, 4, 5, 3, 4, 4, 5, 4, 5, 5, 6, 2, 3, 3, 4, 3, 4, 4, 5, 3, 4, 4, 5, 4, 5, 5, 6, 3, 4, 4, 5, 4, 5, 5,
6, 4, 5, 5, 6, 5, 6, 6, 7, 2, 3, 3, 4, 3, 4, 4, 5, 3, 4, 4, 5, 4, 5, 5, 6, 3, 4, 4, 5, 4, 5, 5, 6, 4, 5, 5, 6,
5, 6, 6, 7, 3, 4, 4, 5, 4, 5, 5, 6, 4, 5, 5, 6, 5, 6, 6, 7, 4, 5, 5, 6, 5, 6, 6, 7, 5, 6, 6, 7, 6, 7, 7, 8]
lookup(x)  $\triangleq$  bitsinbyte[int((x  $\gg$  0) & 0xFF)]
      + bitsinbyte[int((x  $\gg$  8) & 0xFF)]
      + bitsinbyte[int((x  $\gg$  16) & 0xFF)]
      + bitsinbyte[int((x  $\gg$  32) & 0xFF)]
while p f  $\triangleq$  rec while(f) := if p () then f (); while p f else ()
pc_redis(buf)  $\triangleq$  let count = ref(length(buf)) in
  let i = ref(0) in
  let bits = ref(0) in
  while ( $\lambda\_ . 7 \leq !count$ ) (
    bits := !bits + pc7(
      buf_get(buf, !i + 1),
      buf_get(buf, !i + 2),
      buf_get(buf, !i + 3),
      buf_get(buf, !i + 4),
      buf_get(buf, !i + 5),
      buf_get(buf, !i + 6)
    );
    count := !count - 7;
    i := !i + 7;
  );
  while ( $\lambda\_ . 0 < !count$ ) (
    bits := !bits + lookup(buf_get(buf, !i))
    count := !count - 1;
    i := !i + 1;
  );
  !bits

```

We give unfoldings of all definitions excluding while and pc_redis to the SMT solver and ask whether they imply $pc7(x_1, x_2, x_3, x_4, x_5, x_6, x_7) = ones(x_1) + ones(x_2) + ones(x_3) + ones(x_4) + ones(x_5) + ones(x_6) + ones(x_7)$ and $lookup(x) = ones(x)$. Then we do a manual proof by induction on the value stored in count to prove the following Hoare triple:

$$\{ \text{buffer}(\text{buf}) \}$$

$$\text{pc_redis}(\text{buf})$$

$$\{ v. \text{iter}(0, \text{length}(\text{buf}), 0, \lambda i \text{ bits. bits} + \text{ones}(\text{buf_get}(\text{buf}, i))) \Downarrow v \}$$

Priority Bit Map. Here we verify a version of a bit map implementation previously verified in RefinedC⁵. A bitmap $[0; 256] \mapsto [0, 1]$ containing 256 bits is efficiently represented as a vector of 4 64-bit integers and verified against the naïve meta logic representation as a vector of 256 Booleans. The operations are:

- `init` initializes the bitmap to all 0
- `highest` returns the smallest index of a set bit (*i.e.*, of a bit that is 1)
- `set_prio` sets a bit to 1
- `clear_prio` sets a bit to 0

We consider the following implementation:

```

getBit(x, i)  $\triangleq$  (nth(x, int(i  $\gg$  6))  $\gg$  (i & 63)) & 1

ffs(x)  $\triangleq$    if  x & (1  $\ll$  0)  $\neq$  0 then 1 else
              if  x & (1  $\ll$  1)  $\neq$  0 then 2 else
              ...
              if  x & (1  $\ll$  63)  $\neq$  0 then 64 else
              0

init  $\triangleq$  #[0, 0, 0, 0]

highest(xs)  $\triangleq$  if ffs(nth(xs, 0))  $\neq$  0 then ffs(nth(xs, 0)) + -1 else
                if ffs(nth(xs, 1))  $\neq$  0 then ffs(nth(xs, 1)) + 63 else
                if ffs(nth(xs, 2))  $\neq$  0 then ffs(nth(xs, 2)) + 127 else
                if ffs(nth(xs, 3))  $\neq$  0 then ffs(nth(xs, 3)) + 191 else
                -1

set_prio(xs, p)  $\triangleq$  set(xs, int(p  $\gg$  6), (nth(xs, int(p  $\gg$  6))) | (1  $\ll$  (p & 63)))
clear_prio(xs, p)  $\triangleq$  set(xs, int(p  $\gg$  6), (nth(xs, int(p  $\gg$  6))) & ~(1  $\ll$  (p & 63)))

```

We define a representation predicate to relate the operation to vectors of bool in the specs:

$$\text{rep}(xs, bs) \triangleq \text{len}(xs) = 4 \wedge \text{len}(bs) = 256 \wedge \forall i. 0 \leq i < 256 \Rightarrow \text{nth}(bs, i) = (\text{getBit}(xs, i) \neq 0)$$

Then we prove the following Hoare triples by asking the SMT solver similar queries as above, *i.e.*, whether the unfolding of all definitions implies the first-order version of the Hoare triple. We include the exact SMT queries in the accompanying Rocq development.

$$\begin{aligned}
& \{ \text{True} \} \text{init} \{ xs. \text{rep}(xs, \#[\text{false}, \dots, \text{false}]) \} \\
& \{ \text{rep}(xs, bs) \} \text{highest}(xs) \left\{ p. \begin{array}{l} (p = -1 \wedge \forall i < 256. \text{bs}[\text{int}(i)] = \text{false}) \\ \vee (p \neq -1 \wedge \text{bs}[\text{int}(p)] = \text{true} \wedge \forall i < p. \text{bs}[\text{int}(i)] = \text{false}) \end{array} \right\} \\
& \quad \{ \text{rep}(xs, ys) \} \text{set_prio}(xs, p) \{ xs'. \text{rep}(xs', \text{set}(bv, \text{int}(p), \text{true})) \} \\
& \quad \{ \text{rep}(xs, ys) \} \text{clear_prio}(xs, p) \{ xs'. \text{rep}(xs', \text{set}(bv, \text{int}(p), \text{false})) \}
\end{aligned}$$

⁵The original code can be found here: <https://gitlab.mpi-sws.org/iris/refinedc/-/blob/fecd1e6d396dde55e464a71d62e908d60fe920b9/examples/scheduler/include/fdsched/priority.h>

```

new()  $\triangleq$  None
len(l)  $\triangleq$  match l with None  $\Rightarrow$  0 | Some(x)  $\Rightarrow$  1 + len(snd(!x))
nth(l, n)  $\triangleq$  match l with
  | None  $\Rightarrow$  None
  | Some(x)  $\Rightarrow$  if n == 0 then Some(fst(!l)) else nth(snd(!l), n - 1)
push(l, v)  $\triangleq$  let newl = ref(v, l) in Some(newl)
set(l, i, v)  $\triangleq$  match l with
  | None  $\Rightarrow$  ()
  | Some(x)  $\Rightarrow$  if i == 0 then x := (v, snd(!x)) else set(snd(!x), i - 1, v)

```

Fig. 16. Linked list implementation for the iterative verification.

E.3 Iterative Verification (#3)

We iteratively verify the functions `new`, `set`, and `push` of a standard linked-list implementation, depicted in Fig. 16. Each list is either `None` for the empty list or `Some(ℓ)` where ℓ is a reference to a pair of the head and tail of the list. We define the following abstract predicate for the ownership of a list:

$$\text{list}(l) \triangleq l = \text{None} \vee \exists \ell, n, l'. l = \text{Some}(\ell) * \ell \mapsto (n, l') * \text{list}(l')$$

We verify three different kinds of specifications: (1) with respect to ownership of the list for memory safety, (2) with respect to their effect on the length of the list and, (3) with respect to the contents of the list for functional correctness. The resulting specifications are depicted in Fig. 17.

Old expressions. The specifications use so-called “old expressions” `old { e }`, which refer to the value of e in the precondition. To encode them in Daenerys, we define a new Hoare triple

$$\{P\} e \{v, \text{old}. Q(v, \text{old})\}_{\text{old}}$$

It gets access in the postcondition to `old : Expr \rightarrow Val`, a function from expressions to their resulting values if executed in the heap of the precondition. We define it using the following lemma:

LEMMA E.1. *There exists a function $\text{eval} : \text{Heap} \rightarrow \text{Expr} \rightarrow \text{Val}$ such that if $(e, h) \rightsquigarrow_{\text{det}}^* (v, h)$, then $\text{eval } h \ e = v$.*

PROOF. The proof uses classical reasoning. For any heap h and expression e , classically either $\exists v. (e, h) \rightsquigarrow_{\text{det}}^* (v, h)$ or $\neg(\exists v. (e, h) \rightsquigarrow_{\text{det}}^* (v, h))$. In the first case, we can use the axiom of choice to return the value v . In the second case, it does not matter which value eval returns. \square

Given the `eval` function, we define the Hoare-triple using the *ambient heap* (see §D.3) as:

$$\{P\} e \{v, \text{old}. Q(v, \text{old})\}_{\text{old}} \triangleq \Box(\forall h. \Box P \multimap \text{ambient_heap_is } h \multimap \text{wp } e \{v. \Box Q(v, \text{eval } h)\})$$

Compared to the regular Hoare-triples $\{P\} e \{v. Q(v)\}$, this version additionally assumes an ambient heap h , the current heap at the point where we prove the weakest precondition. In the postcondition, we then use `eval` partially instantiated with the heap h .

For postconditions that do not care about old expressions, this Hoare triple is the same as the regular Hoare-triple $\{P\} e \{v. Q(v)\}$, since the weakest precondition can always get access to the ambient heap (see GET-AMBIENT-HEAP in §D.3).

Ownership Specification

$$\{\text{True}\} \text{new}() \{l, _, \text{list}(l)\}_{\text{old}} \quad \{\text{list}(l)\} \text{push}(l, n) \{l', _, \text{list}(l')\}_{\text{old}}$$

$$\{\text{list}(l)\} \text{set}(l, i, n) \{_, _, \text{list}(l)\}_{\text{old}}$$

Length Specification

$$\{\text{True}\} \text{new}() \{l, _, \text{list}(l) * \text{len}(l) \Downarrow 0\}_{\text{old}}$$

$$\{\text{list}(l)\} \text{push}(l, n) \{l', \text{old. list}(l') * \text{len}(l') - 1 \Downarrow \text{old} \{\text{len}(l)\}\}_{\text{old}}$$

$$\{\text{list}(l)\} \text{set}(l, i, n) \{_, \text{old. list}(l) * \text{len}(l) \Downarrow \text{old} \{\text{len}(l)\}\}_{\text{old}}$$

Functional Correctness Specification

$$\{\text{True}\} \text{new}() \{l, _, \text{list}(l) * \text{len}(l) \Downarrow 0 * \forall (i : \mathbb{Z}). \text{nth}(l, i) \Downarrow \text{None}\}_{\text{old}}$$

$$\{\text{list}(l)\}$$

$$\text{push}(l, n)$$

$$\left\{ l', \text{old.} \begin{array}{l} \text{list}(l') * \text{len}(l') - 1 \Downarrow \text{old} \{\text{len}(l)\} * \text{nth}(l', 0) \Downarrow \text{Some}(n) \\ * \forall (j : \mathbb{Z}). 0 \leq_{\text{hp}} j \Rightarrow \text{nth}(l', j + 1) \Downarrow \text{old} \{\text{nth}(l, j)\} \end{array} \right\}_{\text{old}}$$

$$\{\text{list}(l)\}$$

$$\text{set}(l, i, n)$$

$$\left\{ \begin{array}{l} \text{list}(l) * \text{len}(l) \Downarrow \text{old} \{\text{len}(l)\} \\ _, \text{old.} * (0 \leq_{\text{hp}} i \wedge i <_{\text{hp}} \text{len}(l) \Rightarrow \text{nth}(l, i) \Downarrow \text{Some}(n)) \\ * (\forall (j : \mathbb{Z}). 0 \leq_{\text{hp}} j \wedge j <_{\text{hp}} \text{len}(l) \wedge i \neq_{\text{hp}} j \Rightarrow \text{nth}(l, j) \Downarrow \text{old} \{\text{nth}(l, j)\}) \end{array} \right\}_{\text{old}}$$

Fig. 17. Three different types of linked-list function specifications, ordered by strength, where we abbreviate $e_1 \leq_{\text{hp}} e_2 \triangleq \exists n_1, n_2. e_1 \Downarrow n_1 \wedge e_2 \Downarrow n_2 \wedge n_1 \leq n_2$ and $e_1 <_{\text{hp}} e_2 \triangleq \exists n_1, n_2. e_1 \Downarrow n_1 \wedge e_2 \Downarrow n_2 \wedge n_1 < n_2$ and $e_1 \neq_{\text{hp}} e_2 \triangleq \exists v_1, v_2. e_1 \Downarrow v_1 \wedge e_2 \Downarrow v_2 \wedge v_1 \neq v_2$.

Specification

```

list_empty()  $\triangleq$  None
list_cons(val, l)  $\triangleq$  Some(ref(val, l))
list_find(eq, key, l)  $\triangleq$  match l with
  | None  $\Rightarrow$  None
  | Some(p)  $\Rightarrow$ 
    let (hd, tl) = ! p in
    let (k, val) = hd in
    if eq(key, k) then Some(val) else list_find(eq, key, tl)
list_update(eq, key, val, l)  $\triangleq$  match l with
  | None  $\Rightarrow$  (list_cons((key, val), None), false)
  | Some(p)  $\Rightarrow$ 
    let (hd, tl) = ! p in
    let (k, val) = hd in
    if eq(key, k) then p := ((k, val), tl); (Some(p, true)) else
      let (tl', updated) = list_update(eq, key, val, tl) in
      p := (hd, tl'); (Some(p), updated)

```

Representation Predicate

```

islist(x, [])  $\triangleq$  x = None
islist(x, (k, v) :: vs)  $\triangleq$   $\exists l, w. x = \text{Some}(l) \wedge l \mapsto ((k, v), w) \wedge \text{islist}(w, vs)$ 

```

Fig. 18. The linked list implementation parameterized by a custom eq.

E.4 Polymorphic Hashmap (#4)

In this example, we consider a polymorphic hashmap implementation. The hashmap is parameterized by a custom equality function `eq` and a hash function `hash`. The implementation is depicted in Fig. 19. It uses a linked-list implementation depicted in Fig. 18. The hashmap is implemented as an array of “buckets” (*i.e.*, linked-lists storing key-value pairs).

Specifications. The specifications that we prove are depicted in Fig. 20. They are parametric over a key type $K : \text{Val} \rightarrow h\text{Prop}$ and a persistent predicate $K_o : \text{Val} \rightarrow i\text{Prop}$ for ownership of the keys, where we require the following two properties:

$$\forall k. K_o(k) \vdash K(k) \qquad \forall k. \text{persistent}(K_o(k))$$

For an equality function `eq` and a hash function `hash`, the function `hashmap_create` allows us to create a new hashmap (of non-zero size). We will define its precondition `iseqhash` below. It returns an empty hashmap. To track ownership and contents of the hashmap, we use the abstract predicate `ishashmap(eq, h, m)`. It tracks the physical value corresponding to the hashmap `h`, the equality function `eq` that was used to create the hashmap, and the current contents `m`.

Once created, we can resize a hashmap with `hashmap_resize`, which will return a new hashmap (of increased capacity).

```

hashmap_create(eq, hash, cap)  $\triangleq$  (eq, hash, cap, allocN(cap, list_empty()))6
hashmap_get(hm, key)  $\triangleq$  let (eq, hash, cap, arr) = hm in
    let idx = hash(key) % cap in
    let b = ! (arr + idx) in
    list_find(eq, key, b)
hashmap_put(hm, key, val)  $\triangleq$  let (eq, hash, cap, arr) = hm in
    let idx = hash(key) % cap in
    let b = ! (arr + idx) in
    let (b', updated) = list_update(eq, key, val, b) in
    arr + idx := b'; updated
iter_bucket(bucket, f)  $\triangleq$  match bucket with
    | None  $\Rightarrow$  ()
    | Some(b)  $\Rightarrow$ 
        let (hd, tl) = ! b in
        let (k, v) = hd in
        iter_bucket(tl, f); f k v
iter_buckets(arr, i, cap, f)  $\triangleq$  if i < cap then
    let bucket = ! (arr + i) in
    iter_buckets(arr, i + 1, cap, f);
    iter_bucket(bucket, f) else ()
hashmap_resize(hm)  $\triangleq$  let (eq, hash, cap, arr) = hm in
    let map = hashmap_create(eq, hash, 2 * cap) in
    iter_buckets(arr, 0, cap,  $\lambda$  k v. hashmap_put(map, k, v));
    map

```

Fig. 19. The hashmap implementation.

We can get an element for key k in the hashmap with `hashmap_get`. It returns either an element from the hashmap such that comparison with the equality function `eq` results in `true`, or it returns no element and we know that every key in the hashmap is not equal to k according to `eq`.

We can set an element for key k with `hashmap_put`. It returns `true` if the key was already contained (up to `eq`) in the hashmap and an update has been performed. It returns `false` if the key was not present (up to `eq`) and instead the new key-value pair has been added to the map.

⁶The operation `allocN(n , v)` allocates an array of length n , where each element is initialized with v .

$$\begin{array}{ll}
\{0 < c * \text{iseqhash}(\text{eq}, \text{hash})\} & \{\text{ishashmap}(\text{eq}, h, m)\} \\
\text{hashmap_create}(\text{eq}, \text{hash}, c) & \text{hashmap_resize}(h) \\
\{u. \text{ishashmap}(\text{eq}, u, [])\} & \{u. \text{ishashmap}(\text{eq}, u, m)\} \\
\\
\{\text{ishashmap}(\text{eq}, h, m) * K_o(k)\} & \\
\text{hashmap_get}(h, k) & \\
\left\{ \begin{array}{l} (\exists k, w. u = \text{Some}(w) * (k', w) \in m * \Box \text{eq}(k, k') \Downarrow \text{true}) \\ \vee \\ (u = \text{None} * \Box \forall (k', _) \in m. \text{eq}(k, k') \Downarrow \text{false}) \end{array} \right\} \\
\\
\{\text{ishashmap}(\text{eq}, h, m) * K_o(k)\} & \\
\text{hashmap_put}(h, k, v) & \\
\left\{ \begin{array}{ll} \text{then } \exists k', i. \text{ishashmap}(\text{eq}, h, m[i \mapsto (k', v)]) * m[i] = (k', _) * \Box \text{eq}(k, k') \Downarrow \text{true} \\ u. \exists b. u = b * \text{if } b & \text{else } \text{ishashmap}(\text{eq}, h, (k, v) :: m) * \Box \forall (k', _) \in m. \text{eq}(k, k') \Downarrow \text{false} \end{array} \right\}
\end{array}$$

Fig. 20. The hashmap specification.

The equality and hash function. Let us now focus on the properties that we require of the functions `eq` and `hash` when creating a hashmap:

$$\begin{aligned}
\text{refl}(\text{eq}) &\triangleq \text{smt } (\forall x. K(x) \Rightarrow \text{eq}(x, x) \Downarrow \text{true}) \\
\text{sym}(\text{eq}) &\triangleq \text{smt } (\forall x, y. K(x) \Rightarrow K(y) \Rightarrow \text{eq}(x, y) \Downarrow \text{true} \Rightarrow \text{eq}(y, x) \Downarrow \text{true}) \\
\text{trans}(\text{eq}) &\triangleq \text{smt } \left(\forall x, y, z. K(x) \Rightarrow K(y) \Rightarrow K(z) \Rightarrow \right. \\
&\quad \left. \text{eq}(x, z) \Downarrow \text{true} \Rightarrow \text{eq}(x, y) \Downarrow \text{true} \Rightarrow \text{eq}(y, z) \Downarrow \text{true} \right) \\
\text{eqhash}(\text{eq}, \text{hash}) &\triangleq \text{smt } (\forall x, y. K(x) \Rightarrow K(y) \Rightarrow \text{eq}(x, y) \Downarrow \text{true} \Rightarrow \text{hash}(x) \equiv \text{hash}(y)) \\
\text{eqty}(\text{eq}) &\triangleq \Box(\text{eq} \in \mathcal{V} \llbracket K \times K \rightarrow \text{bool} \rrbracket) \\
\text{hashty}(\text{hash}) &\triangleq \Box(\text{hash} \in \mathcal{V} \llbracket K \rightarrow \text{int} \rrbracket) \\
\text{iseq}(\text{eq}) &\triangleq \text{sym}(\text{eq}) * \text{trans}(\text{eq}) * \text{refl}(\text{eq}) * \text{eqty}(\text{eq}) \\
\text{iseqhash}(\text{eq}, \text{hash}) &\triangleq \text{iseq}(\text{eq}) * \text{hashty}(\text{hash}) * \text{eqhash}(\text{eq}, \text{hash}) \\
\text{smt}(P) &\triangleq \Box(\text{ambient_heap} \multimap P)
\end{aligned}$$

The function `eq` must be of type $K \times K \rightarrow \text{bool}$, and the hash function `hash` of type $K \rightarrow \text{int}$. The equality function must be reflexive (`refl`), symmetric (`sym`), and transitive (`trans`). Moreover, equal elements must have the same hash (`eqhash`). Since we wish to be able to establish reflexivity, symmetry, transitivity, and the `eq`-hash relation with an SMT solver, we wrap them with `smt` (`_`). It adds the ambient heap (see §D.3) as a condition to the properties, because this is what the first-order correspondence gives us when we come from first-order logic (see Lemma D.7 in §D.3).

The hashmap predicate. Let us now turn to the hashmap predicate. Each bucket in the hashmap is represented as a linked list, with each element being a pair of a key and a value. Given an `eq` and

hash function, we define the representation predicate for a bucket as follows:

$$\text{bucketeq}(b) \triangleq \forall i, j, k, k', v, v'. b_i = (k, v) \Rightarrow b_j = (k', v') \Rightarrow i \neq j \Rightarrow \text{eq}(k, k') \Downarrow \text{false}$$

$$\text{buckethash}(b, c, i) \triangleq \forall (k, _) \in b. (\text{hash}(k) \% c) \Downarrow i$$

$$\text{isbucket}(b, c, i, l) \triangleq \Box \text{bucketeq}(b) * \Box \text{buckethash}(b, c, i) * \exists v. l + i \mapsto v * \text{islist}(v, b) * \bigstar_{(k, _) \in b} K_o(k)$$

A hashmap is represented as a tuple of the equality function, the hash function, the capacity, and an array of buckets. The representation predicate for the hashmap is defined as follows:

$$\begin{aligned} \text{ishashmap}(\text{eq}, h, m) \triangleq \exists \text{hash}, c, l, C. h = (\text{eq}, \text{hash}, c, l) * |C| = c * c > 0 * \text{concat}(C) \equiv_p m \\ * \text{iseqhash}(\text{eq}, \text{hash}) * \bigstar_{i=1, \dots, c} \text{isbucket}(C_i, c, i, l) \end{aligned}$$

where “ $\text{concat}(C) \equiv_p m$ ” means m is a permutation of the list obtained by concatenating the list of lists C that tracks the contents of the individual buckets in the hashmap.

E.5 Iris Examples (#5)

The modified code and the original version are provided as part of the accompanying Rocq development.

References

- [1] Amal Ahmed, Andrew W. Appel, Christopher D. Richards, Kedar N. Swadi, Gang Tan, and Daniel C. Wang. 2010. Semantic foundations for typed assembly languages. *ACM Trans. Program. Lang. Syst.* 32, 3 (2010), 7:1–7:67. <https://doi.org/10.1145/1709093.1709094>
- [2] Andrew W. Appel and David A. McAllester. 2001. An indexed model of recursive types for foundational proof-carrying code. *ACM Trans. Program. Lang. Syst.* 23, 5 (2001), 657–683. <https://doi.org/10.1145/504709.504712>
- [3] Thibault Dardinier, Michael Sammler, Gaurav Parthasarathy, Alexander J. Summers, and Peter Müller. 2024. Formal Foundations for Translational Separation Logic Verifiers (extended version). arXiv:2407.20002 [cs.PL] <https://arxiv.org/abs/2407.20002>
- [4] Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. 2016. Higher-order ghost state. In *ICFP*. ACM, 256–269. <https://doi.org/10.1145/2951913.2951943>
- [5] Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. 2018. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.* 28 (2018), e20. <https://doi.org/10.1017/S0956796818000151>
- [6] Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *POPL*. ACM, 637–650. <https://doi.org/10.1145/2676726.2676980>
- [7] Peter Müller, Malte Schwerhoff, and Alexander J. Summers. 2016. Viper: A Verification Infrastructure for Permission-Based Reasoning. In *VMCAI (LNCS, Vol. 9583)*. Springer, 41–62. https://doi.org/10.1007/978-3-662-49122-5_2
- [8] Redis. 2024. Redis Pcount Implementation for Potentially Large Buffers. <https://github.com/redis/redis/blob/3fac869f02657d94dc89fab23acb8ef188889c96/src/bitops.c#L40>.
- [9] Michael Sammler, Rodolphe Lepigre, Robbert Krebbers, Kayvan Memarian, Derek Dreyer, and Deepak Garg. 2021. RefinedC: Automating the foundational verification of C code with refined ownership types. In *PLDI*. ACM, 158–174. <https://doi.org/10.1145/3453483.3454036>
- [10] Kasper Svendsen and Lars Birkedal. 2014. Impredicative Concurrent Abstract Predicates. In *ESOP (LNCS, Vol. 8410)*. Springer, 149–168. https://doi.org/10.1007/978-3-642-54833-8_9
- [11] Amin Timany, Robbert Krebbers, Derek Dreyer, and Lars Birkedal. 2024. A Logical Approach to Type Soundness. *J. ACM* (July 2024). <https://doi.org/10.1145/3676954>