



Program Logics à la Carte

MAX VISTRUP, ETH Zurich, Switzerland

MICHAEL SAMMLER, ETH Zurich, Switzerland

RALF JUNG, ETH Zurich, Switzerland

Program logics have proven a successful strategy for verification of complex programs. By providing local reasoning and means of abstraction and composition, they allow reasoning principles for individual components of a program to be combined to prove guarantees about a whole program. Crucially, these components and their proofs can be *reused*. However, this reuse is only available once the program logic has been defined. It is a frustrating fact of the status quo that whoever defines a new program logic must establish every part, both semantics and proof rules, from scratch. In spite of programming languages and program logics typically sharing many core features, reuse is generally not available across languages. Even inside one language, if the same underlying operation appears in multiple language primitives, reuse is typically not possible when establishing proof rules for the program logic.

To enable reuse across and inside languages when defining complex program logics (and proving them sound), we serve program logics *à la carte* by combining program logic fragments for the various effects of the language. Among other language features, the menu includes shared state, concurrency, and non-determinism as reusable, composable blocks that can be combined to define a program logic modularly. Our theory builds on ITrees as a framework to express language semantics and Iris as the underlying separation logic; the work has been mechanized in the Coq proof assistant.

CCS Concepts: • **Theory of computation** → **Logic and verification**; **Separation logic**; *Program semantics*; Programming logic.

Additional Key Words and Phrases: verification, separation logic, Interaction Trees, Iris, Coq

ACM Reference Format:

Max Vistrup, Michael Sammler, and Ralf Jung. 2025. Program Logics à la Carte. *Proc. ACM Program. Lang.* 9, POPL, Article 11 (January 2025), 32 pages. <https://doi.org/10.1145/3704847>

1 Introduction

Program logics are a widely successful approach to program verification [14, 45, 40, 9, 2, 27, 29, 19]. However, they require a non-trivial amount of preparation, especially for programs written in complicated languages. First, a *formal definition of the language semantics* needs to be developed. Then, one must find and state the *rules of the program logic*. Finally, the two need to be connected by a *soundness proof*.

Let us consider what is required to build a program logic for a new language. Probably our language has some form of global mutable state; maybe it has concurrency; maybe it has other forms of non-determinism. The dominant approach to defining language semantics is to use an operational semantics, which has standard ways to model all these language features. To obtain a program logic, again there are common ways of reasoning about such features (the combination of

Authors' Contact Information: [Max Vistrup](mailto:max.vistrup@inf.ethz.ch), ETH Zurich, Zürich, Switzerland, max.vistrup@inf.ethz.ch; [Michael Sammler](mailto:michael.sammler@inf.ethz.ch), ETH Zurich, Zürich, Switzerland, michael.sammler@inf.ethz.ch; [Ralf Jung](mailto:ralf.jung@inf.ethz.ch), ETH Zurich, Zürich, Switzerland, ralf.jung@inf.ethz.ch.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 2475-1421/2025/1-ART11

<https://doi.org/10.1145/3704847>

mutable state and concurrency makes concurrent separation logic a common choice), so we know the rough shape the logic takes, and we just have to make some adjustments to account for the particularities of our concrete language. Maybe the language has some particularly complicated operation that does many things at once; this will require a complicated transition in the operational semantics and an associated complicated proof rule in the program logic. To justify the correctness of the logic, we again have to do proofs that are largely standard.

All of this is a lot of work! If one aims to fully formalize the entire logic and soundness proof (whether on paper or in a mechanized proof), one ends up proving very similar theorems for each new language. The *pattern* is always the same, but there is no reusable theorem that would let us, say, “plug in” a heap with associated points-to assertions $\ell \mapsto v$ to obtain the standard separation logic reasoning principles.

Fundamentally, this is caused by the fact that typical language definitions are monolithic: an operational semantics captures *all* the state that is relevant for the behavior of the program, and a *single* relation describes how all language constructs act on the entire state. Operational semantics provide no clear way to define the operational behavior of a heap once and for all, and turn it into a reusable component with associated reasoning principles. It also provides no good way to compose a single, complicated operation from smaller pieces: every program step is a single state transition. This becomes particularly onerous for concurrent languages where the steps of a small-step operational semantics typically mark the granularity of atomicity, meaning that one cannot simply define a complicated atomic operation as syntactic sugar for many smaller operations.

In this paper, we present *program logics à la carte*: a new approach for defining program semantics and associated program logics from reusable building blocks.

To achieve this, we leverage *ITrees* [48] to represent program semantics in a style that is closer to denotational semantics rather than operational semantics. *ITrees* provide a general monadic encoding of effectful computations in a pure meta-language. They are parametric in the set of effects that the program may invoke. Examples of such effects are non-determinism, state, or concurrency. In this work, we consider these effects to be the *building blocks* that make up a program semantics.

The core of our work is a general program logic for arbitrary *ITrees*. Just like *ITrees* are parameterized over effects, our program logic is parameterized over *logical effect handlers* that define the verification condition for invoking an effect (think: preconditions and postconditions). We have implemented logical effect handlers for a number of common effects, including mutable state, non-deterministic choice (both demonic and angelic), concurrency, and abnormal program failure. These are the building blocks that a language designer has at their disposal, and they all come with an associated program logic fragment that is established once-and-for-all, and an adequacy theorem showing soundness of this program logic fragment.

A language designer can pick the effects of their choice from this menu and describe the language semantics by denoting their language into *ITrees*, composing the primitive operations provided by these effects to build up the core language operations as needed. They can then use the general *ITree* program logic to define a language-specific program logic. The program logic fragments of each effect are automatically available, and their proof rules can be used to build up the reasoning principles for the core language operations. Crucially, we are able to use the compositional power of program-logic-based reasoning to establish the program logic itself. *ITrees* serve as a common foundation for effectful computation—a shared domain with a wide range of applicability.

Our program logic is based on *Iris* [21], which is a natural choice since *Iris* already has a strong focus on modularity and reuse: *Iris* is a “separation logic *framework*” designed to serve as the foundation for domain-specific separation logics [20, 6, 33, 32, 16, 35, 28]. *Iris* already provides reusable building blocks for “ghost state” constructions to capture common reasoning patterns (such as finite maps with fractional per-key permissions, or append-only lists). However, so far

only very limited reuse was possible for the part of the program logic that directly interacts with the language semantics, leading to a lot of duplicated effort across Iris-based projects. With our new approach, this is no longer necessary: language components and their associated program logic rules can be shared and reused with the same ease that Iris users already commonly share and reuse purely logical components.

Contributions. The key contributions of our work are a novel, general-purpose program logic for ITrees and a library of effects with associated reasoning principles expressed in that program logic. We have built effect libraries for concurrency, global state, demonic and angelic non-determinism, safe and unsafe program termination, and partial correctness. The program logic is proven sound (or “adequate”) w.r.t. two notions of execution for ITrees: the typical ITree approach involving gradual interpretation of events, and a novel translation of ITrees into state machines. To demonstrate the potential for reuse and the applicability of this framework, we have ported two existing Iris-based program logics: the default Iris example language, HeapLang, as well as the language used by Islaris [32], a verification approach for machine-code programs against authoritative ISA specifications. The HeapLang program logic supports both total and partial correctness reasoning and comes with a provably-sound interpreter (including a termination proof).

Formalization in Coq. All our work is mechanized in the Coq proof assistant [42]. Our development uses three axioms: To deal with destruction of dependent types, we make use of Axiom K [43]. In relating the interpretation semantics and operational semantics of HeapLang in §6, we make use of the functional form of the (non-extensional) axiom of choice [43]. Finally, we use an axiom that promotes strong ITree bisimulation to equality. Bisimulation-is-equality axioms are sometimes used when working with coinductively defined data types in Coq because Coq’s notion of equality on coinductively defined data types is far too fine [47, 15].

Structure of the paper. The rest of the paper is structured as follows: First, §2 gives an overview of our approach by gradually equipping a language with more and more effects and building up an associated program logic alongside. It also shows how to prove program logics adequate with respect to semantics based on “interpretations”. Then, §3 explains how our program logic and effect libraries are defined in technical detail. The next two sections extend these foundations to cover more kinds of effects: §4 deals with concurrency and §5 handles angelic choice. Finally, §6 and §7 describe the HeapLang and Islaris case studies. We conclude by discussing related work in §8.

2 Key Ideas

In this section, we showcase the key ideas of this paper by building a simple example language and an associated program logic step-by-step. Each step adds one more building block, demonstrating how the language and program logic are built up from reusable components.

2.1 Starting Point: A Pure Lambda Calculus ($\lambda_{\mathbb{Z}}$)

We start with $\lambda_{\mathbb{Z}}$, a basic untyped λ -calculus with integers, addition, and if-expressions:

$$v \in \text{val} ::= z \mid \lambda x. e \quad (z \in \mathbb{Z}) \quad e \in \text{expr} ::= v \mid x \mid e_1 \hat{+} e_2 \mid e_1(e_2) \mid \text{if } e_1 \text{ then } e_2 \text{ else } e_3$$

Background: ITrees. To reason about $\lambda_{\mathbb{Z}}$, we first need to give it a semantics. For this, we use a denotational semantics with *interaction trees* or *ITrees* [48] as our domain. ITrees are a general domain for representing effectful computations. The ITree type **itree** $E R$ is parameterized by a return type $R : \mathbf{Type}$ and an event type $E : \mathbf{Type} \rightarrow \mathbf{Type}$. The type $E A$ should be understood as those events ϵ which have answer type A , that is, which return an answer $a : A$. Intuitively, an ITree is a computation consisting of three kinds of steps: it can either (1) terminate and return a

$$\begin{aligned}
\llbracket v \rrbracket &:= \text{Ret}(v) \\
\llbracket e_1 \hat{+} e_2 \rrbracket &:= v_1 \leftarrow \llbracket e_1 \rrbracket; v_2 \leftarrow \llbracket e_2 \rrbracket; z_2 \leftarrow \text{to_int}(v_2); z_1 \leftarrow \text{to_int}(v_1); \text{Ret}(z_1 + z_2) \\
\llbracket e_1(e_2) \rrbracket &:= v_1 \leftarrow \llbracket e_1 \rrbracket; v_2 \leftarrow \llbracket e_2 \rrbracket; (x, e) \leftarrow \text{to_lam}(v_1); \llbracket e[v_2/x] \rrbracket \\
\llbracket \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \rrbracket &:= v_1 \leftarrow \llbracket e_1 \rrbracket; z_1 \leftarrow \text{to_int}(v_1); \text{if } z_1 \neq 0 \text{ then } \llbracket e_2 \rrbracket \text{ else } \llbracket e_3 \rrbracket \\
\text{to_int}(v) &:= \text{match } v \text{ with } z \implies \text{Ret}(z) \mid _ \implies \text{fail end} \\
\text{to_lam}(v) &:= \text{match } v \text{ with } \lambda x. e \implies \text{Ret}((x, e)) \mid _ \implies \text{fail end}
\end{aligned}$$

Fig. 1. Denotational semantics of $\lambda_{\mathbb{Z}}$.

value of type R , (2) perform a silent step, or (3) perform an effectful computation by emitting an event ϵ that describes the effect that is performed. Formally, ITrees are coinductively defined as

$$t \in \mathbf{itree} \ E R ::=_{\text{coind}} \text{Ret}(r : R) \mid \text{Tau}(t : \mathbf{itree} \ E R) \mid \text{Vis}_A(\epsilon : EA, k : A \rightarrow \mathbf{itree} \ E R)$$

where

- (1) $\text{Ret}(r)$ represents just returning value $r : R$,
- (2) $\text{Tau}(t)$ represents taking a silent step and continuing the computation t , and
- (3) $\text{Vis}_A(\epsilon, k)$ represents emitting an event $\epsilon : EA$, receiving answer $a : A$, and then continuing with $k(a)$. (We shall often elide A and write just $\text{Vis}(\epsilon, a)$.)

Tau steps do not represent a visible action in the program, and as such, it is often desirable to ignore them. However, infinite sequences of Tau steps (which are possible due to ITrees being a *coinductive* data type) are still relevant since they represent diverging computations. To this end, ITrees come with the bisimulation \approx , known as “equivalence up to (finitely many) Taus”, which allows removing/introducing finitely many Taus on either side. This is the canonical extensional notion of equality on ITrees.

$\mathbf{itree} \ E R$ is a monad in R . More specifically, the monadic laws hold up to \approx . The monadic bind enables us to write ITrees in the style of sequential code: $x_1 \leftarrow \dots; x_2 \leftarrow \dots; \dots$.

Denotation. To define the semantics of $\lambda_{\mathbb{Z}}$, we are looking to denote expressions into ITrees. The first step is to pick a suitable set of events E . So far, our language has one effect (not counting non-termination which is natively supported by ITrees): programs can fail, such as when trying to add a number and a function. To represent failure, we use an event type **FailE** with a single event `fail` with answer type \emptyset , thus giving us access to an operation `fail` : $\mathbf{itree} \ \mathbf{FailE} \ \emptyset$.¹ A failure does not return but “crashes” the program and thus the answer type of `fail` is the empty set. To summarize, the set of events for $\lambda_{\mathbb{Z}}$ is:

$$\mathbf{LangE}_{\mathbb{Z}} := \mathbf{FailE}$$

With $\mathbf{LangE}_{\mathbb{Z}}$ at hand, we can define a function $\llbracket _ \rrbracket$ that maps (closed) expressions into the domain $\mathbf{itree} \ \mathbf{LangE}_{\mathbb{Z}} \ \text{val}$ as shown in Figure 1. This is a shallow embedding: pure computations in the language are mapped to computations in the meta-logic, as can be seen for the addition or if-expressions. Operations not supported by the meta-logic are treated monadically. We can also easily represent non-structural recursion (as in the case of function application) thanks to the general fixpoint combinator provided by ITrees. $\llbracket _ \rrbracket$ can be viewed as a form of denotational semantics, but note that events are still uninterpreted at this stage and hence treated purely syntactically.

¹Here, we implicitly coerce events ϵ to ITrees that emit this event, e.g., if we write `fail` in a context where an ITree is expected, it implicitly desugars it to $\text{Vis}(\text{fail}, (\lambda a. \text{Ret}(a)))$. This construction is called `trigger` in the ITree library.

$$\begin{array}{c}
\text{WPConsequence} \\
\frac{\forall r. \Phi(r) \vdash \Psi(r) \quad \text{wp } e \{ \Phi \}}{\text{wp } e \{ \Psi \}}
\end{array}
\qquad
\begin{array}{c}
\text{WPFrame} \\
\frac{P * \text{wp } e \{ \Phi \}}{\text{wp } e \{ v. P * \Phi(v) \}}
\end{array}
\qquad
\begin{array}{c}
\text{WPVal} \\
\frac{\Phi(v)}{\text{wp } v \{ \Phi \}}
\end{array}$$

$$\begin{array}{c}
\text{WPBindPlusL} \\
\frac{\text{wp } e_1 \{ v_1. \text{wp } v_1 \hat{+} e_2 \{ \Phi \} \}}{\text{wp } e_1 \hat{+} e_2 \{ \Phi \}}
\end{array}
\qquad
\begin{array}{c}
\text{WPBindPlusR} \\
\frac{\text{wp } e_2 \{ v_2. \text{wp } v_1 \hat{+} v_2 \{ \Phi \} \}}{\text{wp } v_1 \hat{+} e_2 \{ \Phi \}}
\end{array}
\qquad
\begin{array}{c}
\text{WPPlus} \\
\frac{\Phi(z_1 + z_2)}{\text{wp } z_1 \hat{+} z_2 \{ \Phi \}}
\end{array}$$

$$\begin{array}{c}
\text{WPApp} \\
\frac{\text{wp } e \{ v/x \} \{ \Phi \}}{\text{wp } (\lambda x. e)(v) \{ \Phi \}}
\end{array}
\qquad
\begin{array}{c}
\text{WPIfTrue} \\
\frac{\text{wp } e_1 \{ \Phi \} \quad z \neq 0}{\text{wp if } z \text{ then } e_1 \text{ else } e_2 \{ \Phi \}}
\end{array}
\qquad
\begin{array}{c}
\text{WPIfFalse} \\
\frac{\text{wp } e_2 \{ \Phi \}}{\text{wp if } 0 \text{ then } e_1 \text{ else } e_2 \{ \Phi \}}
\end{array}$$

Fig. 2. Excerpt of the program logic for the $\lambda_{\mathbb{Z}}$.

Program logic. To reason about programs in $\lambda_{\mathbb{Z}}$, we will define a program logic. As the framework for our program logic we use Iris [21], a versatile separation logic that comes with a good foundation of reusable reasoning principles and has already been used as the basis for numerous program logics [20, 6, 33, 32, 16, 28]. Following the usual approach in Iris, we use a *weakest precondition* connective as the core of our program logic. A more traditional Hoare-style program logic can be easily defined on top of this by setting $\{P\} e \{Q\} := P * \text{wp } e \{Q\}$; the typical rules can be derived.

Concretely, $\text{wp } e \{ \Phi \}$ says that in the current state (which can be constrained by assumptions in the logical context), every execution of e is well-behaved and the returned value v satisfies the postcondition $\Phi(v)$. We obtain the expected rules for wp , as shown in Figure 2:² we have the rule of consequence and the specific rules for each language construct. For instance, **WPPlus** says that $z_1 \hat{+} z_2$ satisfies postcondition Φ whenever the corresponding mathematical term $z_1 + z_2$ satisfies Φ . **WPBindPlusL** is a “bind” rule that recurses into the program structure; it says that to reason about $e_1 \hat{+} e_2$, we can first reason about e_1 . Every value v_1 that e_1 can evaluate to must then satisfy the property that any evaluation of $v_1 \hat{+} e_2$ satisfies the desired postcondition. **WPBindPlusR** does the same for the right-hand operand (but this rule only applies if the left-hand operand is a value, indicating a left-to-right evaluation order). We omit similar rules for application and if-expressions. (The frame rule **WPFrame** will only become relevant when we get to reasoning about state.)

These rules are entirely standard. The key idea of our approach lies in *how* wp is defined: instead of defining a new wp for each language, we want to enable reuse across languages. Therefore, we introduce a new general-purpose weakest precondition connective for *arbitrary ITrees*: given $t : \text{itree } E R$, we define $\text{wpi}_H t \{ \Phi \}$ as the weakest precondition that ensures t terminates with a return value that satisfies postcondition Φ . The definition takes as input a *logical effect handler* (or just *handler*) H which provides specifications for the events in E . On top of this, we can define the weakest precondition for $\lambda_{\mathbb{Z}}$ by choosing a suitable handler **LangH $_{\mathbb{Z}}$** and then setting:

$$\text{wp } e \{ \Phi \} := \text{wpi}_{\text{LangH}_{\mathbb{Z}}} \llbracket e \rrbracket \{ \Phi \}$$

The reason this is useful is that wpi satisfies the rules in Figure 3—they come “for free”, without us having to do any language-specific work. Aside from the rule of consequence, we have rules for the bind and return operators of the ITree monad, as well as **WPiEUTT** which states that wpi

²Note that—following the standard Iris approach—these rules are (Iris-level) theorems about the definition of $\text{wp } e \{ \Phi \}$ found below. (In particular, $\text{wp } e \{ \Phi \}$ is *not* inductively defined using these rules as would be typical outside of Iris).

$$\begin{array}{c}
\text{WPICONSEQUENCE} \\
\frac{\forall r. \Phi(r) \vdash \Psi(r) \quad \text{wpi}_H t \{\Phi\}}{\text{wpi}_H t \{\Psi\}}
\end{array}
\qquad
\begin{array}{c}
\text{WPIFRAME} \\
\frac{P * \text{wpi}_H t \{\Phi\}}{\text{wpi}_H t \{v. P * \Phi(v)\}}
\end{array}
\qquad
\begin{array}{c}
\text{WPIEUTT} \\
\frac{t_1 \approx t_2}{\text{wpi}_H t_1 \{\Phi\} \dashv\vdash \text{wpi}_H t_2 \{\Phi\}}
\end{array}$$

$$\begin{array}{c}
\text{WPIBIND} \\
\frac{\text{wpi}_H t \{x. \text{wpi}_H k(x) \{\Phi\}\}}{\text{wpi}_H x \leftarrow t; k(x) \{\Phi\}}
\end{array}
\qquad
\begin{array}{c}
\text{WPIRET} \\
\frac{\Phi(r)}{\text{wpi}_H \text{Ret}(r) \{\Phi\}}
\end{array}$$

Fig. 3. Basic, generic proof rules for weakest preconditions.

is compatible with \approx (“equivalence up to τ ”). This is needed to unfold general recursive ITree definitions, such as our $\llbracket e \rrbracket$, which can only be unfolded up to \approx , not up to full definitional equality.

To complete this definition, we need to define the handler **LangH_Z**. The only event we have to worry about for now is fail, which comes with a handler **FailH** that assigns fail the precondition False. Accordingly, fail can never be called in a verified program. (We will see in §3.2 how exactly handlers capture event specifications.) We can thus pick **LangH_Z** := **FailH**.

The rules in Figure 2 for our weakest precondition wp are now easily derived from the rules in Figure 3 for the underlying ITree weakest precondition wpi. We consider two examples.

PROOF OF WPIFALSE. By definition of $\llbracket _ \rrbracket$ and monadic laws, we have $\llbracket \text{if } 0 \text{ then } e_1 \text{ else } e_2 \rrbracket \approx \llbracket e_2 \rrbracket$. Using **WPIEUTT**, this reduces our goal to $\text{wp } e_2 \{\Phi\}$, and we are done immediately. \square

PROOF OF WPBINDPLUSL. We have $\llbracket e_1 + e_2 \rrbracket \approx v_1 \leftarrow \llbracket e_1 \rrbracket; k(v_1)$ for k a notational shorthand for the continuation. Using **WPIEUTT** and **WPIBIND**, our goal thus turns into $\text{wp } e_1 \{v_1. \text{wp } k(v_1) \{\Phi\}\}$. By monadic laws, $\llbracket v_1 + e_2 \rrbracket \approx v'_1 \leftarrow \text{Ret}(v_1); k(v'_1) \approx k(v_1)$, and we are done by **WPIEUTT**. \square

These proofs demonstrate how the language-agnostic rules for wpi greatly simplify the typically tedious task of establishing proof rules for every single language construct.

2.2 2nd Effect: Mutable State ($\lambda_{Z,!}$)

To demonstrate that we can obtain a language and program logic by composing reusable pieces, we extend our language with another feature, a higher-order heap:

$$v \in \text{val} ::= \dots \mid \ell \quad (\ell \in \mathbb{N}) \qquad e \in \text{expr} ::= \dots \mid \text{ref}(e) \mid !e \mid e_1 \leftarrow e_2$$

Here, ℓ is a heap location. The term $\text{ref}(e)$ allocates a heap cell with content e , $!e$ loads the contents at heap location e , and $e_1 \leftarrow e_2$ stores e_2 at heap location e_1 .

In the setting of operational semantics, making such an extension to the language would require invasive surgery to the semantics: even the type of the stepping relation changes because it has to thread through the global state. However, as we shall see, with ITree-based semantics and program logics, we do not have to labor hard for an extension like this.

In our setting, this extension is provided via the **HeapE_V** event type for a value type V , which admits the following operations. (**HeapE_V** is defined formally in §3.3.)

- (1) **alloc** : $V \rightarrow \text{itree HeapE}_V \mathbb{N}$ which allocates a new heap cell,
- (2) **load** : $\mathbb{N} \rightarrow \text{itree HeapE}_V$ (option V) which loads the value in a heap cell (returning none if it is empty), and
- (3) **store** : $\mathbb{N} \rightarrow V \rightarrow \text{itree HeapE}_V$ (option V) which stores a value in a heap cell and returns the old value (or none if it was empty).

$$\text{WPREF} \quad \text{wp ref}(v) \{v'. \exists \ell. v' = \ell * \ell \mapsto v\}$$

$$\text{WPLoad} \quad \frac{\ell \mapsto v}{\text{wp} ! \ell \{v'. v = v' * \ell \mapsto v\}}$$

$$\text{WPSTORE} \quad \frac{\ell \mapsto v}{\text{wp} \ell \leftarrow v' \{w. w = v * \ell \mapsto v'\}}$$

Fig. 4. Program logic for $\lambda_{\mathbb{Z},!}$.

$$\text{WPIALLOC} \quad \text{wpi}_{\text{HeapH}_V} \text{ alloc}(v) \{\ell. \ell \mapsto v\}$$

$$\text{WPILOAD} \quad \frac{\ell \mapsto v}{\text{wpi}_{\text{HeapH}_V} \text{ load}(\ell) \{v'. v = v' * \ell \mapsto v\}}$$

$$\text{WPISTORE} \quad \frac{\ell \mapsto v}{\text{wpi}_{\text{HeapH}_V} \text{ store}(\ell, v') \{w. w = v * \ell \mapsto v'\}}$$

Fig. 5. Proof rules for $\text{wpi}_{\text{HeapH}_V}$.

The event type for $\lambda_{\mathbb{Z},!}$ is defined as $\mathbf{LangE}_{\mathbb{Z},!} := \mathbf{FailE} \oplus \mathbf{HeapE}_{\text{val}}$, using the sum operator (\oplus) on event types. (We use **blue color** to indicate changes to previous definitions.)

We extend the denotation $\llbracket e \rrbracket : \mathbf{itree} \mathbf{LangE}_{\mathbb{Z},!} \text{ val}$ to account for the new operations:

$$\llbracket \text{ref}(e) \rrbracket := v \leftarrow \llbracket e \rrbracket; \ell \leftarrow \text{alloc}(v); \text{Ret}(\ell)$$

$$\llbracket !e \rrbracket := v \leftarrow \llbracket e \rrbracket; \ell \leftarrow \text{to_loc}(v); w \leftarrow \text{load}(\ell); \text{unwrap}(w)$$

$$\llbracket e_1 \leftarrow e_2 \rrbracket := v_1 \leftarrow \llbracket e_1 \rrbracket; v_2 \leftarrow \llbracket e_2 \rrbracket; \ell \leftarrow \text{to_loc}(v_1); w \leftarrow \text{store}(\ell, v_2); \text{unwrap}(w)$$

$$\text{to_loc}(v) := \text{match } v \text{ with } \ell \implies \text{Ret}(\ell) \mid _ \implies \text{fail end}$$

$$\text{unwrap}(w) := \text{match } w \text{ with some}(v) \implies \text{Ret}(v) \mid \text{none} \implies \text{fail end}$$

Program logic. Our library also provides a handler \mathbf{HeapH}_V for \mathbf{HeapE}_V . To obtain a program logic for the extended language, we can combine the handler for \mathbf{FailE} and for $\mathbf{HeapE}_{\text{val}}$ into a handler $\mathbf{LangH}_{\mathbb{Z},!} := \mathbf{FailH} \oplus \mathbf{HeapH}_{\text{val}}$, and update wp to use this new handler.

The extended logic continues to satisfy all the rules in Figure 2—all proofs continue to proceed as before. In addition, we obtain the rules for the heap primitives as shown in Figure 4. These rules are direct consequences of the rules for the \mathbf{HeapE} operations displayed in Figure 5 (which lift to $\text{wpi}_{\mathbf{LangH}_{\mathbb{Z},!}}$ as we will see in §3.2). Some technicalities are omitted from the latter rules (consult §3.3).

2.3 3rd Effect: Non-Determinism ($\lambda_{\mathbb{Z},!,\text{pick}}$)

The next extension we consider is non-determinism. Namely, we extend the language with an operation to pick an arbitrary integer:

$$e \in \text{expr} ::= \dots \mid \text{pick_int}()$$

This introduces one new effect into the language: $\mathbf{DemonicE}$, modeling demonic non-determinism. This event type contains an event choice_A with answer type A for each inhabited³ type $a : A$. Accordingly, we define $\mathbf{LangE}_{\mathbb{Z},!,\text{pick}} := \mathbf{FailE} \oplus \mathbf{HeapE}_{\text{val}} \oplus \mathbf{DemonicE}$ and extend $\llbracket _ \rrbracket$:

$$\llbracket \text{pick_int}() \rrbracket := z \leftarrow \text{choice}_{\mathbb{Z}}; \text{Ret}(z)$$

Program logic. Our library provides a handler $\mathbf{DemonicH}$ for $\mathbf{DemonicE}$, with a specification for choice_A as shown in Figure 6. With this handler in our quiver, we can take $\mathbf{LangH}_{\mathbb{Z},!,\text{pick}} := \mathbf{FailH} \oplus \mathbf{HeapH}_{\text{val}} \oplus \mathbf{DemonicH}$. Simple as that, we get an extended program logic $\text{wpe} \{ \Phi \}$. Again, the wp rules in Figure 2 and Figure 4 can be carried over. There is also an additional proof

³We will explain the restriction to inhabited types in §2.4.

$$\begin{array}{c}
\text{WPI}_{\text{DEMONIC}} \\
\frac{\forall r \in A. \Phi(r)}{\text{wpi}_{\text{DemonicH}} \text{ choice}_A \{\Phi\}}
\end{array}
\qquad
\begin{array}{c}
\text{WP}_{\text{PICKINT}} \\
\frac{\forall z. \Phi(z)}{\text{wppick_int}() \{\Phi\}}
\end{array}$$

Fig. 6. Proof rules for $\text{wpi}_{\text{DemonicH}}$ and $\lambda_{\mathbb{Z},!,\text{pick}}$.

rule for choice_A , $\text{WP}_{\text{PICKINT}}$, which is a direct consequence of $\text{WPI}_{\text{DEMONIC}}$ and the general laws for wpi that we have already seen in action above.

2.4 Adequacy

We saw how to build up $\text{wpe} \{\Phi\}$ in a modular fashion. But what guarantees do we actually get from proving $\text{wpe} \{\Phi\}$? Intuitively, the answer should be “every execution of e is well-behaved and the returned value satisfies Φ ”. This is formalized by an *adequacy* (or *soundness*) theorem for the program logic. To make this precise, we have to define what an *execution* of a $\lambda_{\mathbb{Z},!,\text{pick}}$ program is and when it is *well-behaved*. For our example language, “well-behaved” will mean “terminates and does not reach fail”. (In §6.1, we will also explain how our approach can in fact deal with partial correctness where non-terminating programs are also considered “well-behaved”).

Execution via relational interpretation. A large part of what makes up an “execution” is already defined by $\llbracket _ \rrbracket$, expressed as a shallow embedding in the meta-logic. However, this denotation does not ascribe a computational meaning to events.

Previous work on ITrees used the concept of *interpretation* to give meaning to events. These interpretations are based on the idea of an effect handler turning events into monadic operations, and then “lifting” that transformation to an entire ITree.

Using this approach, one can obtain an *interpretation function* for **FailE** events:

$$f_{\text{FailE}} : \text{itree} (\text{FailE} \oplus E) R \rightarrow \text{itree} E (\text{val} \cup \{\perp_{\text{fail}}\})$$

f_{FailE} transforms an ITree by removing the **FailE** events and returning \perp_{fail} whenever fail is encountered. By applying this function to $\llbracket e \rrbracket$, we obtain $t_1 : \text{itree} (\text{HeapE}_{\text{val}} \oplus \text{DemonicE}) (\text{val} \cup \{\perp_{\text{fail}}\})$.

Next, we can interpret **HeapE_V** events by threading the state through the computation, starting at the empty heap. This defines for any V, E, R an interpretation function with the following signature:

$$f_{\text{HeapE}} : \text{itree} (\text{HeapE}_V \oplus E) R \rightarrow \text{itree} E R$$

(To simplify the signature, this interpretation function discards the final state.) Applying this to t_1 , we thus obtain $t_2 : \text{itree} \text{DemonicE} (\text{val} \cup \{\perp_{\text{fail}}\})$.

This leaves us with executing the **DemonicE** events. Interpreting non-determinism with an interpretation *function* fails to capture that there is not just a single way to execute choice_A : the entire point is that there is a possible execution for each $a \in A$! This is where we use the concept of a propositional interpretation [49]. Instead of a function, we give an *interpretation relation* to characterize possible executions:

$$\Downarrow_{\text{DemonicE}} : \text{itree} (\text{DemonicE} \oplus E) R \rightarrow \text{itree} E R \rightarrow \text{Prop}$$

Concretely, $\Downarrow_{\text{DemonicE}}$ is defined as a coinductive relation according to the rules in Figure 7. The most interesting rule is **DEMONICIRELCHOICE**: It says that to construct the interpretation for a demonic choice event over A , one has to pick an $a : A$ and then construct the interpretation for the continuation $k(a)$. The rules **DEMONICIRELRET**, **DEMONICIRELTAU**, and **DEMONICIRELVIS** simply forward the actions that are not related to demonic choice. Thus, the different ways to construct the interpretation correspond to the different demonic choices in the ITree while leaving all

$$\begin{array}{c}
 \text{DEMONICIRELCHOICE} \\
 \frac{k(a) \downarrow_{\text{DemonicE}} t'}{\text{Vis}_A(\text{choice}_A, k) \downarrow_{\text{Tau}}(t')} \\
 \\
 \text{DEMONICIRELRET} \\
 \text{Ret}(r) \downarrow_{\text{DemonicE}} \text{Ret}(r) \\
 \\
 \text{DEMONICIRELTAU} \\
 \frac{t \downarrow_{\text{DemonicE}} t'}{\text{Tau}(t) \downarrow_{\text{DemonicE}} \text{Tau}(t')} \\
 \\
 \text{DEMONICIRELVIS} \\
 \frac{\epsilon \notin \text{DemonicE} \quad \forall a. k(a) \downarrow_{\text{DemonicE}} k'(a)}{\text{Vis}_A(\epsilon, k) \downarrow_{\text{DemonicE}} \text{Vis}_A(\epsilon, k')}
 \end{array}$$

Fig. 7. Definition of $\downarrow_{\text{DemonicE}}$.

other events unchanged. Using $\downarrow_{\text{DemonicE}}$, we can turn t_2 into a set of possible interpretations $t' : \mathbf{itree} \emptyset (\text{val} \cup \{\perp_{\text{fail}}\})$.

At this point, there are no more events left: t' is either an infinite loop, or it terminates with some value in $\text{val} \cup \{\perp_{\text{fail}}\}$. We can hence say that t' is a possible execution of the original ITree $\llbracket e \rrbracket$ and therefore of e .

For notational consistency, we view the interpretation functions f_{FailE} and f_{HeapE} as relations as well (i.e., $t \downarrow_E t' \iff f_E(t) = t'$). This lets us define the executions of $t : \mathbf{itree} \text{LangE}_{\mathbb{Z}, !, \text{pick}} R$ as the set of $t' : \mathbf{itree} \emptyset (R \cup \{\perp_{\text{fail}}\})$ that satisfy the following composite interpretation relation:

$$t \downarrow_{\mathbb{Z}, !, \text{pick}} t' := \exists t_1, t_2. t \downarrow_{\text{FailE}} t_1 \downarrow_{\text{HeapE}} t_2 \downarrow_{\text{DemonicE}} t'$$

Proving adequacy effect-by-effect. Having defined our notion of execution, we turn towards what wp has to say about them. As with everything else, we follow a modular approach. We show adequacy of the underlying wpi for one effect type at a time, and then compose those reusable results to obtain adequacy for wp.

The adequacy theorems for the effects we have seen so far are stated in Figure 8. They all take basically the same shape: if $t \downarrow t'$, then $\text{wpi } t \{\Phi\}$ implies $\text{wpi } t' \{\Phi'\}$, showing that every property provable via wpi is preserved under all possible interpretations. The postcondition Φ' remains entirely unchanged for **HEAPADEQUATE** and **DEMONICADEQUATE**, but for **FAILADEQUATE** we have to adjust the postcondition to say that the program will never fail. Finally, **EMPTYADEQUATE** says that proving wpi for an ITree with no effects establishes total correctness: the ITree is equivalent to one that immediately returns a return value r that satisfies the postcondition. Note that the conclusion of these rules is still an Iris proposition, but in the case that Φ is a pure, meta-level predicate, we can use the soundness theorem of Iris to obtain a theorem that lives entirely in the meta-logic without having to trust Iris.

We can compose these adequacy theorems to show that each of the 3 stages (t_1, t_2, t') of an execution $\llbracket e \rrbracket \downarrow_{\mathbb{Z}, !, \text{pick}} t'$ preserves the weakest precondition:

$$\begin{array}{ccc}
 \llbracket e \rrbracket : \mathbf{itree} (\text{FailE} \oplus \text{HeapE}_{\text{val}} \oplus \text{DemonicE}) \text{ val} & & \text{wpi } \llbracket e \rrbracket \{\Phi\} \\
 \downarrow \downarrow_{\text{FailE}} & & \downarrow * \\
 t_1 : \mathbf{itree} (\text{HeapE}_{\text{val}} \oplus \text{DemonicE}) (\text{val} \cup \{\perp_{\text{fail}}\}) & & \text{wpi } t_1 \{v. v \neq \perp_{\text{fail}} * \Phi(v)\} \\
 \downarrow \downarrow_{\text{HeapE}} & & \downarrow * \\
 t_2 : \mathbf{itree} \text{DemonicE} (\text{val} \cup \{\perp_{\text{fail}}\}) & & \text{wpi } t_2 \{v. v \neq \perp_{\text{fail}} * \Phi(v)\} \\
 \downarrow \downarrow_{\text{DemonicE}} & & \downarrow * \\
 t' \approx \text{Ret}(v) : \mathbf{itree} \emptyset (\text{val} \cup \{\perp_{\text{fail}}\}) & & v \neq \perp_{\text{fail}} * \Phi(v)
 \end{array}$$

$$\begin{array}{c}
\text{FAILADEQUATE} \\
\frac{t \downarrow_{\text{FailE}} t' \quad \text{wpi}_{\text{FailH} \oplus H} t \{\Phi\}}{\text{wpi}_H t' \{v. v \neq \perp_{\text{fail}} * \Phi(v)\}}
\end{array}
\qquad
\begin{array}{c}
\text{DEMONICADEQUATE} \\
\frac{t \downarrow_{\text{DemonicE}} t' \quad \text{wpi}_{\text{DemonicH} \oplus H} t \{\Phi\}}{\text{wpi}_H t' \{\Phi\}}
\end{array}$$

$$\begin{array}{c}
\text{HEAPADEQUATE} \\
\frac{t \downarrow_{\text{HeapE}} t' \quad \text{wpi}_{\text{HeapH}_V \oplus H} t \{\Phi\}}{\text{wpi}_H t' \{\Phi\}}
\end{array}
\qquad
\begin{array}{c}
\text{EMPTYADEQUATE} \\
\frac{\text{wpi}_\emptyset t \{\Phi\}}{\exists r. t \approx \text{Ret}(r) * \Phi(r)}
\end{array}
\qquad
\begin{array}{c}
\text{LANGADEQUATE} \\
\frac{\llbracket e \rrbracket \downarrow_{\mathbb{Z}, \text{!pick}} t' \quad \text{wp } e \{\Phi\}}{\exists r \neq \perp_{\text{fail}}. t' \approx \text{Ret}(r) * \Phi(r)}
\end{array}$$

Fig. 8. A selection of adequacy theorems.

This chain is summarized by the adequacy theorem **LANGADEQUATE** for our program logic, formalizing the intuitive reading of $\text{wp } e \{\Phi\}$ from the beginning of this subsection.

This proof is entirely compositional and factors into intermediate stages, each focusing on one effect at a time. Just like we constructed the program logic from smaller building blocks, this confers an advantage of reusability: if one wants to derive a program logic for a language with more kinds of effects, one can define these new effects and prove adequacy theorems for them, and then use these together with the reusable components our library provides without having to monolithically improve the entire logic to be adequate.

Interpreter soundness proof. As already mentioned, ITrees come with a concept of interpretation functions that make events executable. For non-deterministic choice, there is more than one possible execution, and thus we considered not an interpretation function but an interpretation relation. However, it is still possible to define an interpretation function f_{DemonicE} that computes a legal instantiation of demonic choice: $t \downarrow_{\text{DemonicE}} f_{\text{DemonicE}}(t)$. This crucially relies on the constraint that choice_A can only be used for *inhabited* types A . Following standard ITree patterns, we can compose these functions for all our effects into a single end-to-end interpreter $f_{\mathbb{Z}, \text{!pick}}$ that can execute $\lambda_{\mathbb{Z}, \text{!pick}}$ programs. Thanks to the ITree-based formulation of **LANGADEQUATE**, one can easily show the following soundness property: if $\vdash \text{wp } e \{\Phi\}$ (where Φ is a pure proposition), the interpreter applied to e will terminate in a value $v \neq \perp_{\text{fail}}$ satisfying $\Phi(v)$. This guarantees not only safety but also termination for our interpreter.

3 Weakest Preconditions for ITrees and Logical Effect Handlers

We have seen the high-level idea of how a language-specific program logic can be defined in terms of a general weakest precondition connective for ITrees, wpi , alongside a menu of reusable effect libraries. We also saw some examples of such effect libraries. In this section, we show how wpi and these effect libraries are defined. We start with a definition of wpi in §3.1 (although this definition will later be extended in §4). We show how to define logical effect handlers with the simple examples of **FailH** and **DemonicH** (§3.2), and then discuss the more complicated handler **HeapH** (§3.3).

3.1 Defining wpi_H

We can now define the weakest precondition $\text{wpi}_H t \{\Phi\}$ for ITrees at the core of our theory:

$$\text{wpi}_H t \{\Phi\} := \text{!} \begin{cases} \Phi(r) & \text{if } t = \text{Ret}(r) \\ \text{wpi}_H t' \{\Phi\} & \text{if } t = \text{Tau}(t') \\ H_A(\epsilon, (\lambda a. \text{wpi}_H k(a) \{\Phi\})) & \text{if } t = \text{Vis}_A(\epsilon, k) \end{cases}$$

The first two cases are straightforward: $\text{Ret}(r)$ asserts the postcondition Φ and $\text{Tau}(t')$ continues with t' . The more interesting case is the one for events, $\text{Vis}_A(\epsilon, k)$, which is deferred to a *logical effect*

handler: $H(\epsilon, \Psi)$ determines the verification condition for each event ϵ in E and logical continuation Ψ in predicate transformer style. The logical continuation will be the weakest precondition for the ITree continuation k . If the event ϵ with answer type A has precondition P and postcondition $Q(a)$ for answer a , we would set $H(\epsilon, \Psi) := P * (\forall a : A. Q(a) * \Psi(a))$. In other words, triggering the event requires first proving P and then proving the logical continuation assuming Q . (We will discuss handlers in more detail in the next subsection. We will also adapt wpi to support concurrency in §4.)

All of this is wrapped in the Iris *update modality* $\text{I} \Rightarrow$ [21]. Intuitively, the update modality allows one to perform standard Iris reasoning for updating ghost state when proving a wpi .

All the rules for wpi_H in Figure 3 (on page 6) hold for free, or rather, for cheap: we only need a single property of H . Namely, we require that handlers satisfy *monotonicity*. For all events ϵ with answer type A , and all $\Phi, \Psi : A \rightarrow \text{iProp}$, the following must hold:

$$(\forall a. \Phi(a) * \Psi(a)) * H_A(\epsilon, \Phi) * H_A(\epsilon, \Psi) \quad (\text{HANDLERMONO})$$

This property trivially holds for all handlers presented in this paper.

Thanks to this monotonicity property, the recursive definition of wpi_H is well-formed by taking the least fixpoint. Choosing the least fixpoint as opposed the greatest fixpoint means that (by default) our weakest precondition is *termination sensitive* or *total*, that is, it implies termination of programs. However, we shall see *later* in §6.1 how this definition also subsumes termination insensitive reasoning, thus uniting both total and partial verification in one, common framework.

3.2 Logical Effect Handlers: Failure, Non-Deterministic Choice

We now have a more in-depth look at handlers H . Handlers codify what one needs to prove when verifying an event, *i.e.*, we have:

$$H(\epsilon, \Phi) \vdash \text{wpi}_H \epsilon \{ \Phi \}$$

(On the right hand side, the event ϵ is implicitly coerced to an ITree as described in §2.1.)

We define a corresponding handler for each event type. For example, for the **FailE** and **DemonicE** events (introduced in §2.1 and §2.3), the handlers take the following shape:

$$\text{FailH}_\emptyset(\text{fail}, \Phi) := \text{False} \quad \text{DemonicH}_A(\text{choice}_A, \Phi) := \forall a. \Phi(a)$$

FailE has a single event `fail` that we want to prove never happens. We can encode this by defining **FailH**(`fail`, Φ) as `False` and thus ensuring that we can never prove $\text{wpi}_H \text{fail} \{ \Phi \}$. For demonic choice over a type A (given by the `choiceA` event), we want to verify that the program is correct *for all* possible choices. We encode this by using a universal quantifier in the handler **DemonicH**. From this, we obtain the rule **WPIDEMONIC** in §2.3.

Composing handlers. To define a handler for a programming language with lots of different effects, we compose handlers for the individual effects. If H_1 is a handler for E_1 and H_2 is a handler for E_2 , we can define a handler $H_1 \oplus H_2$ for the sum $E_1 \oplus E_2$ in the obvious way.

Subsumption. A program using only a subset of the available effects can be verified in the corresponding fragment of the program logic. Suppose H is a handler for E and H' is a handler for E' such that E' is contained in E . We write $H' \subseteq H$ if $H(\epsilon, \Phi) \dashv\vdash H'(\epsilon, \Phi)$ for every event ϵ in E' and every Φ . This gives rise to the rule:

$$\text{WPI}_{\text{SUBSUME}} \frac{E' \subseteq E \quad H' \subseteq H \quad t : \text{itree } E' R}{\text{wpi}_H t \{ \Phi \} \dashv\vdash \text{wpi}_{H'} t \{ \Phi \}}$$

The most salient examples are $H_1 \subseteq H_1 \oplus H_2$ and $H_2 \subseteq H_1 \oplus H_2$ for any handlers H_1, H_2 . For instance, proof rules from $\text{wpi}_{\text{DemonicH}}$ thus lift to proof rules for $\text{wpi}_{\text{DemonicH} \oplus \dots}$.

3.3 Handling Mutable State

In §2.2, we introduced the \mathbf{HeapE}_V event type for mutable heaps. This is in fact a derived construction, based on the \mathbf{StateE}_S event type for global state of type S . \mathbf{HeapE}_V is defined as $\mathbf{StateE}_{\mathbb{N} \xrightarrow{\text{fin}} V}$.

The \mathbf{StateE}_S event type is a standard ITree construction. It consists of two events: `get`, with answer type S , returns the current state; `put(s)`, with answer type $()$, overwrites the current state.

The handler for \mathbf{StateE}_S follows the usual Iris recipe for dealing with global state: it is parameterized by a *state interpretation* $S : S \rightarrow iProp$ which is used to relate the physical state to Iris’s logical state. Intuitively, $S(s)$ says “we own the state and it is currently s ”.

Based on the state interpretation, we define the state handler as follows:

$$\begin{aligned} \mathbf{StateH}_S^S(\text{get}, \Phi) &:= \forall s. S(s) * \models (S(s) * \Phi(s)) \\ \mathbf{StateH}_S^S(\text{put}(s'), \Phi) &:= \forall s. S(s) * \models (S(s') * \Phi(s)) \end{aligned}$$

The handler for `put` says that to prove $\text{wpi}_{\mathbf{StateH}_S^S, \mathcal{E}} \text{put}(s') \{\Phi\}$, we can briefly take ownership $S(s)$ of the old state, and then we have to give back ownership $S(s')$ of the new state and establish $\Phi(s)$. (We will get back to the update modality \models shortly.) The handler for `get` is similar, except that the state interpretation has to be given back unchanged.

Building HeapE on top of StateE. On top of `get` and `put`, we can define heap operations:

$$\begin{aligned} \text{load}(\ell) &:= \sigma \leftarrow \text{get}; \text{Ret}(\sigma(\ell)) && : \mathbf{itree} \mathbf{HeapE}_V \text{ (option } V) \\ \text{store}(\ell, v) &:= \sigma \leftarrow \text{get}; \text{put}(\sigma[\ell := \text{some}(v)]); \text{Ret}(\sigma(\ell)) && : \mathbf{itree} \mathbf{HeapE}_V \text{ (option } V) \\ \text{alloc}(v) &:= \sigma \leftarrow \text{get}; \ell := \text{find_free}(\sigma); \text{put}(\sigma[\ell := \text{some}(v)]); \text{Ret}(\ell) && : \mathbf{itree} \mathbf{HeapE}_V \mathbb{N} \end{aligned}$$

To derive the rules for these constructs in Figure 5 (on page 7) from the basic rules for \mathbf{StateE} , we follow basically the same recipe as the standard Iris program logic [24]: we set up ghost state that tracks the current contents of physical state in S (this is why we need an update modality in the state handler) and use the same ghost state to give meaning to $\ell \mapsto v$. The one technical wrinkle is that we have to put a third view of this ghost state into a shared invariant to permit the proof to “remember” facts about the global state in between invocations of state operations.⁴ The use of an invariant gives rise to a technical side-condition, requiring the namespace of this invariant to be in the mask for each heap operation. Up to such venial side-conditions, this construction lets us then derive the desired rules in Figure 5.

Similarly, the interpretation relation $\downarrow_{\mathbf{HeapE}}$ and the adequacy theorem $\mathbf{HEAPADEQUATE}$ in Figure 8 (on page 10) can be derived from a general interpretation relation and associated adequacy theorem for \mathbf{StateH}_S^S , which we omit for lack of space.

4 Extension: Concurrency

Having pinned down basic features such as program failure and state, we turn our attention to extending the theory with a more advanced feature: concurrency. Modeling and reasoning about concurrency is yet another reusable component in our theory. However, it is sufficiently non-standard that we have to generalize our notion of wpi to accommodate the new kinds of control-flow that can arise with concurrency. To demonstrate this, let us return to $\lambda_{\mathbb{Z}, !, \text{pick}}$ from §2 and consider another language extension, arriving at $\lambda_{\mathbb{Z}, !, \text{pick}, \text{spawn}}$. The syntax becomes:

$$e \in \text{expr} ::= \dots \mid \text{spawn } \{e\}$$

⁴For the Iris experts: we use the “authoritative” construction. Its authoritative part supports fractional permissions, so we can put one half in S and one half in a global invariant. The points-to connective is defined, as usual, as a fragment of the same ghost state. The invariant itself is then also made part of S so that it does not have to be explicitly threaded through.

$\text{spawn } \{e\}$ represents spawning a new thread executing e .

4.1 Denoting a Concurrent Language into ITrees

We saw in §2.3 how to denote $\lambda_{\mathbb{Z},!,\text{pick}}$ into ITrees. We now show how to denote the extended $\lambda_{\mathbb{Z},!,\text{pick},\text{spawn}}$ into ITrees. Since we are adding a new effect (concurrency), we will need to extend the set of events **LangE**:

$$\mathbf{LangE}_{\mathbb{Z},!,\text{pick},\text{spawn}} := \mathbf{ConcE} \oplus \mathbf{FailE} \oplus \mathbf{HeapE}_{\text{val}} \oplus \mathbf{DemonicE}$$

While our language has preemptive concurrency, we model it using *cooperative concurrency* on the level of ITrees. Specifically, the event type **ConcE** unlocks the following new operations:

- (1) $\text{spawn} : \mathbf{itree } \mathbf{LangE}_{\mathbb{Z},!,\text{pick},\text{spawn}} () \rightarrow \mathbf{itree } \mathbf{LangE}_{\mathbb{Z},!,\text{pick},\text{spawn}} ()$ spawns a new thread executing some ITree.
- (2) $\text{yield} : \mathbf{itree } \mathbf{LangE}_{\mathbb{Z},!,\text{pick},\text{spawn}} ()$ yields control to an arbitrary thread in the thread pool (including possibly the current thread).

With this, we can extend our semantic interpretation with a denotation for spawn (using 0 as return value since our language does not have a “unit” value):

$$\llbracket \text{spawn } \{e\} \rrbracket := \text{spawn}(\llbracket e \rrbracket; \text{Ret}(()); \text{Ret}(0))$$

But this is not enough. We must also augment the denotation of other program terms to insert a yield at any point when control may be passed to another thread. It is necessary to exercise some care in doing so to ensure that we model the preemptive semantics in the intended way:

- (1) We want to yield “between” any two computation steps. For instance, when evaluating $! \ell + ! \ell$, it is crucial that we yield in between the two loads.
- (2) However, some expressions such as $\ell \leftarrow v$ are *atomic* which means they should execute “in a single step” without interleaving with other threads. We do not want yields in the denotations of such expressions.
- (3) Expressions such as $e_1(e_2)$ and $\text{if } e_1 \text{ then } e_2 \text{ else } e_3$ that evaluate by first transforming to another expression e' should yield before continuing with computing e' . For example, β -reductions may not terminate and we do not want one thread to block the thread pool.

An excerpt of the placement of yields in $\llbracket _ \rrbracket$ according to these considerations is displayed in Figure 9. To honor (1), we define a notational shorthand $\llbracket e \rrbracket_{\text{yield}}$ that places a yield after the evaluation of e ; we use this to evaluate the subexpressions of an expression (except for the branches of an if-expression and the argument to $\text{spawn } \{ _ \}$, which are not eagerly evaluated). However, to honor (2), we make use of a helper function $\text{yield_if_not_val}(e)$ which exhibits a yield only if the expression e is not a value. Finally, to honor (3), we ensure that denotations $\llbracket e \rrbracket$ that end on a recursive instance $\llbracket e' \rrbracket$ have a $\text{yield_if_not_val}(e')$ before this instance to mark the computational step from e to e' .

Example: Compare-and-store. To explain why we are using cooperative concurrency in the denotation, we consider adding an atomic compare-and-store operation to $\lambda_{\mathbb{Z},!,\text{pick},\text{spawn}}$:

$$e \in \text{expr} ::= \dots \mid \text{CAS}(e, e_1, e_2)$$

$\text{CAS}(\ell, v_1, v_2)$ loads the current value from ℓ , compares it to v_1 , and stores v_2 at ℓ if the two values are equal. Crucially, this operation should happen in a single atomic step: No other thread should execute during the load-compare-store sequence performed by the CAS. With our yield -based

$$\begin{aligned}
\text{yield_if_not_val}(e) &:= \text{match } e \text{ with } v \implies \text{Ret}(\text{()}) \mid _ \implies \text{yield end} \\
\llbracket e \rrbracket_{\text{yield}} &:= v \leftarrow \llbracket e \rrbracket; \text{yield_if_not_val}(e); \text{Ret}(v) \\
\llbracket e_1 \leftarrow e_2 \rrbracket &:= v_1 \leftarrow \llbracket e_1 \rrbracket_{\text{yield}}; v_2 \leftarrow \llbracket e_2 \rrbracket_{\text{yield}}; \ell \leftarrow \text{to_loc}(v_2); \\
&\quad w \leftarrow \text{store}(\ell, v_1); \text{unwrap}(w) \\
\llbracket e_1(e_2) \rrbracket &:= v_1 \leftarrow \llbracket e_1 \rrbracket_{\text{yield}}; v_2 \leftarrow \llbracket e_2 \rrbracket_{\text{yield}}; (x, e) \leftarrow \text{to_lam}(v_1); \\
&\quad \text{yield_if_not_val}(e[v_2/x]); \llbracket e[v_2/x] \rrbracket \\
\llbracket \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \rrbracket &:= v_1 \leftarrow \llbracket e_1 \rrbracket_{\text{yield}}; z_1 \leftarrow \text{to_int}(v_1); \\
&\quad \text{if } z_1 \neq 0 \text{ then } (\text{yield_if_not_val}(e_2); \llbracket e_2 \rrbracket) \\
&\quad \text{else } (\text{yield_if_not_val}(e_3); \llbracket e_3 \rrbracket)
\end{aligned}$$

Fig. 9. Excerpt of the placement of yields.

cooperative concurrency, we can express the desired semantics of CAS as the following ITree:

$$\begin{aligned}
\llbracket \text{CAS}(e, e_1, e_2) \rrbracket &:= v_1 \leftarrow \llbracket e_1 \rrbracket_{\text{yield}}; v_2 \leftarrow \llbracket e_2 \rrbracket_{\text{yield}}; v \leftarrow \llbracket e \rrbracket_{\text{yield}}; \ell \leftarrow \text{to_loc}(v); \\
&\quad v'_? \leftarrow \text{load}(\ell); v' \leftarrow \text{unwrap}(v'_?); \\
&\quad \text{if } v_1 = v' \text{ then } (\text{store}(\ell, v_2); \text{Ret}(\text{true})) \text{ else } \text{Ret}(\text{false})
\end{aligned}$$

First, this ITree evaluates the subexpressions of CAS. Then the load-compare-store sequence is expressed using the basic load and store operations provided by the **HeapE** event type. We can reuse their specifications (Figure 5) when establishing proof rules such as those for CAS:

$$\begin{array}{c}
\text{HLWP} \text{CAS} \text{SUC} \\
\hline
\ell \mapsto v \\
\hline
\text{wp}_{\mathcal{E}} \text{CAS}(\ell, v, w) \{v'. v' = \text{true} * \ell \mapsto w\}
\end{array}
\qquad
\begin{array}{c}
\text{HLWP} \text{CAS} \text{FAIL} \\
\hline
v \neq v' \quad \ell \mapsto v \\
\hline
\text{wp}_{\mathcal{E}} \text{CAS}(\ell, v', w) \{v'. v' = \text{false} * \ell \mapsto v\}
\end{array}$$

This is in stark contrast to an operational semantics, where the proof of the CAS rules typically has to duplicate the reasoning performed in the rules for loads and stores. In an attempt to overcome this, one could try to define CAS inside the language instead of having it as a primitive operation:

$$\text{CAS}(x, a, b) := \text{if } a = !x \text{ then } (x \leftarrow b; \text{true}) \text{ else false}$$

However, this is not equivalent to the intended definition: CAS is supposed to be atomic whereas this definition is not; other threads could take steps between the load and the store. With cooperative concurrency, on the other hand, we have a level of abstractions below that of our language where we are able to define the helper functions load and store with no yields so that no additional interleavings are introduced. This reuse simplifies both the language semantics and the correctness proof for the program logic.

4.2 Program Logic for Concurrent Programs

With this denotation at our disposal, we can begin to build a program logic for $\lambda_{\mathbb{Z},!,\text{pick},\text{spawn}}$ in the style of §2. This requires generalizing wpi, so we first talk about how Iris deals with concurrency.

Intermezzo: Concurrency in Iris. As usual for concurrent separation logics [30], the case of disjoint concurrency is handled via a separating conjunction: if the forked-off thread has precondition P , then the parent thread needs to prove $P * Q$ where P is handed to the forked-off thread and only Q remains in the parent thread.

$$\begin{array}{c}
\text{WPI}_{\text{SPAWN}} \\
\frac{\text{wpi}_{\text{ConcH}\oplus H;\top} t \{ \text{True} \} * \Phi(\cdot)}{\text{wpi}_{\text{ConcH}\oplus H;\mathcal{E}} \text{spawn}(t) \{ \Phi \}}
\end{array}
\qquad
\begin{array}{c}
\text{WPI}_{\text{YIELD}} \\
\frac{\Phi(\cdot)}{\text{wpi}_{\text{ConcH};\top} \text{yield} \{ \Phi \}}
\end{array}$$

Fig. 10. Weakest precondition proof rules for lTrees with concurrency.

But what if the two threads are sharing state? The answer to that is to use an *invariant*: the logical assertion $\boxed{P}^{\mathcal{N}}$ expresses that P is permanently maintained as an invariant on the shared state. Threads can get access to P *atomically*, for an instant in time, with a proof rule like this:⁵

$$\frac{P * \text{wp}_{\mathcal{E}\setminus\mathcal{N}} e \{ r. P * \Phi(r) \} \quad \boxed{P}^{\mathcal{N}} \quad \mathcal{N} \subseteq \mathcal{E} \quad e \text{ is atomic}}{\text{wp}_{\mathcal{E}} e \{ \Phi \}}$$

Ignoring the \mathcal{E} and \mathcal{N} for now, this rule expresses that to prove correctness of e , we may instead prove $P * \dots$, *i.e.*, P is made available as an extra assumption. This proof can temporarily break the invariant, but when e finishes execution, we have to re-establish P , thus ensuring that the invariant is maintained again. Crucially, since e is an atomic expression, no other thread can notice that the invariants was temporarily broken—it always holds between any two atomic steps of the program.

However, atomicity alone is not sufficient to ensure soundness of this rule. The other potential problem is *reentrancy*: if the rule could be used twice on the same invariant, the program would gain access to $P * P$, and that would be unsound. This is where the *mask* \mathcal{E} comes in: every invariant lives in a *namespace* \mathcal{N} , and the weakest precondition connective keeps track of which invariants are still available in its mask. To open an invariant, the entire namespace must be in the mask ($\mathcal{N} \subseteq \mathcal{E}$), and it is subsequently removed from the mask (e is verified with the reduced mask $\mathcal{E} \setminus \mathcal{N}$).

wpi rules for ConcE. To reason about concurrent programs with wpi, we introduce $\text{wpi}_{H;\mathcal{E}} t \{ \Phi \}$ which carries a mask to indicate which Iris invariants are currently available. We will omit the mask when it is not relevant for the current discussion; in that case the mask is arbitrary but fixed (*i.e.*, it must be the same for all wpi in a rule or theorem statement). In fact, this also applies to all the wpi proof rules shown so far.⁶

Let us see how these masks appear in the rules for **ConcE**. Our library for **ConcH** provides the rules in Figure 10 for spawn and yield. Both of these rules are more subtle than meets the eye. Unlike other rules stated so far, **WPI_{SPAWN}** is stated over an extended handler **ConcH** $\oplus H$ to allow the spawned thread t to also exhibit events that are not related to concurrency. The separating conjunction in the premise expresses that the parent thread needs to split its resources in two *disjoint* parts: one is passed to the new thread t , the other is used to prove the postcondition Φ . Remember that after using “bind”-like rules, Φ will typically itself be a weakest precondition, capturing the verification condition for the continuation in the parent thread. That way, this separating conjunction exactly captures that parent thread and the child thread may only access disjoint state. But of course, this happens in the context of Iris, so one can use invariants to get around that. This is where the other rule comes in: **WPI_{YIELD}** only holds for the full mask \top . This forces all invariants to be closed, thus ensuring that whenever execution switches from one thread to another, the new thread can rely on all invariants being satisfied.

⁵We ignore slight technicalities that have to do with either later modalities or timeless propositions; see Jung et al. [21].

⁶For **EMPTYADEQUATE** and **LANGADEQUATE**, we add an Iris *update modality* in the conclusion that carries the mask.

$$\begin{array}{c}
\text{WPSPAWN} \\
\frac{\text{wp}_{\top} e \{\text{True}\}}{\text{wp}_{\mathcal{E}} \text{spawn} \{e\} \{v. v = \text{unit}\}}
\end{array}
\qquad
\begin{array}{c}
\text{WPIINVOPEN} \\
\frac{P * \text{wp}_{\mathcal{E} \setminus \mathcal{N}} e \{r. P * \Phi(r)\} \quad \boxed{P}^{\mathcal{N}} \quad \mathcal{N} \subseteq \mathcal{E}}{\text{wp}_{\mathcal{E}} e \{\Phi\}}
\end{array}$$

$$\begin{array}{c}
\text{WPBINDSTOREL} \\
\frac{\text{wp}_{\top} e_1 \{v_1. \text{wp}_{\top} v_1 \leftarrow e_2 \{\Phi\}\}}{\text{wp}_{\top} e_1 \leftarrow e_2 \{\Phi\}}
\end{array}
\qquad
\begin{array}{c}
\text{WPBINDSTORER} \\
\frac{\text{wp}_{\top} e_2 \{v_2. \text{wp}_{\top} v_1 \leftarrow v_2 \{\Phi\}\}}{\text{wp}_{\top} v_1 \leftarrow e_2 \{\Phi\}}
\end{array}$$

Fig. 11. Excerpt of the program logic for the $\lambda_{\mathbb{Z},!,\text{pick},\text{spawn}}$.

Furthermore, wpi permits the following rule for opening invariants:

$$\begin{array}{c}
\text{WPIINVOPEN} \\
\frac{P * \text{wpi}_{H;\mathcal{E} \setminus \mathcal{N}}^t \{r. P * \Phi(r)\} \quad \boxed{P}^{\mathcal{N}} \quad \mathcal{N} \subseteq \mathcal{E}}{\text{wpi}_{H;\mathcal{E}}^t \{\Phi\}}
\end{array}$$

The attentive reader may notice the lack of an atomicity side-condition, which could seem like it would render our rule unsound. This is, however, not the case. Yes, we can in fact open invariants around arbitrary code t , even when t exhibits yields. But in that case, we will not be able to show the $\text{wpi } t \{\dots\}$ in the assumption because, as we saw, we can only step over yield using WPIYIELD if we are at full mask \top .

Program logic. Now we can finally get back to $\lambda_{\mathbb{Z},!,\text{pick},\text{spawn}}$. The program logic for our concurrent language is defined largely as before, with a new handler **ConcH** for the **ConcE** events—and with a mask, to handle opening and closing of invariants:

$$\begin{aligned}
\text{LangH}_{\mathbb{Z},!,\text{pick},\text{spawn}} &:= \text{ConcH} \oplus \text{FailH} \oplus \text{HeapH}_{\text{val}} \oplus \text{DemonicH} \\
\text{wp}_{\mathcal{E}} e \{\Phi\} &:= \text{wpi}_{\text{LangH}_{\mathbb{Z},!,\text{pick},\text{spawn}};\mathcal{E}} \llbracket e \rrbracket \{\Phi\}
\end{aligned}$$

After replacing each wp by $\text{wp}_{\mathcal{E}}$, the extended logic continues to satisfy the rules seen so far: the rules in Figure 2, the rules in Figure 4, and the rule WPICKINT . Additionally, it satisfies a number of new proof rules, some of which are shown in Figure 11. Most rules, such as WPPLUS and WPLoad , support an arbitrary mask. However, the “bind” rules (e.g., WPBINDSTOREL and WPBINDSTORER) as well as the rule for β -reduction (WPAPP) and reduction of if-expressions (WPIFTRUE and WPIFFALSE) require all invariants to be closed. The reason for this is that these operations contain a yield, so WPIYIELD forces the mask to be \top .

In other words, while we can apply WPIINVOPEN around expressions like $\ell_1 \leftarrow v_1; \ell_2 \leftarrow v_2$, apparently treating the non-atomic sequence of two stores as “atomic”, we cannot complete that proof because we would have to apply WPAPP to reduce away the semicolon (which desugars to a λ -abstraction in the usual way). Instead of the typical Iris approach of ensuring atomicity up-front via a side-condition in WPIINVOPEN , atomicity is ensured “semantically” by making it impossible to complete the proof when an invariant has been opened around an operation that yields.

4.3 Inside the Abstraction: Extending wpi with Support for Concurrency

We have seen, by example, how to use **ConcE** and **ConcH**, and we now show how we *define* them. As it turns out, concurrency is a sufficiently “different” effect that we have to extend our definition of handlers to support it.

Defining ConcE. Similar to Choice Trees [7], we encode cooperative concurrency in ITrees using an event type **ConcE** with three kinds of events:

- (1) fork with answer type $\{\text{cur}, \text{new}\}$. This rather unusual event causes the continuation to be run twice: the current thread continues its execution with the answer cur , and a new thread is created that continues with answer new .
- (2) yield with answer type $()$. This yields control to some running thread (which may be the current thread again).
- (3) endthread with answer type \emptyset . This safely ends the current thread and yields control to another thread.

From these ingredients, the more familiar operation that spawns a thread running t is defined by:

$$\text{spawn}(t) := x \leftarrow \text{fork}; \text{if } x = \text{cur} \text{ then } (\text{yield}; \text{Ret}(())) \text{ else } (t; \text{endthread})$$

The newly created thread runs t and then invokes endthread , thus ensuring that spawn itself only returns once (in the parent thread).

Logical effect handlers with concurrency. To obtain a program logic, we have to define a handler for these three events. Let us start by considering fork . One candidate definition would be:

$$\mathbf{ConcH}(\text{fork}, \Phi) := \Phi(\text{cur}) * \Phi(\text{new})$$

This models multithreading through the separating conjunction, as is usual in concurrent separation logic. However, this definition breaks monotonicity. Note that **HANDLERMONO** is stated as a *separation logic* version of monotonicity, *i.e.*, using magic wands. This means the implication from Φ to Ψ can only be used once, but the definition of **ConcH** above would have to use it twice. As a consequence, this definition is incompatible with fundamental rules such as the frame rule. Intuitively, the problem is that fork returns twice, so its continuation gets duplicated, which is not compatible with the basic premise of separation logic where resources can only be used once.

To overcome this, we extend the notion of handlers to account for concurrency. We add a new parameter Φ_s , the thread-spawning logical continuation, to handlers: $H_A(e, \Phi, \Phi_s)$. Here, $\Phi_s(a)$ represents the weakest precondition for spawning a new thread executing the continuation for $a : A$. All the existing, sequential handlers will simply ignore this argument. We shall impose the following extended monotonicity condition on handlers:

$$(\forall a. \Phi(a) * \Psi(a)) * \Box(\forall a. \Phi_s(a) * \Psi_s(a)) * H_A(\epsilon, \Phi, \Phi_s) * H_A(\epsilon, \Psi, \Psi_s)$$

In particular, the implication from Φ_s to Ψ_s is given under Iris' *persistence modality* \Box , which means that it can be used multiple times.

Updating wpi. We also have to update the weakest precondition for ITrees to account for this new parameter (and for masks, as discussed above), arriving at its final, actual definition:

$$\text{wpi}_H t \{\Phi\} := \text{H} \begin{cases} \Phi(r) & \text{if } t = \text{Ret}(r) \\ \text{wpi}_H t' \{\Phi\} & \text{if } t = \text{Tau}(t') \\ H_A(\epsilon, (\lambda a. \text{wpi}_H k(a) \{\Phi\}), (\lambda a. \top \text{H} \text{wpi}_H k(a) \{\text{False}\})) & \text{if } t = \text{Vis}_A(\epsilon, k) \end{cases}$$

We use postcondition False for the thread spawning continuation to enforce that new threads never end in a $\text{Ret}(r)$, *i.e.*, only the main thread can return. This is critical to ward off the issues related to fork returning twice. Essentially, this imposes a proof obligation to show that only the original thread can return. spawn 's use of endthread ensures that this is the case.

On top of this, we define wpi with a mask as follows:

$$\text{wpi}_{H;\epsilon} t \{\Phi\} := \epsilon \text{H} \text{wpi}_H t \{r. \emptyset \text{H} \Phi(r)\}$$

This definition permits all invariants in \mathcal{E} to be opened for the entire computation represented by t : $\mathcal{E} \Vdash_{\emptyset}$ is a *mask-changing update modality*, which says that starting with mask \mathcal{E} , arbitrary invariants can be opened. (Iris masks support framing, so the empty mask does not *force* all invariants to be opened.) Then, at the end of the computation $\emptyset \Vdash_{\mathcal{E}}$ demands that all the invariants in \mathcal{E} are closed before showing the postcondition $\Phi(r)$.

Defining *ConcH*, correctly. With this out of the way, we can finally define **ConcH**:

$$\mathbf{ConcH}(\text{fork}, \Phi, \Phi_s) \quad := \Phi(\text{cur}) * \Phi_s(\text{new})$$

$$\mathbf{ConcH}(\text{yield}, \Phi, \Phi_s) \quad := \emptyset \Vdash_{\top} \top \Vdash_{\emptyset} \Phi(())$$

$$\mathbf{ConcH}(\text{endthread}, \Phi, \Phi_s) := \emptyset \Vdash_{\top} \text{True}$$

Crucially, the handler for fork uses Φ only once and thus satisfies monotonicity. The handler for yield uses a mask-changing update $\emptyset \Vdash_{\top}$ to force *all* invariants to be closed, and then immediately switches back to the empty mask which lets the invariants be opened again. However, this is enough to ensure that for one instant, all invariants are satisfied, and thus we can soundly switch from one thread to another. The handler for endthread is similar, except that it never returns to the program, so after closing all invariants there is nothing left to be proven. From these handlers, standard Iris reasoning can derive the rules in [Figure 10](#) (on [page 15](#)).

4.4 Adequacy for Concurrency

Concurrency fits into the adequacy story from [§2.4](#) as yet another reusable component. Given an ITree $t : \mathbf{itree} (\mathbf{ConcE} \oplus E) R$, we specify what are the valid executions (or *interleavings*) $t' : \mathbf{itree} E R$ by means of an interpretation relation:

$$\downarrow_{\mathbf{ConcE}}: \mathbf{itree} (\mathbf{ConcE} \oplus E) R \rightarrow \mathbf{itree} E R \rightarrow \text{Prop}$$

This relation, whose definition we shall return to in a moment, satisfies an adequacy theorem like those before. Namely, the weakest precondition is preserved under any execution of **ConcE** events:

$$\frac{\text{CONCADEQUATE} \quad t \downarrow_{\mathbf{ConcE}} t' \quad \text{wpi}_{\mathbf{ConcH} \oplus H; \top} t \{ \Phi \}}{\text{wpi}_{H; \top} t' \{ \Phi \}}$$

Adequacy for $\lambda_{\mathbb{Z}, !, \text{pick}, \text{spawn}}$. The new interpretation relation for our language uses a composite relation:

$$\begin{aligned} \downarrow_{\mathbb{Z}, !, \text{pick}, \text{spawn}} &: \mathbf{itree} \mathbf{LangE}_{\mathbb{Z}, !, \text{pick}, \text{spawn}} R \rightarrow \mathbf{itree} \emptyset (R \cup \{\perp_{\text{fail}}\}) \rightarrow \text{Prop} \\ t \downarrow_{\mathbb{Z}, !, \text{pick}, \text{spawn}} t' &:= \exists t_1, t_2, t_3. t \downarrow_{\mathbf{ConcE}} t_1 \downarrow_{\mathbf{FailE}} t_2 \downarrow_{\mathbf{HeapE}} t_3 \downarrow_{\mathbf{DemonicE}} t' \end{aligned}$$

Composing **CONCADEQUATE** with the adequacy theorems in [Figure 8](#), we obtain adequacy for our language with concurrency:

$$\frac{\text{LANGADEQUATE}' \quad \llbracket e \rrbracket \downarrow_{\mathbb{Z}, !, \text{pick}, \text{spawn}} t' \quad \text{wp}_{\top} e \{ \Phi \}}{\exists r \neq \perp_{\text{fail}}. t' \approx \text{Ret}(r) * \Phi(r)}$$

Although the order of the other events did not matter, it is important that we put **ConcE** in the beginning of $\mathbf{LangE}_{\mathbb{Z}, !, \text{pick}, \text{spawn}}$. Indeed, **ConcE** is a very special kind of event whose semantics does not commute with, for example, the semantics of **HeapE**: if we interpreted **HeapE** before **ConcE** (that is, $t \downarrow_{\mathbf{HeapE}} t_1 \downarrow_{\mathbf{ConcE}} t_2 \cdots$), each thread would have its own independent copy of the heap. To have the right interactions with other events, **ConcE** must always be interpreted first.

$$\begin{array}{c}
\text{CONCIREL} \\
\frac{[t] \downarrow_{\text{ConcE}}^0 t'}{t \downarrow_{\text{ConcE}} t'} \\
\\
\text{CONCIRELVIS} \\
\frac{\epsilon \notin \text{ConcE} \quad \forall a. \text{tp}[i := k(a)] \downarrow_{\text{ConcE}}^i k'(a)}{\text{tp}[i := \text{Vis}(\epsilon, k)] \downarrow_{\text{ConcE}}^i \text{Vis}(\epsilon, k')} \\
\\
\text{CONCIRELENDTHREAD} \\
\frac{\text{delete}(i, \text{tp}) \downarrow_{\text{ConcE}}^j t'}{\text{tp}[i := \text{Vis}(\text{endthread}, k)] \downarrow_{\text{ConcE}}^i \text{Tau}(t')} \\
\\
\text{CONCIRELTAU} \\
\frac{\text{tp}[i := t] \downarrow_{\text{ConcE}}^i t'}{\text{tp}[i := \text{Tau}(t)] \downarrow_{\text{ConcE}}^i \text{Tau}(t')} \\
\\
\text{CONCIRELRET} \\
\text{tp}[i := \text{Ret}(r)] \downarrow_{\text{ConcE}}^i \text{Ret}(r) \\
\\
\text{CONCIRELYIELD} \\
\frac{\text{tp}[i := k(())] \downarrow_{\text{ConcE}}^j t'}{\text{tp}[i := \text{Vis}(\text{yield}, k)] \downarrow_{\text{ConcE}}^i \text{Tau}(t')} \\
\\
\text{CONCIRELFORK} \\
\frac{\text{tp}[i := k(\text{cur})] \# [k(\text{new})] \downarrow_{\text{ConcE}}^i t'}{\text{tp}[i := \text{Vis}(\text{fork}, k)] \downarrow_{\text{ConcE}}^i \text{Tau}(t')}
\end{array}$$

Fig. 12. Definition of $\downarrow_{\text{ConcE}}$.

The interpretation relation for ConcE. We now define the interpretation relation $\downarrow_{\text{ConcE}}$ for **ConcE**. The first step is a *thread pool evaluation relation* $\text{tp} \downarrow_{\text{ConcE}}^i t'$. Here, tp is a list of ITrees representing the current threads, and i is the index of the thread that is currently running. The relation describes the ITrees t' that can arise by interleaving thread executions in an arbitrary way.

Thread pool evaluation is defined coinductively as shown in Figure 12 (we implicitly assume that every index i, j is in-bounds). The key rules are **CONCIRELYIELD**, **CONCIRELENDTHREAD**, and **CONCIRELFORK**, which define what happens when the active thread $\text{tp}(i)$ runs one of the **ConcE** events. On a yield, we pick an arbitrary new thread j with which to continue the execution. The current thread i is updated to reflect that the yield has been executed and returned $()$. We also add a **Tau** event to t' ; this ensures that the coinduction is well-formed. On an endthread, we delete the current thread from the thread pool and continue the execution at some new thread j .⁷ And finally, fork updates the current thread i to continue with the $k(\text{cur})$ continuation and adds a new thread to the thread pool that executes the $k(\text{new})$ continuation. This is the key rule and the source of all the complications we had to deal with above since it duplicates k .

The remaining rules say that the interleaved ITree t' mirrors the behavior of the active thread: **CONCIRELTAU** forwards silent steps, **CONCIRELRET** terminates execution when the active thread reaches a **Ret**,⁸ and **CONCIRELVIS** forwards visible events. Note how the latter requires the premise to be shown for *all* possible answers a ; the “interleaving” of an ITree is not just a single execution but resolves all scheduling questions for all possible answers to uninterpreted events.

With this definition in place, we can tie it all together and prove **CONCADEQUATE**. This proof is highly non-trivial due to our use of a least fixpoint in the definition of wpi : we are showing that the per-thread termination proof that is implicit in the premise carries over to all possible interleavings.

The proof also relies on a technical side-condition omitted in the paper: the handler H for the remaining events needs to be *sequential*, which means that it must be constant in the argument Φ_s . **HEAPADEQUATE** also carries this side-condition. Aside from **CONCH**, all handlers discussed in this paper are sequential.

Interpreter. Using a round-robin scheduler, we can define an interpretation function

$$f_{\text{ConcE}} : \text{itree}(\text{ConcE} \oplus E) R \rightarrow \text{itree} E R$$

⁷We omit some technical details related to what happens if the *last* thread ends, which is something that can never happen in the languages we consider since the denotations into ITrees never put an endthread into the main thread.

⁸The denotation of a language into ITrees generally ensure that only the main thread ever reaches **Ret**; all the other threads are ended with **endthread**.

that instantiates the relation $\downarrow_{\text{ConcE}}$ in the sense that $\forall t. t \downarrow_{\text{ConcE}} f_{\text{ConcE}}(t)$. As before, this can be composed to an end-to-end interpreter for $\lambda_{\mathbb{Z}, \perp, \text{pick}, \text{spawn}}$ and a soundness result showing that if one proves $\text{wp}_{\top} e \{\Phi\}$, the interpreter applied to e will terminate in a value $v \neq \perp_{\text{fail}}$.

5 Extension: Angelic Choice and State Machine Adequacy

We now discuss a second extension that enables our approach to support *non-computational effects*. But first, let us explain what we mean by “non-computational effects” via a concrete example of such an effect: *angelic choice*.

Angelic choice. Angelic choice [12] (or angelic non-determinism) is similar to the demonic choice introduced in §2.3, except that it behaves dually: instead of having to prove correct behavior for all choices during verification, we (the verifier) can pick one of the choices. The program is correct if there exists any way to make a choice that leads to the desired outcome. Concretely, we can represent angelic choice with an event type **AngelicE** that provides an operation $\text{angelic_choice}_A : \text{itree } \mathbf{AngelicE} \ A$ with corresponding handler **AngelicH** such that we obtain the following rule:

$$\frac{\text{WPI}_{\mathbf{AngelicE}} \quad \exists r. \Phi(r)}{\text{wpi}_{\mathbf{AngelicH}} \text{angelic_choice}_A \{\Phi\}}$$

Note that this rule is the same as $\text{WPI}_{\mathbf{DemonicE}}$ from Figure 6, except that the universal quantifier in the premise is replaced by an existential quantifier.

Angelic choice is not commonly found in programming languages and cannot, in general, be compiled to executable machine code. However, angelic choice does have a wide variety of use-cases such as modelling partial programs [4] and concurrency [17, 10], reasoning about interaction with external code [34, 16], and encoding concise specifications [12, 37]. As a concrete example, Sammler et al. [34] use angelic choice to translate values between a C-like language and an assembly-like language. Here, the challenge is that all values at the assembly level are represented by integers, but values at the C level have more structure (*i.e.*, values can be integers, pointers, or boolean). Thus, when passing a value from an assembly program to a C program, one needs to recover this structure. Since there is no information in the value to do this, Sammler et al. [34] rely on angelic choice. With this approach, the verifier gets to decide whether an assembly integer corresponds to an C integer, pointer, or boolean, as long as one can prove that the remaining program is correct for this choice. (This is in contrast to demonic choice, where one would need to prove that the program is correct for all choices.) The exact same approach is used in Melocoton [16] for reasoning about translation of values between C and OCaml.

Adequacy for non-computational effects. A problem arises when we try to apply the adequacy approach from §2.4 to **AngelicE**. So far all effects were *computational*, meaning the program executions could be described by interpretation relations with corresponding interpretation functions. But angelic choice does not fit this pattern: we cannot remove angelic choice by interpretation since the witness of the angelic choice is not chosen by the interpretation but during verification.

To equip our framework with the ability to handle such non-computational effects, we introduce a novel notion of “execution” for ITrees that turns an ITree into a state machine in a modular way, side-stepping event interpretation altogether. We then prove adequacy of wpi w.r.t. that state machine. While supporting more kinds of effects, this approach loses the connection to the typical concept of event interpretation in ITrees, and in particular it does not give rise of a soundness proof relating wpi to an interpreter.

Turning ITrees into state machines. Concretely, for an ITree with events E we pick a type of states Σ and construct a state machine given by the following multi-step multi-relation:

$$(\Downarrow) : \mathbf{itree} E R \rightarrow \Sigma \rightarrow (\mathbf{itree} E R \rightarrow \Sigma \rightarrow \text{Prop}) \rightarrow \text{Prop}$$

The meaning of the execution relation $(t, \sigma) \Downarrow T$ is subtle to explain, so we first consider some examples. If t does not contain any angelic choice, then for each possible sequence of demonic choices, there will be a corresponding execution of the shape $(t, \sigma) \Downarrow \{(t', \sigma')\}$. This corresponds to a small-step reduction sequence (of 0 or more steps) between these two states. The set of outcomes T (represented by a predicate) can always be made bigger, so the full meaning of $(t, \sigma) \Downarrow T$ in the absence of angelic choice is that there exists a sequence of demonic choices giving rise to a reduction sequence starting at (t, σ) and ending in some state in T .

Conversely, if t *only* contains angelic choice, then there will be a single derivation of $(t, \sigma) \Downarrow T$ with a large T reflecting all the possible final values, plus further derivations due to closure under larger T , and derivations reflecting partial executions that do not reach a final value.

However, the key power in this relation lies in how it characterizes programs that mix demonic and angelic choice. At this point we think of demonic and angelic choice as being two opponents playing against each other, with their choices resolved by a strategy (in the game-theoretic sense). An “execution” corresponds to a strategy of the demonic player, with the set of outcomes T bounding the states that an arbitrary angelic strategy can reach against this player. As an example, consider the following ITrees, where $t_{a,d}$ first performs an angelic then a demonic choice, and $t_{d,a}$ vice versa:

$$\begin{aligned} t_{a,d} &:= a \leftarrow \text{angelic_choice}_{\mathbb{Z}}; d \leftarrow \text{choice}_{\mathbb{Z}}; \text{Ret}(a, d) \\ t_{d,a} &:= d \leftarrow \text{choice}_{\mathbb{Z}}; a \leftarrow \text{angelic_choice}_{\mathbb{Z}}; \text{Ret}(a, d) \end{aligned}$$

For $t_{a,d}$, we have $\forall f : \mathbb{Z} \rightarrow \mathbb{Z}. t_{a,d} \Downarrow \{\text{Ret}(a, f(a)) \mid a \in \mathbb{Z}\}$. This reflects the fact that for every demonic strategy f reacting to the initial angelic choice a , there is an execution that forces the return value to be of the shape $(a, f(a))$. In contrast, for $t_{d,a}$ we only have $\forall d : \mathbb{Z}. t_{d,a} \Downarrow \{\text{Ret}(a, d) \mid a \in \mathbb{Z}\}$. This reflects the fact that the demonic choice cannot depend on the angelic choice. In particular, $t_{a,d} \Downarrow \{\text{Ret}(a, a) \mid a \in \mathbb{Z}\}$ holds but the same does not hold for $t_{d,a}$.

In summary, $(t, \sigma) \Downarrow T$ means that from t with initial state σ , there exists a demonic strategy such that no matter the angelic strategy, the execution will reach a state in T . This representation of state machines with both kinds of non-determinism follows prior work [31, 34, 16, 8] (but note that some prior work swaps how angelic and demonic choice are represented).

We define \Downarrow modularly by defining a step relation $(\rightsquigarrow) : EA \rightarrow \Sigma \rightarrow (A \rightarrow \Sigma \rightarrow \text{Prop}) \rightarrow \text{Prop}$ for every event provided by the event type E . This relation shows how the operation takes a “small” step to its result, a set of the possible angelic choices. For example, the \rightsquigarrow relations for the **AngelicE**, **DemonicE**, and **HeapE_V** are given as follows:

$$\frac{\exists x. T(x, \sigma)}{(\text{choice}_A, \sigma) \rightsquigarrow T} \quad \frac{\forall x. T(x, \sigma)}{(\text{angelic_choice}_A, \sigma) \rightsquigarrow T} \quad \frac{T(\sigma[\ell], \sigma)}{(\text{load}(\ell), \sigma) \rightsquigarrow T} \quad \frac{T(\sigma[\ell], \sigma[\ell \mapsto v])}{(\text{store}(\ell, v), \sigma) \rightsquigarrow T}$$

Note how to construct an execution for demonic choice, we get to pick an answer x , just like we would when constructing a trace in a regular small-step semantics. As part of a larger derivation witnessing an execution in our multi-relation, this choice of x defines a strategy for making demonic choices. Conversely, constructing an execution for angelic choice means proving that we can react to all possible angelic answers while still remaining in the outcome set T . Operations like load and store (from §2.2) use the state σ to read resp. write the value for the given location.

We can define \rightsquigarrow for $E_1 \oplus E_2$ from \rightsquigarrow for E_1 and E_2 by building the product of states and using the step relation corresponding to the event. Overall, this allows us to automatically obtain \rightsquigarrow for a combined event type like **LangE**.

With \rightsquigarrow at hand, we define \Downarrow by coinductively lifting \rightsquigarrow to ITrees such that:

$$\frac{T(t, \sigma)}{(t, \sigma) \Downarrow T} \quad \frac{(t, \sigma) \Downarrow T}{(\text{Tau}(t), \sigma) \Downarrow T} \quad \frac{(\epsilon, \sigma) \rightsquigarrow (\lambda x, \sigma'. (k(x), \sigma') \Downarrow T)}{(\text{Vis}(\epsilon, k), \sigma) \Downarrow T}$$

Note how the first rule corresponds to terminating the current execution, but applies any time, not just when t is a Ret.

Adequacy in terms of state machines. After defining \Downarrow , we need to prove that its handling of events agrees with the handler H used by wpi_H . We encode this as a condition $\text{sound}(H, I)$ where I is a invariant on the state Σ . (The definition of $\text{sound}(H, I)$ can be found in the accompanying Coq development.) Note that this condition can be proven once and for all for each handler since it is independent of the verified ITree. We prove $\text{sound}(H, I)$ for all handlers presented in this paper and show that it lifts to $H_1 \oplus H_2$. (We even support concurrency via **ConcE** using an extended version of \Downarrow not covered in the paper.) We prove the following adequacy theorem for wpi_H :⁹

$$\text{STATEMACHINEADEQUATE} \frac{\text{sound}(H, I) \quad (t, \sigma) \Downarrow T \quad I(\sigma) * \text{wpi}_H t \{\Phi\}}{\exists t', \sigma'. T(t', \sigma') * I(\sigma') * \text{wpi}_H t' \{\Phi\}}$$

This theorem states that, for a sound handler H , we can “step in” wpi_H along a \Downarrow execution after proving I for its initial state. We obtain a result (t', σ') in the set of final states T of the execution. The invariant I holds for the final state σ' and the “remaining” ITree t' satisfies the weakest precondition wpi_H . In particular, if T ensures that t' is equal to $\text{Ret}(x)$, we obtain $\Phi(x)$. Note that this holds for every possible T (demonically), but only for one (t', σ') in T (angelically).

The fact that \Downarrow is coinductively defined makes this theorem stronger than if it were inductively defined. In particular, the premise can be established without proving that t terminates. As a consequence, this adequacy can even be used to prove termination by choosing a suitable T indicating that the original t terminates (for example, $T(t, \sigma) := \exists r. t \approx \text{Ret}(r)$).

Overall, we obtain another approach to defining the concept of “executing” an ITree and a corresponding adequacy theorem, both of which are built compositionally by combining reusable pieces for individual effects.

6 Case Study: HeapLang

In this section, we demonstrate the applicability of our approach by using it to build a program logic for HeapLang, the default language for program verification in Iris. HeapLang [21, §6.1] is an untyped lambda calculus with an ML-like higher order heap and concurrency. Our objective is to recover the original HeapLang program logic with all its basic rules [21, Figure 13] in a compositional style using the theory developed in earlier sections. The syntax of the language is:

$$v \in \text{val} ::= () \mid z \mid \text{true} \mid \text{false} \mid \ell \mid \text{rec } f(x) := e \mid \dots \quad (z \in \mathbb{Z})$$

$$e \in \text{expr} ::= v \mid x \mid e_1(e_2) \mid \text{spawn } \{e\} \mid \text{ref}(e) \mid !e \mid e_1 \leftarrow e_2 \mid \text{CAS}(e, e_1, e_2) \mid \dots$$

(Arithmetic operations and the usual operations on pairs and sums are omitted for brevity.)

Program logic for HeapLang. Following the same pattern as §2 and §4, we specify a semantics for HeapLang by denoting expressions into **itree HeapLangE** val where

$$\text{HeapLangE} := \text{ConcE} \oplus \text{FailE} \oplus \text{HeapE}_{\text{val}} \oplus \text{DemonicE}$$

For most of the expressions, the semantic interpretation $\llbracket _ \rrbracket$ is defined as for $\lambda_{\mathbb{Z}, !, \text{pick}, \text{spawn}}$ in §4, except that HeapLang uses right-to-left evaluation order and its closures have a binder f for

⁹We omit masks from this theorem to avoid clutter.

recursive calls. Also, we tweak `alloc` (see §3.3) to instead pick the free location non-deterministically (thus necessitating **DemonicE** among our events).

To build a program logic, we combine the various handlers from §3 using \oplus to obtain a handler **HeapLangH** for **HeapLangE**. As in §2, we then define $\text{wp}_{\mathcal{E}} e \{\Phi\} := \text{wpi}_{\text{HeapLangH};\mathcal{E}} \llbracket e \rrbracket \{\Phi\}$.

$\text{wp}_{\mathcal{E}} e \{\Phi\}$ satisfies the various rules that were discussed in §2. Using the notion of evaluation contexts K [21, §6.1], the bind rules such as **WPBindPlusL** can be summarized in a single rule:

$$\frac{\text{HLWPBind} \quad \text{wp}_{\top} e \{v. \text{wp}_{\top} K[v] \{\Phi\}\}}{\text{wp}_{\top} K[e] \{\Phi\}}$$

(The ramifications of the full mask \top were already discussed in §4.2.) This is a consequence of a “semantic bind lemma” which says that the *syntactic* bind $K[e]$ interprets to the *monadic* bind:

LEMMA 6.1 (SEMANTIC BIND LEMMA). *If $K \neq \bullet$ then $\llbracket K[e] \rrbracket \approx v \leftarrow \llbracket e \rrbracket_{\text{yield}}; \llbracket K[v] \rrbracket$.*

6.1 Termination-Insensitive Reasoning

So far, we have only considered a very strong kind of program logic, namely a total weakest precondition that ensures termination. This does not match the usual logic used for **HeapLang**, which just proves partial correctness. In particular, the total weakest precondition has no coinductive reasoning principles that can be used to verify, say, recursive functions without proving termination.

To address this, we also define a partial weakest precondition $\text{wp}_{\mathcal{E}}^{\triangleright} e \{\Phi\}$ that does not guarantee termination. We can use our existing framework to define this weakest precondition without having to redefine `wpi`. Iris approaches partial verification by means of the later modality \triangleright which comes with a powerful coinductive reasoning principle: *Löb induction* [21, §5.6]. Typical Iris program logics are set up such that the weakest precondition involves at least one \triangleright per program step, which can be used with Löb induction to derive the usual partial correctness reasoning principle for recursive functions. To achieve the same with our approach, we define an event type **StepE** with a single event, `step`, with answer type $()$. The handler for **StepE** is then defined as **StepH**(`step`, Φ , Φ_s) := $\triangleright \Phi()$.

We define **HeapLangE** ^{\triangleright} := **HeapLangE** \oplus **StepE**, and **HeapLangH** ^{\triangleright} := **HeapLangH** \oplus **StepH**, and finally define $\llbracket e \rrbracket^{\triangleright}$ like $\llbracket e \rrbracket$ but dredging step events at every point in $\llbracket e \rrbracket$ that corresponds to a step in the operational semantics (see the Coq code for the exact placement). The result is a partial program logic $\text{wp}_{\mathcal{E}}^{\triangleright} e \{\Phi\} := \text{wpi}_{\text{HeapLangH}^{\triangleright};\mathcal{E}} \llbracket e \rrbracket^{\triangleright} \{\Phi\}$ validating the standard **HeapLang** rules.

In our Coq mechanization, we take this one step further and make **StepH** and `wp` parametric in whether a later modality is emitted. This allows uniform treatment of partial and total correctness and their proof rules in one framework without duplicated proof effort.

Adequacy for StepH. **StepH** fits into the compositional adequacy story of §2.4 as yet another reusable piece. To describe the semantics of **StepE**, we define an interpretation relation

$$\downarrow_{\text{StepE}}^n: \text{itree}(\text{StepE} \oplus E) R \rightarrow \text{itree} E (R \cup \{\perp_{\text{step_timeout}}\}) \rightarrow \text{Prop} \quad \text{for } n \in \mathbb{N}$$

by coinduction according to the rules in Figure 13. For consistency, we write it as a relation, but it can equivalently be written as a function f_{StepE} . Intuitively, $f_{\text{StepE}}(n, t)$ executes the first n step events in t as no-ops. Once this “fuel” is used up, the next step event terminates program execution by returning $\perp_{\text{step_timeout}}$.

The adequacy theorem (**STEPADEQUATE**) turns the `wpi` of t into a `wpi` of every partial execution t' of t . It requires later credits ℓn (Spies et al. [38]) which provide the right to strip n later modalities. The Iris soundness theorem provides any fixed number of later credits, so this is sufficient to prove that for every n , if the program terminates in n steps, the postcondition holds.

$$\begin{array}{c}
\frac{n > 0 \quad k(\cdot) \Downarrow_{\text{StepE}}^{n-1} t'}{\text{Vis}(\text{step}, k) \Downarrow_{\text{StepE}}^n \text{Tau}(t')} \quad \text{Vis}(\text{step}, k) \Downarrow_{\text{StepE}}^0 \text{Ret}(\perp_{\text{step_timeout}}) \quad \text{Ret}(r) \Downarrow_{\text{StepE}}^n \text{Ret}(r) \\
\\
\frac{t \Downarrow_{\text{StepE}}^n t'}{\text{Tau}(t) \Downarrow_{\text{StepE}}^n \text{Tau}(t')} \quad \frac{\epsilon \notin \text{StepE} \quad \forall a. k(a) \Downarrow_{\text{StepE}}^n k'(a)}{\text{Vis}(\epsilon, k) \Downarrow_{\text{StepE}}^n \text{Vis}(\epsilon, k')} \\
\\
\frac{\text{STEPADEQUATE} \quad t \Downarrow_{\text{StepE}}^n t' \quad \text{wpi}_{\text{StepH}\oplus\text{H};\emptyset} t \{\Phi\} * \mathcal{L}n}{\text{wpi}_{\text{H};\emptyset} t' \{x. \text{match } x \text{ with } \perp_{\text{step_timeout}} \implies \text{True} \mid r \implies \Phi(r)\}}
\end{array}$$

Fig. 13. Rules defining the coinductive relation $\Downarrow_{\text{StepE}}^n$, and adequacy for the corresponding handler.

6.2 A Verified Interpreter

As before, we can compose the interpretation functions for the compounding events to obtain an interpreter for `HeapLang` and an associated correctness proof. Unlike the existing `HeapLang` interpreter (which only exists as an experimental feature in the Iris development repository), this does not require a full second, executable definition of the language semantics; we can just reuse the `ITree` denotation. We also obtain a stronger soundness result for the interpreter, showing in particular that if a program was proven to terminate using the total weakest precondition, then the interpreter eventually terminates.

6.3 Correctness of `ITree` Semantics w.r.t. Operational Semantics

By composing interpretation relations as in §2.4, we obtain an interpretation relation $\Downarrow_{\text{HeapLangE}^\triangleright}^{n;\sigma}$ which allows us to describe the executions of $\llbracket e \rrbracket^\triangleright$. To provide assurance that this semantics is meaningful, we prove a result relating it to the more well-established operational semantics of `HeapLang` [21, Figure 12]. First, we define two notions of “adequacy” for a program w.r.t. a postcondition, one in terms of $\Downarrow_{\text{HeapLangE}^\triangleright}$ and one in terms of operational semantics:

Definition 6.2. e is *interpretationally adequate* w.r.t. postcondition $\phi : R \rightarrow \text{Prop}$ if for every σ, n, t', r such that $\llbracket e \rrbracket^\triangleright \Downarrow_{\text{HeapLangE}^\triangleright}^{\sigma;n} t' \approx \text{Ret}(r)$, then $r \neq \perp_{\text{fail}}$ and either $r = \perp_{\text{step_timeout}}$ or $\phi(r)$.

Definition 6.3. A thread pool tp is *progressive* at heap σ if no thread is stuck: for each $e \in \text{tp}$, there is some e', σ', \vec{e}_f so that $e; \sigma \rightarrow_{\text{t}} e'; \sigma'; \vec{e}_f$.

e is *operationally adequate* w.r.t. postcondition $\phi : \text{val} \rightarrow \text{Prop}$ if

- (1) for any $\sigma, \sigma', \text{tp}'$ such that $[e]; \sigma \rightarrow_{\text{tp}}^* \text{tp}'; \sigma'$, tp' is progressive at heap σ' , and
- (2) for any $\sigma, \sigma', v, \text{tp}'$ such that $[e]; \sigma \rightarrow_{\text{tp}}^* [v] \# \text{tp}'; \sigma'$, we have $\phi(v)$.

The following chain of implications holds:

$$\text{wp}_\top^\triangleright e \{\phi\} \implies e \text{ is interpretationally adequate w.r.t. } \phi \implies e \text{ is operationally adequate w.r.t. } \phi$$

The first implication is obtained modularly by composing the adequacy theorems seen so far.

The second implication relates the `ITree` semantics and the operational semantics. For reasons of space, we cannot spell out the details here and refer the reader to our Coq formalization. The proof is a simulation involving an intermediate notion of *ITree traces* that we define in our library. For an operational semantics trace of length n starting at e , the proof constructs by induction an `ITree` trace in $\llbracket e \rrbracket^\triangleright$. The proof concludes by constructing the relational interpretation $\llbracket e \rrbracket^\triangleright \Downarrow_{\text{HeapLangE}^\triangleright}^{\sigma;n} t'$

from this ITree trace. The latter step is entirely modular and reusable: in our Coq formalization, we provide a number of composable trace lemmata that allow the user to construct relational interpretations from ITree traces.

6.4 Evaluation of the Program Logic

Having established a new program logic for HeapLang, we now turn to comparing it to the existing program logic for HeapLang and evaluating the usefulness of our approach.

Expressive power of the logic. We obtain for our program logic nearly the same rules as HeapLang’s existing program logic. The only significant difference is the fact that in our logic, invariants can be opened around any block of code, at the expense of the bind rule `HLWPBIND` needing the full mask \top ; cf. §4.2. Furthermore, while we still use the Iris proof mode [25, 23], we have only reimplemented some of the additional proof mode integration and automation available in the existing HeapLang implementation. We have also omitted support for the more recent extension of HeapLang introducing prophecy variables [22].

In particular, our total weakest precondition corresponds to the existing total weakest precondition of Iris¹⁰ and thus shares its expressive power. Concretely, it can be used to show termination, but not liveness, of concurrent programs and can use first-order invariants (thanks to Iris’ “timeless” mechanism [21, §5.7]) such as sharing a points-to assertion between threads, but not higher-order invariants (due to the lack of step-indexing).

To exemplify that our (partial) program logic gives the same expressive power as the original HeapLang program logic, we have ported the proof of a join primitive and of a higher-order lock with an impredicative invariant [39] from the original HeapLang to our logic.

Comparison of proof effort. There is a number of qualitative advantages in terms of proof effort for establishing the HeapLang program logic using our ITree-based approach compared to the standard Iris approach. For one, in the original HeapLang logic, it was necessary to state every rule twice, once for partial correctness and once for total correctness, whereas we state and prove the rules in a way that is parametric in partiality/totality (*i.e.*, whether later modalities are emitted).

Furthermore, ITrees enable reuse of abstractions in the definition of the semantics, which in turn leads to more compositional proofs. For example, the definition of $[[_]]$ relies on helper functions `load` and `store`, and their specifications can be reused in establishing the program logic. As we discussed in the compare-and-swap example in §4.1, this kind of reuse is typically not available in the setting of operational semantics.

Overall, our proofs have about the same size as the original, highly optimized proofs but require less specialized proof engineering and have reasoning that is higher-level and compositional (namely application of `wpi` lemmata rather than inverting the small-step relation).

7 Case Study: Islaris

As our second large case study, we redefine the program logic used by Islaris [32] using the approach presented in this paper. Islaris provides an Iris-based program logic for traces that describe the semantics of assembly programs based on authoritative models of real-world assembly languages like Armv8 and RISC-V.

These traces are formally described by the *Isla trace language (ITL)* shown in Figure 14. An ITL “program” is given by a trace t , which correspond to the SMTLIB-traces generated by the Isla tool [3] from partially evaluating the ISA models. Traces consist of events j (not to be confused with the

¹⁰This part of Iris was never described in a paper; it can be found at https://gitlab.mpi-sws.org/iris/iris/-/blob/iris-4.2.0/iris/program_logic/total_weakestpre.v.

$$\begin{aligned}
\tau &::= \text{BitVec}(n) \mid \text{Boolean} \mid \dots & e &::= v \mid \text{not}(e) \mid \text{bvadd}(e_1, e_2) \mid \dots \\
j &::= \text{ReadReg}(r, v) \mid \text{WriteReg}(r, v) \mid \text{ReadMem}(v_d, v_a, n) \mid \text{WriteMem}(v_a, v_d, n) \\
&\quad \mid \text{AssumeReg}(r, v) \mid \text{DeclareConst}(x, \tau) \mid \text{DefineConst}(x, e) \mid \text{Assert}(e) \mid \text{Assume}(e) \\
t &::= [] \mid j :: t \mid \text{Cases}(t_1, \dots, t_n)
\end{aligned}$$

Fig. 14. Syntax of the Isla trace language (ITL) from Sammler et al. [32].

events ϵ used by ITrees) that describe how the ITL program manipulates the machine state (*i.e.*, registers and memory). Events can in turn use standard SMTLIB expressions e for manipulating bit vectors and booleans.

Islaris is an interesting case study for this paper since it uses a variety of effects that exercise the ability of our approach to handle non-standard programming languages. While the original work had to rely on various tricks to fit the Isla trace language into the fixed interface provided by Iris, we will see how our approach allows a direct encoding of Islaris using a combination of standard and non-standard events. Concretely, Islaris uses the following event type:

$$\mathbf{IslarisE} := \mathbf{DemonicE} \oplus \mathbf{StateE}_{\text{Islaris}} \oplus \mathbf{FailE} \oplus \mathbf{StepE} \oplus \mathbf{HaltE} \oplus \mathbf{SpecE}$$

We first have the standard events for demonic choice, state, and failure introduced in §2. We also use the **StepE** event (§6.1) to obtain partial correctness reasoning principles for recursive programs. Beyond this, Islaris uses two non-standard events that we discuss next: **HaltE** and **SpecE**.

HaltE. The **HaltE** event type provides the $\text{halt} : \mathbf{itree} \ \mathbf{HaltE} \ \emptyset$ operation that (safely) halts the execution and trivially finishes verification (*i.e.*, $\text{wpi} \ \text{halt} \ \{\Phi\} \dashv\vdash \text{True}$). This operation is necessary since ITL has an unusual way to read values from registers (and memory), akin to prophecy variables: First, it declares a variable that non-deterministically guesses the value that will be read from the register (using $\text{DeclareConst}(x, \tau)$ from Figure 14). Then, the read operation $\text{ReadReg}(r, x)$ uses halt to prune all executions where the value of the variable x does not correspond to the actual value of the register. (This unusual encoding of reads comes from the fact that the events j in ITL are SMT constraints that can only restrict existing variables but not assign them a new value.) Iris’s language interface does not provide a dedicated mechanism to support the halt operation, so Islaris uses a notion of a value that represents a halted program. With logical event handlers, we do not need to rely on such encodings since we can just directly encode halt as its own event.

SpecE. The second non-standard event **SpecE** comes from the fact that Islaris does not just prove the standard Iris adequacy that no program gets stuck (*i.e.*, no fail occurs), but also proves that the memory accesses to specially marked memory regions (representing MMIO regions) satisfy a user-defined safety property. To reason about such externally visible events, Islaris uses an encoding based on the observation mechanism that Jung et al. [22] introduced to reason about prophecy variables. Instead, our approach directly supports defining a **SpecE** event type with an operation $\text{emit}(\kappa)$ where κ represents a visible event (*i.e.*, read or write from resp. to MMIO memory) and a handler ensuring the safety properties are upheld.

No concurrency. Note that **IslarisE** does *not* use the **ConcE** event. This is on purpose since Islaris targets a sequential setting (verifying concurrent assembly programs against authoritative semantics is a research topic on its own). However, while Sammler et al. describe the program logic as sequential in the paper, in the actual Coq formalization it is concurrent (with a sequentially consistent semantics that does not match the concurrency of the actual assembly languages) since

```

1 Inductive trace_step :=
2 | AssertS e b ann es regs:
3   eval_exp e = Some (Val_Boo1 b) ->
4   trace_step (Smt (Assert e) ann :t: es) regs
5   (Some (LAssert b)) es
6 ...
7 Inductive seq_step :=
8 | SeqStep  $\sigma$   $\theta$   $\kappa$   $t'$   $\kappa'$   $\theta'$   $\sigma'$ :
9    $\theta$ .(seq_nb_state) = false ->
10  trace_step  $\theta$ .(seq_trace)  $\theta$ .(seq_regs)  $\kappa$   $t'$  ->
11  match  $\kappa$  with
12  | Some (LAssert b) =>  $\sigma' = \sigma \wedge \kappa' = \text{None} \wedge$ 
13     $\theta' = \theta <| \text{seq\_trace} := t' |> <| \text{seq\_nb\_state} := \text{negb } b |>$ 
14  ...

```

(a) Semantics of Assert from Sammler et al. [32]

(b) Semantics of Assert from this work

Fig. 15. Comparison of semantic definitions

the Iris language interface only supports concurrent languages. This means that the original work loses the reasoning principles for sequential languages (e.g., those pertaining to invariants). Our formalization of IsIris does not suffer from this drawback since we can easily model a sequential language by simply not using **ConcE**. You only pay for what you use.

Recreating the IsIris semantics and program logic. Using **IsIrisE**, it is straightforward to define a denotation from ITL traces into ITrees. In fact, the ITree-based definition is significantly simpler than the original definition since it can avoid the complex encodings described above. As a concrete example, Figure 15 shows how the semantics of $\text{Assert}(e)$ are encoded in the original version by Sammler et al. [32] and in the ITree-based version. The original version in Figure 15a splits the evaluation into two relations: First, trace_step checks that the expression e evaluates to the boolean b . Then, seq_step encodes the actual semantics of $\text{Assert}(e)$ by setting seq_nb_state to true if the assert fails (i.e., e evaluates to false). This seq_nb_state field encodes whether the execution should halt. (This is also seen on line 9 of Figure 15a, which ensures that the semantics can only step if seq_nb_state is false.) The ITree-based version in Figure 15b is a single recursive function that turns an ITL trace into an ITree. After checking that the expression e evaluates to the boolean b , it uses the `assume` function provided by the **HaltE** library to halt execution if b is equal to false. This shows how our ITree-based approach can replace the complex encoding based on seq_nb_state with a reusable library.

With the `compile_trace'` function at hand, it is straightforward to derive the IsIris program logic on top of it. We define a wp_{asm} for IsIris using `wpi` and reprove all program logic rules of the original IsIris. We also prove that our program logic satisfies the same adequacy statement as the original work. For this, we leverage the state machine adequacy described in §5, showing that this adequacy method can also scale to complex programming languages.

Even though we derive the same program logic with the same adequacy statement, the definition and proofs of the ITree-based version are significantly smaller than the original version. Concretely, the original definition of the trace_step and seq_step relations has around 130 LOC, while our ITree-based `compile_trace'` function has around 90 LOC. Also the proofs establishing the program logic become significantly smaller (around 270 LOC instead of around 450 LOC). What is even more important is that these proofs are qualitatively simpler: The original proofs require the use of a complex stepping lemma, followed by an inversion of the seq_step and trace_step relations from Figure 15a that requires careful manipulation of the proof state by hand. In contrast, the proofs for the ITree version consist of straightforward applications of the `wpi` lemmata for the ITree defining the semantics.

8 Related Work

We discuss related work along two axes: program logics that permit reuse across languages, and the state of modeling and reasoning about various computational effects with ITrees.

8.1 Program Logics with Reuse

The most prominent example of a “language-agnostic” program logic intended to be instantiated with a wide range of user-defined languages is Iris. The Iris technical manual [44, §8] describes their *language interface*: an arbitrary type of expressions and global state, and an associated per-thread small-step operational semantics—basically, a state transition system. Given an instance of this interface, Iris provides a weakest precondition connective and associated program logic rules. However, only the basic structural rules such as a bind lemma and a rule of consequence can be shared. Each language needs to re-define almost every aspect of its operational semantics, and then use “lifting lemmas” to provide corresponding reasoning principles in the program logic. Moreover, the language interface can be quite rigid: for instance, the only way a program may terminate is by returning a value, which means supporting an operation that halts the machine abruptly (but safely) from anywhere in the program requires non-trivial modeling. Iris also has to define two entirely separate weakest precondition connectives for total and partial correctness reasoning. There is a reusable library for defining the standard points-to connective, but its relation to the operational semantics needs to be re-proven for each new language.

In contrast, our approach allows more flexibility and more reuse: the same basic program logic can support both total and partial correctness reasoning, a **HaltE** effect for safe machine termination is easily supported, and the **HeapE** effect library can provide a ready-to-use points-to connective that is already integrated with the ITree semantics, to name a few. However, so far we have not implemented support for HeapLang’s prophecy variables [22] in our approach; that remains an interesting candidate for future work.

Abstract Separation Logic [5] defines a separation logic for a language that is denoted into a trace of “local actions” that each describe how the global state is altered. However, these actions do immediately act on the global state; there is no layering that would allow building up the global state from smaller, reusable pieces. Similarly, the Views framework [11] is defined for any language generated by a set of “atomic commands” that are given as (global) state transformers. In contrast, our approach supports composing **HeapE** with another instance of **StateE** that governs a separate piece of state (such as a file system or a network).

Dijkstra Monads [41, 18, 1, 26] provide a foundation for deriving weakest precondition connectives for arbitrary monadic computations. By applying a suitable sequence of monad transformers (which are in particular endofunctors on the category of monads) to a base effect observation (a morphism from a computation monad to a specification monad), one can build up the weakest precondition connective effect-by-effect. Their theory therefore encapsulates a different flavor of compositionally built program logics: instead of a single weakest precondition connective that generalizes to a wide range of effects, they systematically derive a new weakest precondition connective for each combination of such effects. However, since monads and concurrent computation are ill-fitted, they do not offer support for reasoning about concurrent programs. Furthermore, reasoning about state in this framework involves directly talking about the entire state; there is no separation logic support for local reasoning about memory.

Outcome Logic [50] provides a unified logic for reasoning about termination and non-termination (among other things). This gives a similar uniform treatment of partial and total correctness as our **StepH** handler described in §6.1. However, Outcome Logic focuses on unifying even more different reasoning style (*e.g.*, correctness and incorrectness reasoning), while we aim to provide a modular

approach for building program logics. While outcome logic can handle both demonic and angelic choice, it does not currently support using both of them in the same program, unlike our approach.

8.2 ITrees

The ITree line of work has so far mostly been focused on being able to formally capture the semantics of languages in a uniform framework and performing equational reasoning based on those semantics. As such, there has not been a lot of work on program logics for ITrees. The most notable exception is the work by Silver and Zdancewic [36] which applies Dijkstra monads to reason about ITrees. Using the recipe set out by Dijkstra monads, they derive a program logic for a specific ITree-based language (a simple imperative language called IMP) but do not discuss the idea of reusing program logic components and rules across languages.

The by far biggest application of ITrees is the VellVM project [49] which models a significant fraction of the LLVM IR specification using ITrees. They use a wide range of effects for that, most of which are also supported by our framework: several kinds of state, non-deterministic choice, fatal failure, and non-fatal machine termination. The one effect we have not implemented is external function calls; this is an interesting direction for future work. That would then allow us to build a program logic for the VellVM semantics of LLVM IR.

Choice Trees [7] extend ITrees with a *native* form of (demonic) non-deterministic choice. This is quite different from VellVM and our own approach where non-deterministic choice is just yet another effect. The payoff for special-casing choice is that the equational theory can be extended to properly support reasoning about choice. Our use of a program logic can be seen as an alternative approach for reasoning about ITrees with non-determinism that avoids special-casing. One case study for Choice Trees is a model of cooperative concurrency very similar to ours: thread forking is defined by duplicating the continuation. They give semantics to this model via a non-deterministic scheduler expressed directly in Choice Trees. We believe that our relational interpretation of the concurrency effects produce the same result. The new contribution of our framework is that we build a fully-featured concurrent separation logic for reasoning about these ITrees, enjoying all the concurrent reasoning principles that Iris provides. As part of our HeapLang case study, we also proved that this cooperative model of concurrency indeed soundly models all possible behaviors of a language defined with a small-step operational semantics and preemptive concurrency. Our approach is also able to support angelic choice as yet another effect, in contrast to Choice Trees that only support demonic choice.

Guarded interaction trees [13] (GITrees) provide a fully denotational model of a language with higher-order state into a variant of ITrees with support for higher-order events. This is different from our model of HeapLang (which also supports higher-order state): our heap stores *syntactic* HeapLang values, representing closures as expressions rather than their denotations. This makes our model much less suited for equational reasoning, but also much less technically demanding. Furthermore, event interpretation in GITrees hard-codes the state monad and therefore does not support other effects such as non-determinism and concurrency. It would be interesting to combine these lines of work and extend our program logic to support reasoning about GITrees.

Acknowledgments

We would like to thank the anonymous reviewers for their helpful feedback, which significantly improved the paper.

Data Availability Statement

The Coq development for this paper can be found in [46]. For the latest version of the paper and Coq development, see <https://plf.inf.ethz.ch/research/pop125-itree-program-logic.html>.

References

- [1] Danel Ahman, Catalin Hritcu, Kenji Maillard, Guido Martínez, Gordon D. Plotkin, Jonathan Protzenko, Aseem Rastogi, and Nikhil Swamy. 2017. Dijkstra monads for free. In *POPL*. ACM, 515–529. <https://doi.org/10.1145/3009837.3009878>
- [2] Andrew W. Appel. 2016. Modular Verification for Computer Security. In *CSF*. IEEE Computer Society, 1–8. <https://doi.org/10.1109/CSF.2016.8>
- [3] Alasdair Armstrong, Brian Campbell, Ben Simmer, Christopher Pulte, and Peter Sewell. 2021. Isla: Integrating Full-Scale ISA Semantics and Axiomatic Concurrency Models. In *CAV (LNCS, Vol. 12759)*. Springer, 303–316. https://doi.org/10.1007/978-3-030-81685-8_14
- [4] Rastislav Bodík, Satish Chandra, Joel Galenson, Doug Kimelman, Nicholas Tung, Shaon Barman, and Casey Rodarmor. 2010. Programming with angelic nondeterminism. In *POPL*. ACM, 339–352. <https://doi.org/10.1145/1706299.1706339>
- [5] Cristiano Calcagno, Peter W. O’Hearn, and Hongseok Yang. 2007. Local Action and Abstract Separation Logic. In *LICS*. IEEE Computer Society, 366–378. <https://doi.org/10.1109/LICS.2007.30>
- [6] Tej Chajed, Joseph Tassarotti, Mark Theng, Ralf Jung, M. Frans Kaashoek, and Nickolai Zeldovich. 2021. GoJournal: a verified, concurrent, crash-safe journaling system. In *15th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2021, July 14-16, 2021*, Angela Demke Brown and Jay R. Lorch (Eds.). USENIX Association, 423–439. <https://www.usenix.org/conference/osdi21/presentation/chajed>
- [7] Nicolas Chappe, Paul He, Ludovic Henrio, Yannick Zakowski, and Steve Zdancewic. 2023. Choice Trees: Representing Nondeterministic, Recursive, and Impure Programs in Coq. *PACMPL* 7, POPL (2023), 1770–1800. <https://doi.org/10.1145/3571254>
- [8] Arthur Charguéraud, Adam Chlipala, Andres Erbsen, and Samuel Gruetter. 2023. Omnisemantics: Smooth Handling of Nondeterminism. *ACM Trans. Program. Lang. Syst.* 45, 1 (2023), 5:1–5:43. <https://doi.org/10.1145/3579834>
- [9] Haogang Chen, Daniel Ziegler, Tej Chajed, Adam Chlipala, M. Frans Kaashoek, and Nickolai Zeldovich. 2016. Using Crash Hoare Logic for Certifying the FSCQ File System. In *USENIX Annual Technical Conference*. USENIX Association. https://www.usenix.org/conference/atc16/technical-sessions/presentation/chen_haogang
- [10] Santiago Cuellar, Nick Giannarakis, Jean-Marie Madiot, William Mansky, Lennart Beringer, Qinxiang Cao, and Andrew W. Appel. 2020. Compiler correctness for concurrency: from concurrent separation logic to shared-memory assembly language. <https://www.cs.princeton.edu/~appel/papers/ccc.pdf>
- [11] Thomas Dinsdale-Young, Lars Birkedal, Philippa Gardner, Matthew J. Parkinson, and Hongseok Yang. 2013. Views: compositional reasoning for concurrent programs. In *POPL*. ACM, 287–300. <https://doi.org/10.1145/2429069.2429104>
- [12] Robert W. Floyd. 1967. Nondeterministic Algorithms. *J. ACM* 14, 4 (1967), 636–644. <https://doi.org/10.1145/321420.321422>
- [13] Dan Frumin, Amin Timany, and Lars Birkedal. 2024. Modular Denotational Semantics for Effects with Guarded Interaction Trees. *PACMPL* 8, POPL (2024), 332–361. <https://doi.org/10.1145/3632854>
- [14] Ming Fu, Yong Li, Xinyu Feng, Zhong Shao, and Yu Zhang. 2010. Reasoning about Optimistic Concurrency Using a Program Logic for History. In *CONCUR (Lecture Notes in Computer Science, Vol. 6269)*. Springer, 388–402. https://doi.org/10.1007/978-3-642-15375-4_27
- [15] Jason Gross. 2024. Answer to ‘Is CoInductive “extensionality” sound in Coq? Is it generalizable?’. <https://stackoverflow.com/a/69905520>. Accessed: 2024-10-24.
- [16] Armaël Guéneau, Johannes Hostert, Simon Spies, Michael Sammler, Lars Birkedal, and Derek Dreyer. 2023. Melocoton: A Program Logic for Verified Interoperability Between OCaml and C. *PACMPL* 7, OOPSLA2 (2023), 716–744. <https://doi.org/10.1145/3622823>
- [17] Aquinas Hobor, Andrew W. Appel, and Francesco Zappa Nardelli. 2008. Oracle Semantics for Concurrent Separation Logic. In *ESOP (LNCS, Vol. 4960)*. 353–367. https://doi.org/10.1007/978-3-540-78739-6_27
- [18] Bart Jacobs. 2015. Dijkstra and Hoare monads in monadic computation. *Theor. Comput. Sci.* 604 (2015), 30–45. <https://doi.org/10.1016/j.tcs.2015.03.020>
- [19] Jules Jacobs, Jonas Kastberg Hinrichsen, and Robbert Krebbers. 2024. Deadlock-Free Separation Logic: Linearity Yields Progress for Dependent Higher-Order Message Passing. *PACMPL* 8, POPL (2024), 1385–1417. <https://doi.org/10.1145/3632889>
- [20] Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2018. RustBelt: Securing the foundations of the Rust programming language. *PACMPL* 2, POPL (2018), 66:1–66:34. <https://doi.org/10.1145/3158154>
- [21] 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>
- [22] Ralf Jung, Rodolphe Lepigre, Gaurav Parthasarathy, Marianna Rapoport, Amin Timany, Derek Dreyer, and Bart Jacobs. 2020. The future is ours: Prophecy variables in separation logic. *PACMPL* 4, POPL (2020), 45:1–45:32. <https://doi.org/10.1145/3371113>

- [23] Robbert Krebbers, Jacques-Henri Jourdan, Ralf Jung, Joseph Tassarotti, Jan-Oliver Kaiser, Amin Timany, Arthur Charguéraud, and Derek Dreyer. 2018. MoSeL: A general, extensible modal framework for interactive proofs in separation logic. *PACMPL* 2, ICFP (2018), 77:1–77:30. <https://doi.org/10.1145/3236772>
- [24] Robbert Krebbers, Ralf Jung, Ales Bizjak, Jacques-Henri Jourdan, Derek Dreyer, and Lars Birkedal. 2017. The Essence of Higher-Order Concurrent Separation Logic. In *ESOP (LNCS, Vol. 10201)*. 696–723. https://doi.org/10.1007/978-3-662-54434-1_26
- [25] Robbert Krebbers, Amin Timany, and Lars Birkedal. 2017. Interactive proofs in higher-order concurrent separation logic. In *POPL*. 205–217. <https://doi.org/10.1145/3009837.3009855>
- [26] Kenji Maillard, Danel Ahman, Robert Atkey, Guido Martínez, Catalin Hritcu, Exequiel Rivas, and Éric Tanter. 2019. Dijkstra monads for all. *PACMPL* 3, ICFP (2019), 104:1–104:29. <https://doi.org/10.1145/3341708>
- [27] William Mansky, Andrew W. Appel, and Aleksey Nogin. 2017. A verified messaging system. *PACMPL* 1, OOPSLA (2017), 87:1–87:28. <https://doi.org/10.1145/3133911>
- [28] William Mansky and Ke Du. 2024. An Iris Instance for Verifying CompCert C Programs. *PACMPL* 8, POPL (2024), 148–174. <https://doi.org/10.1145/3632848>
- [29] Roland Meyer, Thomas Wies, and Sebastian Wolff. 2022. A concurrent program logic with a future and history. *PACMPL* 6, OOPSLA2 (2022), 1378–1407. <https://doi.org/10.1145/3563337>
- [30] Peter W. O’Hearn. 2007. Resources, concurrency, and local reasoning. *Theor. Comput. Sci.* 375, 1-3 (2007), 271–307. <https://doi.org/10.1016/j.TCS.2006.12.035>
- [31] Ingrid Rewitzky. 2003. Binary Multirelations. In *Theory and Applications of Relational Structures as Knowledge Instruments*. LNCS, Vol. 2929. Springer, 256–271. https://doi.org/10.1007/978-3-540-24615-2_12
- [32] Michael Sammler, Angus Hammond, Rodolphe Lepigre, Brian Campbell, Jean Pichon-Pharabod, Derek Dreyer, Deepak Garg, and Peter Sewell. 2022. Islaris: verification of machine code against authoritative ISA semantics. In *PLDI*. ACM, 825–840. <https://doi.org/10.1145/3519939.3523434>
- [33] Michael Sammler, Simon Spies, Robbert Krebbers, Kayvan Memarian, Derek Dreyer, and Deepak Garg. 2021. RefinedC: Automating the Foundational Verification of C Code with Refined Ownership Types. In *PLDI*. 158–174. <https://doi.org/10.1145/3453483.3454036>
- [34] Michael Sammler, Simon Spies, Youngju Song, Emanuele D’Osualdo, Robbert Krebbers, Deepak Garg, and Derek Dreyer. 2023. DimSum: A Decentralized Approach to Multi-language Semantics and Verification. *PACMPL* 7, POPL (2023), 775–805. <https://doi.org/10.1145/3571220>
- [35] Upamanyu Sharma, Ralf Jung, Joseph Tassarotti, M. Frans Kaashoek, and Nikolai Zeldovich. 2023. Grove: a Separation-Logic Library for Verifying Distributed Systems. In *SOSP*. ACM, 113–129. <https://doi.org/10.1145/3600006.3613172>
- [36] Lucas Silver and Steve Zdancewic. 2021. Dijkstra monads forever: termination-sensitive specifications for interaction trees. *PACMPL* 5, POPL (2021), 1–28. <https://doi.org/10.1145/3434307>
- [37] Youngju Song, Minki Cho, Dongjae Lee, Chung-Kil Hur, Michael Sammler, and Derek Dreyer. 2023. Conditional Contextual Refinement. *PACMPL* 7, POPL (2023), 1121–1151. <https://doi.org/10.1145/3571232>
- [38] Simon Spies, Lennard Gäher, Joseph Tassarotti, Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. 2022. Later credits: resourceful reasoning for the later modality. *PACMPL* 6, ICFP (2022), 283–311. <https://doi.org/10.1145/3547631>
- [39] Kasper Svendsen and Lars Birkedal. 2014. Impredicative Concurrent Abstract Predicates. In *ESOP (LNCS, Vol. 8410)*. 149–168. https://doi.org/10.1007/978-3-642-54833-8_9
- [40] Kasper Svendsen, Lars Birkedal, and Matthew J. Parkinson. 2013. Joins: A Case Study in Modular Specification of a Concurrent Reentrant Higher-Order Library. In *ECOOP (Lecture Notes in Computer Science, Vol. 7920)*. Springer, 327–351. https://doi.org/10.1007/978-3-642-39038-8_14
- [41] Nikhil Swamy, Joel Weinberger, Cole Schlesinger, Juan Chen, and Benjamin Livshits. 2013. Verifying higher-order programs with the dijkstra monad. In *PLDI*. ACM, 387–398. <https://doi.org/10.1145/2491956.2491978>
- [42] The Coq Team. 2024. The Coq proof assistant. <https://coq.inria.fr/>.
- [43] The Coq Team. 2024. The Logic of Coq. <https://github.com/coq/coq/wiki/The-Logic-of-Coq>. Accessed: 2024-10-18.
- [44] The Iris Team. 2024. The Iris 4.2 Reference. <https://plv.mpi-sws.org/iris/appendix-4.2.pdf>
- [45] Aaron Turon, Viktor Vafeiadis, and Derek Dreyer. 2014. GPS: navigating weak memory with ghosts, protocols, and separation. In *OOPSLA*. ACM, 691–707. <https://doi.org/10.1145/2660193.2660243>
- [46] Max Vistrup, Michael Sammler, and Ralf Jung. 2024. *Artifact of "Program logics à la carte"*. <https://doi.org/10.5281/zenodo.14180355> Development version: <https://gitlab.mpi-sws.org/iris/itree-program-logic>.
- [47] Vladimir Voevodsky. 2024. Forum discussion on ‘coinductives’. <https://groups.google.com/g/homotopytypetheory/c/tYRTcL2OpYo/m/PIrl6t5me-oj>. Accessed: 2024-10-24.
- [48] Li-yao Xia, Yannick Zakowski, Paul He, Chung-Kil Hur, Gregory Malecha, Benjamin C. Pierce, and Steve Zdancewic. 2020. Interaction trees: representing recursive and impure programs in Coq. *PACMPL* 4, POPL (2020), 51:1–51:32. <https://doi.org/10.1145/3371119>

- [49] Yannick Zakowski, Calvin Beck, Irene Yoon, Ilia Zaichuk, Vadim Zaliva, and Steve Zdancewic. 2021. Modular, compositional, and executable formal semantics for LLVM IR. *PACMPL* 5, ICFP (2021), 1–30. <https://doi.org/10.1145/3473572>
- [50] Noam Zilberstein, Derek Dreyer, and Alexandra Silva. 2023. Outcome Logic: A Unifying Foundation for Correctness and Incorrectness Reasoning. *Proc. ACM Program. Lang.* 7, OOPSLA1 (2023), 522–550. <https://doi.org/10.1145/3586045>

Received 2024-07-11; accepted 2024-11-07